The Ghost of Socrates: Exploring Philosophical Issues in Information Systems

By John M. Artz, PhD
Associate Professor of
Information Systems
The George Washington University

Table of Contents

Introduction

Philosophy is, to a large extent, about the meaning of words. What do you mean when you say words such as truth, reality, justice, or fairness? Philosophers attempt to find the essential or deeper meanings of these and other similar words that refer to important concepts that we use to guide us in making important decisions that affect our lives. And the goal of philosophy, again to a large extent, is to organize these meanings into coherent frameworks that help us make sense out of the world around us.

"But, if philosophy is just about the meanings of words," one might ask, "then what is the difference between philosophy and a dictionary?" And, with that question, one has already embarked, perhaps unknowingly, on a philosophical quest.

We have the word 'philosophy' which refers to a concept. This concept, in turn, refers to a collection of meanings. To some people, philosophy might be the product of an ongoing process of questioning the world around them. To others, philosophy might be a corpus of knowledge that is the source of intellectual enlightenment. Another might see philosophy as a collection of incomprehensible jibber jabber used to confuse people who believe that they already know what they need to know and resent the confusion that philosophy creates. And, yet others may see philosophy simply as that dense and boring stuff that they were required to read in a sophomore class, most of which they have mercifully forgotten. Depending upon whom you talk to, philosophy can be any of these things and, perhaps, many more. But we can summarize by saying that philosophy is some sort of an intellectual quest, including the results of that quest and perhaps the techniques employed in that quest. And understanding the meaning of words is perhaps one of the products of that quest.

A dictionary, on the other hand, is a book, far too infrequently referenced, that provides guidelines for spelling, pronunciation and the meanings of words. So philosophy is some sort of an ongoing intellectual investigation and a dictionary is a reference book. They are actually different in every way except the fact that they both have some concern about the

meaning of words. So now we can focus on this concern about meaning and see how it may be different.

The meanings addressed by the dictionary reflect common usage and, in some better dictionaries, the derivations of the words also. This shows where the word came from and how it is commonly used. The dictionary will help you avoid using a word incorrectly in a sentence but it will not reflect the fullness of the meaning of the concept referenced by the word. Philosophy, on the other hand, will examine the deeper meanings with the goal of providing insights or organizing knowledge to advance understanding. Some philosophies examine meanings in order to improve the world. Philosophy is different from a dictionary in that a dictionary does not try to integrate the various meanings. Nor does it provide contextual meanings or deeper meanings.

It may seem a little silly to be exploring the difference between philosophy and a dictionary, but the exercise has already yielded two benefits. First, we see that words have deeper meanings beyond what the dictionary provides. These deeper meanings are not universally agreed upon. They may evolve over time. And it is possible to have meanings within a single concept that conflict with each other or, perhaps even, contradict each other. One of the tasks of the philosopher is to attempt to integrate conflicting meanings and to integrate new meanings that evolve over time with previous understandings of a particular concept.

Second, and more importantly for the current discussion, this comparison of philosophy and dictionaries has provided us with an example of an important technique called concept analysis, which is the basis of most of the discussions in this book. If you could follow, albeit not necessarily appreciate, the preceding discussion, you should be able to follow most of the analysis and discussion provided in the following chapters.

When you ask a question such as, “Is privacy a right?” You are embarking upon a philosophical quest. In order to answer this question, you must understand the meaning of

the concepts privacy and right. Once you understand these concepts, you have to decide if the concept of privacy is included in the collection of concepts that we call rights.

Philosophy, taken as a whole, can be an overwhelming and often confusing body of wisdom, knowledge, insights, comments and the like. There are even some false starts and contradictions to be contended with. Just remembering the names of all the important philosophers would be a daunting task. Remembering the names and a three line synopsis of their contributions would be even more daunting without even requiring that one understand the three-line synopsis and how it compares with the three line synopsis of other philosophers.

In order to make this enormous collection of human intellectual achievements more accessible to the average person, there are several simplifying or organizing approaches that one may use. First, one can study philosophy from a historical or chronological perspective. In this approach you examine the pre-Socratic philosophers, then Socrates, Plato, Aristotle and so on until you reach the modern day philosophers. This approach makes sense because most major philosophers were aware of what those who went before them had to say and this approach reflects the evolution of ideas. There are several problems, however, with this approach. First, it assumes that philosophy just gets better every day and if this is the case why not just read modern philosophers since their insights must incorporate the wisdom and insights of those who went before them. A second problem is that if a person were confronted with a philosophical problem that they would like to resolve, they would probably not want to study the thoughts of a couple hundred long dead philosophers - most of whom had nothing to say on the topic about which they are concerned – before making any progress on the problem at hand.

Another approach is to study specific branches of philosophy such as metaphysics, moral philosophy, aesthetics, social philosophy or political philosophy. A person wondering if Andy Warhol's Pop Art was really art, might do well to study some aesthetics while social and political philosophy might be of little use in answering this question. Studying specific branches of philosophy might also employ a chronological approach, but it

would focus on the development of specific ideas rather than the full range of observations of each philosopher. This approach also has some shortcomings.

To paraphrase a widely cited observation on intellectual overkill, if somebody wants to know what time it is, don't tell them how to make a watch. Sometimes developing all this intellectual machinery is simply not worth it for somebody who wants to figure out an answer to a philosophical question that has immediate relevance to him or her and must be satisfying primarily to him or her.

Somebody who embarked on either of the above approaches might be able to say – Plato said this, Aristotle said that, and Locke or Hume said another thing. But if you ask – What do you think? – you might get a puzzled look.

This is where concept analysis comes in. If one wishes to know if Andy Warhol's Pop Art is really art, they may begin by first asking – "What is art?" and then by asking if Andy Warhol's Pop Art fits within that concept. It is unlikely that this query will result in a yes or no answer. It is more likely that it will result in some understanding of art and some understanding of how close to the center of the concept the art the work in question resides. It will also result in some understanding of the attributes of Pop Art that keep it from being central to the concept if that were indeed the case.

This is not to suggest that concept analysis is the only way or the best way to study philosophy. Hopefully, the practice of concept analysis would, eventually, lead one to embark on one of the other approaches. For example, in order to answer the question – What is art? – one might consult the literature on aesthetics. Or one might ask – What is art, according to Aristotle? So concept analysis is not exclusive of the other approaches, and in fact it might lead to the other approaches.

The benefit of concept analysis is that it gets the curious person on the way to philosophical inquiry much more quickly. If a person asks – What is art? And you tell them to spend four years studying the works of the great philosophers before they can

address the question; they are likely to be deterred from their quest. However, if you ask them for a few examples of what they think is art (as Socrates would have), then perhaps consult the dictionary to see if the definition provides any insight, and finally move on to some formal concept analysis, you would likely begin making progress right away.

This approach is even more relevant today within the field of information systems, which is the focus of this book. Technology and the humanities are viewed as different and frequently mutually exclusive paths in life. So the technologists who are facing philosophical issues probably have no background in philosophy and hence no way to address these issues. They may not even realize that they are facing philosophical problems. People with a background in the humanities might very well be in a position to proceed on some of these philosophical questions, but feel barred from the inquiry due to their lack of understanding of technology. And, of course, a person who has some background in philosophy and a rich understanding of technology might not see the connections between the two. How we got to this state of affairs bears some further examination.

Several versions of the modern computer were invented in the early to mid 1940's and the name 'computer' reveals the common understanding of that time regarding what this device would be good for. Computers were electronic computational devices and were useful for complex mathematical computations such as trajectory tables or the orbits of heavenly bodies. However, if that were the ultimate fate of these computational machines, it is likely that few people today would care much about them.

But, innovative computer engineers saw that these computational devices had possibilities beyond complicated arithmetic. They saw that these devices could be used also for information processing. At the time, information processing was done using punched cards, a technology that dated back to the turn of the century. Punched cards were a physical medium being processed by a mechanical device. Information could be processed only as fast as cards could be sent through a machine and breakdowns were common either from mechanical failure or from cards getting jammed in the machine. By

recording and processing information electronically it would be possible to avoid card jams and mechanical failures. Further, the information could be processed at the speed of light.

There is an observation, frequently attributed to Thomas Watson who was the president of IBM at the time that the computer was invented. When asked how many computers he thought he could sell by end of the century, he replied 'five'. This apocryphal story circulates in the lore of computer technology as one of those statements that the speaker would like to take back. And whether or not Watson ever really said this, the story is at least revealing. The key to understanding why this statement probably made sense at the time can be found in the word 'computer'. A computer is a computational device. When Watson was predicting how many computers he could sell, he was probably speculating on how many customers there would be for electronic devices that could perform complex mathematical calculations. And, at the time, there probably were not that many customers who needed that level of computational power.

However, clever engineers were capable of seeing other ways to use this device for information processing. The modern computer does very little computation. Most computer power is applied to tasks such as moving, searching, validating, retrieving or storing information. In fact, if we wanted to be more accurate in our nomenclature we would call computers processors rather than computers. But the initial title stuck and we still call them computers. However, it is interesting to note that computer chips have been called processors for many years now. Again we see how looking a little more closely into the meanings of words yields deeper insights into issues that appear to be purely technological.

Over the next half a century computers became increasingly more powerful and the impact upon the life of the average person increased. This impact became so great that many philosophical questions arose regarding the proper use of this technology. Issues such as privacy, social impacts, impact on workers, and universal access to computer networks arose and needed to be addressed. Over that same period of time new academic

disciplines such as computer science and information systems also arose attempting to keep up with and sometimes advance this rapidly evolving information technology.

Unfortunately, these new academic disciplines arose with an arcane vocabulary and an esoteric set of concepts that separated them from the average person whose life was being affected by these new technologies. This separation created an exclusive technological priesthood, which in turn led to two significant problems. First, people who are not members of the technological priesthood, whether they be average people or learned scholars, were excluded from addressing these newly arising philosophical issues by their unfamiliarity with the arcane vocabulary and esoteric concepts. Second, people within the priesthood tended to stay within the priesthood attempting to solve all problems in terms of the technological concepts within the field and rarely looking outside of the field for guidance. The belief was that these are computer problems and they should be solved within the computer field. Anything that anybody had to say before the computer was invented couldn't possibly be useful in solving computer problems.

Perhaps, at first, computer problems were new problems requiring new solutions. But eventually, information systems led us to new incarnations of old problems. Maybe the question of *how* to achieve a particular result with regard to information technology is indeed a technical problem. However, the question of *what* to achieve is a problem as old as philosophy itself. And most of the questions that we ask today regarding what should be done are questions that have been asked many times before.

I call this book The Ghost of Socrates because I believe that if it were possible to conjure up the Ghost of Socrates, he might be able to provide us with an enormous amount of insight into the central problems in information systems even though he lived and died twenty-four centuries before the computer was invented. Socrates could probably help us in two ways. First, his method of inquiry, generally known as the Socratic Method, would help us clarify what we do know and help us formulate the questions that we need to ask. The Socratic Method relies heavily on concept analysis thus making Socrates a major contributor to the current philosophical investigation. Second, the philosophical positions

of those that followed him, his student Plato and Plato's student Aristotle, provide philosophical frameworks that can help us frame some of the problems. While we are conjuring up ghosts, I might also conjure up David Hume and John Locke, and possibly Bertrand Russell and Ludwig Wittgenstein. If this little exercise were possible, this esteemed group might be on the leading edge of the problems in information systems in a relatively short amount of time, and would probably be able to help us formulate our philosophical problems in a way that may lead to some solutions. Although I may not always mention these guys by name, their thinking influenced my thinking in some very productive ways.

The central problems in information systems are probably quite different from those that the average person might suspect. For example, how to process a record faster is of far less importance that what does the data on that record mean, how do we know it is correct, who should be allowed to see that information, what should that information be used for, and how can we use that data to make better decisions. In information security, for example, we know how to make information secure. Just lock it up and don't let anybody have access to it. But the value of information is derived from the uses to which that information is put, so in order to have value information must be accessible. So the question becomes who should have access to information, which is no longer a technical problem.

There were many decisions I had to make in developing the framework for this book. Who is the target audience? How much effort might they be willing to make to grasp some of these concepts? Do I want to advance the understanding of a few leading edge researchers or do I want to advance the understanding of a wide audience. And finally, do I want to raise questions that need to be addressed or do I want to provide answers that I am comfortable with on some of these issues.

As an aside, it is interesting to note that the word decision comes from the same root as the word incision. Incision means to cut into something, while decision means to cut something off. What are we cutting off when we make a decision? We are cutting off

alternatives. And that is the agony of having to decide. As I choose a target audience, I am quite likely cutting off other target audiences. As I choose a focus for the book, I am cutting off other possible focuses. I would like to write a book that everyone would like to read, and I would like to raise important issues and provide satisfying answers. However, one book cannot do everything. So here is what I decided.

The average practitioner in information systems has little background in philosophy. To that extent, the average practitioner is on a par with the average user. My target audience is the average person who feels that computers and information systems play an important role in their lives. So if you are a person who develops information systems or a person who relies heavily on information systems in your job or personal life, this book should be of interest to you. If you are more advanced, say a researcher in information systems, I believe I will provide you with a new look at some old problems and some perspectives that you may not have previously considered. If you are a philosopher, I am probably not going to provide any new philosophical insights, but I might provide some novel applications of existing ideas in the area of information systems.

I cannot provide any insight to someone who does not read the book. Therefore readability and accessibility become of paramount importance. Some of the ideas in this book are quite profound. This is not because they are my ideas. I actually borrowed the philosophical ideas from other places. But to compensate or perhaps to cover for this philosophical depth, I have chosen to write in a somewhat informal and breezy style to make the reading more enjoyable even if the reader does not grasp the full significance of each point. Hopefully, the reading will be easy and interesting enough that the reader will reread a chapter or two after pondering the contents.

One of the themes in the Platonic dialogs is that Socrates does not set up a hurdle that all his interlocutors must get over. Instead, he meets each person at their own level and attempts to raise their understanding a notch or two. I agree with that objective. People of all different backgrounds and perspectives may read this book. If each comes away with a better understanding of these issues than they had when they began the book, then I

would consider the book a success and would feel that my efforts in writing it have been worthwhile. Following is an outline of the book and the topics that will be addressed.

Introduction

The Ghost of Socrates

This chapter will provide some background on Socrates, his student Plato and Plato's student Aristotle. It will show why consulting these philosophers is highly relevant to philosophical inquires in information systems.

Knowledge Management

In recent years we have seen a lot of resources invested in knowledge management systems. By asking the question what is knowledge and is it knowledge that is being managed by these systems we can get some insight into the weaknesses and misconceptions in knowledge management.

Concept analysis

This chapter will introduce concept analysis as a technique for philosophical inquiry. It will provide specific suggestions so that one may be employ immediately to begin investigating general ideas or specific questions.

Protecting Privacy in Cyberspace

Most people today feel that the right to privacy must be protected at all costs. However, a closer inspection of the issues here reveals that privacy is a tradeoff and when one examines those trade offs it is less clear that privacy is paramount.

Computers and the Quality of Life

Most people would agree that modern information technology should somehow improve our quality of life. Following this is generally the claim that everybody should have access to technology and specifically there should be universal access to the World Wide Web. But these are just surrogate solutions. That is people believe that more technology means a better life. A closer examination of quality of life reveals that the issue is more complex and more technology may or may not mean an improved quality of life.

Machine Intelligence

Is artificial intelligence a real possibility? This debate has gone back and forth for several decades now. This chapter will revisit some of the classic arguments from the perspective of concept analysis and explore such questions as what do we mean by intelligence and does the things that machine do fit within that concept.

Redistribution of Social Power

Users of information technology are sometime, maybe even frequently, resistant to the introduction of new technology or new information systems. It is generally believed that there is some sort of resistance to change among users of information technology. But perhaps it is not change that people are resistant to. Rather it may be redistributions of social power and the uncertainty that comes with those redistributions. This chapter looks at ways in which information technology may affect the distribution of power in society and organizations.

Science vs. Technology

Science can be used to advance technology and technology can be used to advance science. So it appears that these two concepts are hopeless intertwined when, in fact, they

are quite different concepts. This chapter will pick apart the differences between science and technology and show how these two synergistic concepts could not be more different.

Teleology in Information Systems Development

The first step in the development process for information systems is a step called Requirements Analysis. In this step the analysts determine the things that the information system must do in the form of requirements. But what are requirements and where to they come from? This chapter will examine a very basic metaphysical flaw in information systems development and suggest a way to shore up the problem.

Thinking Clearly About Concepts

This final chapter will revisit concept analysis with the benefit of the proceeding chapters and future elaborate on the idea of clarifying what we mean. It will also discuss how to avoid being drawn in to poorly formed concepts.

Annotated Bibliography

Very few references are made in the course of the book to avoid unnecessary distractions to the reader. However, in order to fairly give credit to the source of important ideas in this book an annotated bibliography will be provided. The books cited in the bibliography can also provide further reading for those readers who wish to investigate these topics and ideas in greater depth.

Chapter 1: The Ghost of Socrates

It is hard to imagine how anyone who lived twenty-four centuries ago could possibly have anything of value to say about problems that arose in the late twentieth century as a result of computers and information technology. Certainly Socrates could never have dreamed of computers, networks, databases or artificial intelligence. And in his wildest imaginings he certainly never would have been able to envision the World Wide Web. But the problems we are facing are not new. Neither is the means of solving them.

There are two misconceptions that must be overcome in order to benefit from the wisdom of Socrates. This first is that problems created by technology must be solved by technology. Many people in the computer field act as though the world began in the early 1940's with the invention of the modern computer and any wisdom from before that date could not possibly be relevant to the problems we are facing in the information age.

Certainly, information technology has brought certain problems such as personal privacy to the foreground. But personal privacy is just one of the more recent instances of fairness when considering the rights of the individual versus the rights of society, or the rights of one individual versus the rights of another individual. And the issue of fairness when allocating rights is certainly not new.

The second misconception is that Socrates was involved in exploring lofty philosophical ideals and consequently is too far removed from the practical problems of today to have

any insight into them. Of course, whenever we think of a great philosopher, like Socrates, whose name and teachings have lasted for over twenty-four centuries, we often do so with a natural reverence that puts as much distance between us and his ideas as there is between the ideal we have created and the man himself.

While the teachings of Socrates are worthy of great respect, this respect should never be so great that we promote them to a distant world of ideals that has little to do with our daily lives. Socrates was a person just like you and I are people. And Socrates was concerned with complex issues of his day, just like you and I are concerned with complex issues of our day. What Socrates offered was a way of examining these issues that is as useful and relevant today as it was twenty-four centuries ago. And if we could, somehow, conjure up the Ghost of Socrates, he might be able to give us a hand sorting out the complex issues that we are facing.

This chapter will make the figure of Socrates more familiar by showing how the world in which Socrates lived was not as different from our world as one might think. It will also, briefly, introduce the Socratic method and show how this can be applied to problems arising from information technology. Finally, Plato and Aristotle will also be introduced in order to provide some key philosophical concepts that will be used throughout the rest of the book. It is important to remember that Socrates, Plato and Aristotle should not be deified. They were smart guys with some great insights. And those insights can be extremely useful as we sort out the complex issues arising from our increasing reliance on information systems technology. But they were people just like we are people and

turning to them for advice in resolving our confusion is probably a very prudent idea.

Turmoil in Ancient Greece

During the 5th century BCE, ancient Greece experienced tremendous intellectual advances and made lasting contributions to western civilization in the areas of art, government, mathematics and, of course, philosophy. This Golden Age of Greece (sometimes called the Golden Age of Pericles) lasted forty-six years, from 477 BCE until 431 BCE. The Golden Age ended when war broke out between Athens and Sparta. This war, known as the Peloponnesian war, lasted until 404 BC and left Athens in defeat.

After four and a half decades of intellectual advances and social enlightenment, Athens endured twenty-seven years of war, only to be defeated by Sparta, one of its least enlightened neighbors. During the Golden Age, new ideas such as democracy were introduced. And the defeat by the military state of Sparta raised questions about whether or not democracy was such a good idea. In fact, the entire period was characterized by many questions on central issues of philosophy and politics. Some questions challenged traditional ideas and values, while other questions challenged new ideas and new values. Socrates lived from about 469 BCE until 399 BCE. So the entire span of Socrates' life was characterized by great intellectual advances accompanied by great social and ideological turmoil; not altogether unlike the 20th century. If you sometimes feel like you want the world to just stop everything while you take some time to figure things out, then you probably know how Socrates felt back in the 5th century BCE.

Contributing to the volatile intellectual climate during the Golden Age was a group of itinerant teachers and public speakers known as the Sophists. The Sophists were clever, but cynical, rhetoricians who believed that they could effectively argue any side of an argument and that personal virtue was demonstrated by material success. One of the most famous of these Sophists was Protagoras who make the well-known statement that "Man is the measure of all things". The sophists did not believe in enduring truths. They claimed that truth was whatever people believed truth to be, not unlike the social relativists of the 20th century who claim that truth is what a society decides that it is. Certainly the Sophists of ancient Greece and the social relativists of the late twentieth century cannot easily be lumped together philosophically. However, they both demonstrate that in a time of social and intellectual turmoil and change, some people are led to the conclusion that there is no such thing as enduring truth.

All this is to say that the world of Socrates in ancient Greece was not as different from the 20th century as one might first believe. And the method that Socrates brought to bear on his problems might well be brought to bear on the problems of today.

Socratic Method

Socrates' main contribution to philosophy was not a body of philosophical doctrine. Rather it was a method of exploring philosophical concepts. As we have noted, Socrates lived during a period of great social turmoil in ancient Greece when traditional values

(particularly personal and political) were changing. In this period of changing values, Socrates contributed to the public debate by asking questions of the form - "What is _____?" where the blank could be filled in with terms like justice, virtue, beauty, or any other concept that was worthy of scrutiny.

These questions would be asked in a group setting and the other participants in the discussion, known as interlocutors, would respond with their ideas on the concept in question. Socrates would, in turn, respond to their contributions by asking further questions, which were intended to get closer to the heart of the issue. This process of question and answer exploration was know as Socratic dialectic and is the basis of the Socratic Method. The Socratic Method is used widely today as a teaching method particularly in fields such as law and business in which subtleties are teased out of cases by an instructor who leads the students through the key principles by asking a continuing series of questions. However, there is more to Socratic Method than just asking questions. The essence of Socratic Method is asking the right questions to get to the heart of a concept under examination. Thus, Socratic Method is the basic for concept analysis which will be examined in more detail in Chapter 3.

It bears emphasizing that Socrates did not preach or promote any philosophical doctrine. Instead, he asked questions to illuminate the issues under discussion. Granted, that the questions asked often led to certain conclusions or interpretations, but more often than not they just led to the conclusion that Socrates and the interlocutors did not understand the concept well enough to have any opinions and the most they could conclude was that

they were ultimately ignorant. This philosophical neutrality makes the Socratic Method valuable in examining any complex issue because rather than lead one to a specific conclusion it leads one to a deeper understanding of the issues and a realization of how little one really knows about the question at hand.

It is also worthy of mentioning that Socrates frequently viewed his interlocutors in terms of their personal development using a metallic metaphor of bronze, silver and gold. A bronze person might be selfish and concerned only with his own needs. A silver person might be socially responsible, but ultimately motivated by fame or glory. A gold person would be an enlightened individual who pursued truth for the sake of truth and did the right things because they *were* the right things. In his discussions Socrates would try to raise the bronze person to a silver level, a silver person to a gold level, and a gold person to yet greater self-awareness. Thus, the goal of his dialectic was the individual improvement of the interlocutors rather than the successful promotion of a particular philosophical doctrine.

Since this is a book on philosophical issues in information systems and not a book on philosophy, once we leave this chapter the focus will be more on philosophical explorations rather than philosophy itself. Hence, before leaving this chapter, it is necessary to introduce a couple more philosophers and a few more important concepts.

Socrates, Plato and Aristotle

When we look and the depth and breadth of ancient Greek philosophy it is easy to assume that it must have been developed and refined by hundreds of scholars over many centuries. In fact, the majority of ancient Greek philosophy can be attributed to three philosophers over the course of about a hundred and fifty years. These three philosophers were Socrates, his student Plato, and Plato's student Aristotle. Each made a unique contribution to western philosophy. Socrates contributed a philosophical method of inquiry. Plato contributed a view of philosophy based on ideals. And Aristotle contributed a view of philosophy based on rigorous observation and physical evidence. The tensions and synergies between Plato's view of the world and Aristotle's view of the world can still be seen today in the tensions and synergies between theoretical and empirical sciences.

Socrates was born around 469 BCE and died (or more accurately was put to death) in 399 BCE, which means that his life encompassed the entire period of the Golden Age followed by the humiliating defeat of the Athenians by the Spartans. He was a fat ugly little man who hung around in the marketplace posing difficult questions and, in the view of many civic minded Athenians, creating a civil disturbance. Socrates was brought before the Athenian senate on charges of impiety and corrupting the youth of Athens and was sentenced to death by drinking hemlock for his punishment.

We do not know much about Socrates' views because he did not write anything down. In fact, Socrates was opposed to writing because he thought that writing would weaken mental skills such as memory and verbal agility. This is particularly poignant in our age of technology when many people feel that reliance on technology is a bad thing because other important skills, like being able to add in your head, will deteriorate. Calculators and spell checkers are considered by some to be evils of modern technology because they replace important mental skills. But writing turned out to be O.K. and maybe these other technologies will be O.K. too.

Since Socrates did not write anything down, we have to rely on Plato for the content of Socrates' philosophy. Plato wrote a series of thirty-five dialogs in which Socrates sparred with a variety of opponents on a variety of topics. The most famous of these dialogs is The Republic, in which Plato presents his views on the perfect form of government. However, having to rely on the dialogs of Plato for our understanding of Socrates creates a small problem since we do not know if the views presented in Plato's writings are the views of Socrates as recorded by Plato or the views of Plato presented through the character of Socrates. However, this problem is less significant that it may seem at first because Socrates' main contribution to philosophy is his dialect method of inquiry and not the conclusions that we can draw from his dialogs.

The important thing that we learned from Socrates is that philosophy begins with relentlessly questioning the meanings to which concepts

refer. Clarifying the meanings of concepts is important because philosophical conclusions can be no more clear than the understanding of the concepts upon with they rely.

Plato, Socrates most famous student, was born in 427 BC and died in 347 BC. This is to say that Socrates died (or was put to death) when Plato was 28 years old. Plato lived to be 80 years old and during his long life wrote many dialogs showcasing (and perhaps idealizing) Socrates in his pursuit of the truth. As was mentioned before, it is not entirely clear, when we speak of Socrates, whether we are speaking of the historical person or the fictional character of the dialogs. Nor do we know exactly how well the fictional character faithfully represents the real Socrates. But all this uncertainty is less problematic that it may appear at first. If we attribute the method of inquiry presented in the dialogs to Socrates and the philosophical content to Plato then it becomes a little easier to sort things out.

Plato responded to the intellectual turbulence of ancient Greece by looking for a philosophical source of certainty and enduring truth. He longed for the old days when society was more orderly and disliked the turmoil that eventually led to the death of his teacher, Socrates. In his search for intellectual certainty he found mathematics and used that as a model for much of his thinking. A triangle would always have three sides, no matter what the sophists might say. And the sum of the angles would always be 180 degrees whether the senate of Athens liked it or not.

Looking at the physical world, Plato noticed that physical objects such as trees are all different from each other and thus no individual tree could be presented as representing the essential concept of treeness. He posited that there is an essential concept of treeness, however, and each individual tree is an imperfect instance of that essential concept. Thus, Plato preferred to work with idealized essential concepts rather than the imperfect instances found in the physical world. When we develop mathematical ideas, scientific theories, or even simulation models we are relying on the Platonic view that ideal concepts are somehow superior to the instances that they represent. Idealized concepts do provide us with two primary benefits. First, they are much easier to work with than the imperfect and inconsistent instances of the real world. Hence, idealized concepts provide intellectual economy in organizing our knowledge and thinking through our ideas. Second, idealized concepts allow us to envision our world in ways in which it does not currently exist. Thus, the idealized concepts of Plato provide us with a means of getting from how things are to how things could be or how things should be.

The important thing that we learn from Plato is that the world is an imperfect, inconsistent and sometimes messy place. If we idealize concepts they can be much easier to work with and if we can describe a better world intellectually we may be able to bring about a better world.

Aristotle, who was, in turn, Plato's most famous student, was born in Macedonia in 384 BC and did not move to Athens until 366 BC. This means that he never met Socrates and when he arrived in Athens much of the turmoil had died down. Aristotle stayed in Athens until Plato's death in 347 BC and four years later became the tutor of Alexander the Great. Aristotle viewed Plato's world of idealized concepts with some skepticism. Since Plato could not point to the idealized concept of treeness, it simply did not exist for Aristotle. Aristotle preferred to organize knowledge based upon physical characteristics. He would say that trees have properties in common with each other that make them trees. Aristotle's emphasis on observation and physical properties makes his philosophical view the basis for modern empirical science. However, the drawback with Aristotle's approach is that in order to develop a concept of trees based upon common properties one would have to have all instances of trees available in order to abstract the common properties. If Aristotle had never seen a pine tree, for example, when coming up with a definition of tree, he might conclude that it wasn't a tree because it didn't have leaves. Nonetheless, Aristotle's emphasis on observation and evidence for organizing our knowledge and testing our ideas has been the bedrock of modern scientific thought.

The important thing that we learn from Aristotle is that we need evidence to support our conclusions and that evidence should be objectively observed rather than contrived to support our preconceptions.

In order see the differences in philosophical perspectives between Plato and Aristotle more clearly it is useful to examine a couple of somewhat exaggerated examples. In The Republic, Plato presents his view of the ideal form of government, which is led by a philosopher king. If asked whether or not the ruler could be a woman he would probably reply that if she were sufficiently enlightened with regard to philosophy there would be no reason why she couldn't be the ruler. If Aristotle were asked if women could be rulers he would probably point out that there were no female rulers and hence that is not an appropriate role for women.

Aristotle comes out looking a little silly in that example so in the next example we will pick on Plato. Assume that Plato and Aristotle are looking at a horse and a by stander asks – How do you know it is a horse? Aristotle would describe the physical characteristics of the horse including number of legs, head shape, average weight and the like. He would then point out that the horse in question conforms to that description. Plato would said that we know it is a horse because it conforms, somewhat imperfectly, to an idealized concept of horseness. If the student pressed Plato to find out how he knows about this idealized concept of horesness, Plato would say that this idealized concepts exists only in a nonmaterial world known as the World of Forms and that the constituents of that nonmaterial world can only be apprehended by the intellect. If pressed, Plato would have to admit that this nonmaterial World of Forms is more real than the physical world because the physical world is imperfect whereas the World of

Forms is perfect, eternal and unchanging. The more Plato was pressed on this point, the sillier he would sound.

It is easy to take the views of Plato and Aristotle to extremes thus making their positions look silly, and this was certainly not the intent of the two previous examples. But their philosophical views, at the core if not at the boundaries, are the mainstreams of western thought. In summary, Socrates contributed a method of relentless questions, which is useful in clarifying our understanding of concepts. Plato preferred to work with idealized concepts and possibilities for perfection, while Aristotle preferred to work with the evidence as it presented itself. We will come back to these ideas over and over again as we explore philosophical issues in information systems.

Turmoil in the Information Age

In the twentieth century, we have seen enormous social upheaval, not the least of which is the advent of the so-called "Information Age". Information technology has changed many of our habits and routines, while forcing us to face new personal, social and political issues. The collection of large amounts of personal data in databases has led to the questions - What is personal privacy? and Is privacy a right ? The distribution of software and the availability of information over the Internet has given rise to the question - What is ownership?; and more specifically - What is stealing ? The promises of artificial intelligence have led people to ask - What is thinking? and What is intelligence ?

The introduction of information technology in organizations causes power shifts leaving one group with greater power and another group with less. This shift in power leads us back to the classic Socratic questions - What is fair? and What is just ? Unreliable software, bad data in databases, and irresponsible behaviors of computer hackers have caused us to revisit personal responsibility and recall the classic questions of moral philosophy - What is right? and What is virtue ?

At a larger social level, information technology has impacted our lives in significant ways. It is creating a skills gap between the technically literate and the technologically illiterate. Instead of calling in an order for pizza, we fax it in. Instead of calling someone on the telephone we send him or her an email. Instead of sitting around the television at an appointed time to watch our favorite television show, we record it and watch it at some other time. "Don't tell me the score from last night's game, I haven't watched it yet !" Imagine a courtship ritual that begins with "I'd like to get to know you better, do you have a home page?" Sometimes we sit back and wonder where all this technology is going? On the balance is life better or worse? This reflection leads us back to basic questions such as "What is the purpose in life ", "What is quality of life?", and, more recently, "How does technology affect the quality of life ?"

Finally, information technology affects us at even the most metaphysical levels by casting doubt on the extent to which we know things. Conventional wisdom tells us that a rock is real because we can drop it on our toe and feel pain. However, in today's world

of concepts, the backdrop of physical reality is no longer adequate. Computers do not obey the laws of physics and thus mock the notion of things being real because they have physical substance. Suppose a company is suffering from declining sales. Suppose further that the 'sales picture' is represented by a dozen different means from reports, to spreadsheets, to simulation models. Which representation reflects 'the reality' of the situation? Recently advocates of computer simulation have argued (somewhat persuasively) that simulation is superior to experimentation for theory building because real data is messy and experiments are usually faulty due to issues of repeatability and variable control. Simulations, on the other hand, can control variables and can be repeated as often as is necessary without any confounding influences, suggesting the Platonic view that that theory is somehow 'more real' than the data. The role that information technology plays in mediating our interpretations of reality has, once again, opened up the ancient philosophical questions What is real? and How do we know ? Pushing this question to the limits, where moral philosophy and metaphysics intersect, we can ask the ultimate questions (the design questions), not What is? but What should be ? and How do we know what should be ?

I called this book *The Ghost of Socrates* because, if we could conjure up the Ghost of Socrates, he may well be able to help us gain insight into the most vexing problems in information systems. In the following chapters, we will apply the Socratic approach to concept analysis and critical questioning to a variety of philosophical issues in information systems.

Chapter 2 – Knowledge Management

Knowledge is Power – Francis Bacon

As organizations pursue their concerns, the individuals within those organizations gain experience regarding how the business performs most effectively. As people move around within organizations or from one organization to another, much of that experience is lost. It seems reasonable that the accumulated experiences of the individuals within an organization constitute a significant component of the knowledge of the organization. It further seems reasonable that this knowledge should be viewed as a resource of the organization and should be managed to the benefit of the organization. In fact, this all sounds so reasonable that many organizations have embarked on projects to develop knowledge management systems with exactly these objectives in mind.

On the face of it, it all seems obvious. The more knowledgeable a company is with regard to its business, the more productive it would be. Managing the knowledge of the company effectively would make the company more knowledgeable and hence more productive. So it should be fairly easy to justify knowledge management. Unfortunately, after just a few well-chosen questions the whole concept begins to unravel. You begin by asking – What are you going to manage in the knowledge management systems? To this question you should get the obvious answer – knowledge. But when you ask the next question – What is knowledge – the focus of the systems becomes a little more elusive. You might get answers including memos, project reports, or case studies. You might get

answers involving historical data. But it is quite likely that you will not get a definition of knowledge. Instead you will get examples of things that are going to be stored in the knowledge management system. You could observe that everything mentioned seems to be a document of some kind so why don't they just call it a document retrieval system. Or you could ask how one knows whether or not the items that will be managed in the system really constitute knowledge. But neither of these questions is likely to get an answer, because knowledge management is based on some rather poorly understood concepts. In order to get to the bottom of this fuzzy understanding we will begin with the misleadingly simply question - What is knowledge?

An Illustrative Example

Imagine you just heard the news that your local ice-skating rink was taken over by Eskimos. How would you validate this claim? The first thing that most people would do is to turn on the local news to see if there is any mention of the event. If the event is not mentioned on the local news, that doesn't mean that it didn't occur. It only raises doubt about the validity of rumor. If the news program verifies the claim, then it increases the likelihood that the claim is true, but it certainly does not make the incident an absolute truth. If it is April 1st or if the news program is on a local campus news station, there is always the possibility that the news story is a hoax or a joke. Even if the news program is real, there is a chance that a misinterpretation may be occurring. The invaders might not be Eskimos. They may be somebody else dressed up as Eskimos. The ice rink may no longer be an ice rink. It may be a building that used to be the ice rink but is now used to

conduct animal research or some other endeavor that may draw protestors. Perhaps, it wasn't really taken over. Perhaps, these people showed up for some other reason and their presence was viewed as a take over. So even if the news program agrees with your rumor, you cannot be absolutely sure that the rumor is true. There could be misinformation, or it could be out of context.

You could check with the local police. If they confirmed the information on the news, then that would give it greater credibility. If they did not confirm, or if they provided different information, or outright denied the story, then the truth of the rumor would be more difficult to resolve.

You might try to reason it out. Certainly, Eskimos are from a place where there is a lot of ice so the ice rink makes some sense. But your local ice rink may be thousands of miles from the Eskimo's native land, so that might cast doubt about the story. If Eskimo's had threatened this action in the past, or had actually done this before, that might give some credibility to the story. But if they had never actually done this before and there was no precedent for this happening, it might detract from the credibility.

Unable to get solid confirmation from the evening news or the police, and unable to reason it out for yourself, you might decide to pay a visit to the ice rink. So you park a block away, and sneak up to a window on the dark side of the building. You see people on the inside dressed in Eskimo outfits, so the rumor isn't completely without merit, but you still cannot be sure that the ice rink has been taken over by Eskimos. The people in

the Eskimo garb may or may not be Eskimos. You can't really tell if they have taken over the ice rink or are simply meeting there. So you have given further credibility to the story, but your knowledge of the situation is not absolute.

Over the next few days you talk with your colleagues about the incident. You watch news shows, listen to the radio and read articles in the newspaper. Over time the reported stories become more consistent and there are fewer and fewer challenges. Eventually you feel comfortable that you know the truth about the rumor as well as you are ever going to know it.

The Testing of Knowledge

None of the techniques used above to validate the rumor managed to provide absolute confirmation. Indeed, this is also the case with knowledge. Absolute knowledge is unattainable, but by testing knowledge with a variety of techniques over time the knowledge becomes more reliable. Following are a variety of knowledge testing techniques that were alluded to in the above story.

Appeal to Authority

Using the appeal to authority technique of testing knowledge says that you do not know what is true, but there is somebody who does know; or at least somebody who knows better than you do. In the example above, the news program or the local police are

viewed as credible sources for the veracity of the rumor and appealing to their authority tests the rumor.

In other situations the appeal to authority test appears in other forms. If somebody says "Aristotle said this," or "Plato said that," they are appealing to the authorities of Aristotle and Plato respectively. Research papers that cite copious references are appealing to the authority of earlier writers, to some extent, for the foundations upon which their claims rest. Frequently, we defer to experts or widely accepted sources or information such as the dictionary, encyclopedia or the law to provide justifications for our claims.

It is important to remember that an appeal to authority is not an absolute confirmation of knowledge. The authority could be wrong. They could be misunderstood. Their comments may have been taken out of context or applied incorrectly to a different situation. However, if some authority does support the veracity of a piece of knowledge, it does lend further credibility to it.

Appeal to Reason

Sometimes we just reason things out. We consider the evidence and try to determine if such a claim is logically possible. More formally, we attempt to deduce, from the premises, whether or not the claim can be logically justified. In the above example, the connection between the ice rink and the cold climate where Eskimos come from seems to make sense. But the distance they would have had to travel did not seem to make sense.

Appeal to reason can take on a variety of forms. In its simplest form we just try to figure out whether or not the claim makes sense given the constituent components of the claim. Or we might see if the claim makes sense given other things we know. We like to maintain some level of cognitive consistency so we attempt to insure that new things we learn are consistent with the things we already know.

At a more formal level, appeal to reason becomes theoretical research where knowledge is organized into categories and relationships are specified between or among the categories. At the extreme of formality, appeal to reason becomes pure mathematics and symbolic logic. But few people go that far in order to validate their knowledge. Plato was a strong proponent of appeal to reason. He felt that observations were error prone and the material world was too inconsistent to support true understanding.

Appeal to Observation

You have probably heard people say, "Seeing is believing." This is the essence of the appeal to observation. If you want to know if something is true, you need to see it for yourself. In the example above, going to the ice rink and looking in the window to see if there are Eskimo's in there is an appeal to observation. You need to see something with your own eyes in order to believe it.

An appeal to observation can take on a variety of forms. At the most rudimentary level you need to see something with your own eyes in order to believe it. At more formal levels we find empirical science with elaborately constructed experimental and data gathering techniques attempting to derive knowledge from observation.

Of course, appeal to observation is not fool proof. You might be viewing a cleverly constructed ruse. Or you might draw the wrong conclusion from the observations that you make. Nonetheless, the appeal to observation says that if the map and the territory don't agree, believe the territory. Aristotle was a strong proponent of appeal to observation. He believed that knowledge must be based on supporting observable evidence. Abstract reasoning can go awry. But, the facts, are all, are the facts.

Repeatability

Whether a piece of knowledge is verified by appealing to authority, reason or observation, it must be possible for the verification to be repeated by another independent party. For example, if a person were to claim that the veracity of the Eskimo rumor was revealed to him or her through divine intervention, then the verification would be in doubt because it would not be easily repeatable. Similarly, if a person claimed that the truth came to them in a dream, the repeatability would be in question. When appealing to authority, a person must be able to cite the authority and others must be able to consult that authority independently in order to give greater credibility to the knowledge.

If the knowledge were verified with an appeal to reason, then it should be possible that the reason be presented to others and it would seem reasonable to them also. Frequently, what is reasonable to one person is not at all reasonable to another so having others examine the reasonability of a claim weeds out a lot of claims that are the results of faulty thinking, biases, or wishful thinking.

In the case of appealing to observation, it must be possible for others to experience the same observations. If somebody claims, for example, that they viewed an unidentified flying object, then it must be possible for others to make the same observation in order for the claim to have credibility. Carried to an extreme this is the essence of experimental and empirical sciences where, theoretically, the observation made can be repeatedly observed by any number of independent and objective parties.

Repeated Testing Over Time

Even if we have used all available techniques to test a knowledge claim, it is still possible that we have introduced our own biases into the process. For example, in the scenario above, assume that a person is particularly suspicious of Eskimos and believes, without justification that they will someday take over the county. Assume further, that the local news station is putting on an April Fool's story and that some volunteers have dressed up as Eskimo's and have converged at the ice rink. This person would appeal to authority see the newscast and would probably believe it. When this person reasoned through the claim, it would be consistent with their somewhat paranoid view of the world. And if

they appealed to observation by going to the ice rink, they would receive further confirmation of the credibility of the claim. Since one person cannot see beyond their biases, it is necessary to have the knowledge tested by more people as the repeatability criterion requires. However, if this paranoid person were to check with a similarly paranoid friend, it is unlikely that they would see through the ruse. Hence, not only is it important to have knowledge verified by multiple people, but it is important that the additional people have no other goal than the objective verification of the knowledge.

Consequently, it is important that the knowledge claim be repeatedly tested over time. Facts or evidence that was not available at one time might be available at another time. Further, any group of people may share an interpretive bias of some kind. Hence, the more the knowledge is tested over time, the more likely it is that new evidence or the lack of biases will yield a different conclusion. Consequently, if a piece of knowledge is continually tested without being over turned, the more likely it is to be true.

Triangulation and Convergence

None of the techniques discussed so far have been able to provide us with absolute knowledge. The best they can do it to increase the credibility of a piece of knowledge. To further increase the credibility we should test that piece of knowledge using as many techniques as we possibly can. For example, if appeals to authority, reason, and observation all lead to the same conclusions, then we can feel more comfortable that the knowledge in question is credible. And if disparate views begin to converge over time

then we can feel more comfortable that our knowledge is solid and is not likely to change.

Social Aspects

Since the veracity of knowledge claims must be verified by other people it is clear that knowledge is a social phenomenon. If a person cannot articulate a claim to others then there is no way in which it can be verified and hence it cannot be knowledge. If person can articulate a claim but others cannot repeat the verification procedure, then the claim cannot be verified and cannot achieve the status of knowledge. Implicit in the need for others to repeat the verification procedure is that it must be possible to explain the verification procedure. This says that in order to be knowledge it must also be possible to justify a claim to others and this gets us to heart of what it means for a claim to be considered knowledge.

What is Knowledge?

So far we have looked at ways in which we can test the credibility of knowledge but have not yet addressed the fundamental question of what is knowledge. This question can be quite complicated and has occupied philosophers for centuries. In fact, there is an entire branch of philosophy, or more specifically and an entire branch of metaphysics devoted the to problem of what constitutes knowledge. This branch of metaphysics called epistemology is concerned with determining how we separate knowledge from other

things that appear to be knowledge but are not. Fortunately, we don't have to get into the more subtle questions of epistemology because we are not searching for ultimate truth. We are only looking for knowledge that remains true within timeframe of our concerns.

What are some of the things that appear to be knowledge but are not? There are dreams, hallucinations, opinions, beliefs, and illusions among the things that are not generally considered to be knowledge. We will examine each of these and see how each differs from credible knowledge. But first, we need a working definition of knowledge.

In this case, the dictionary is not particularly helpful because the dictionary provides definitions that reflect the variety of usages of the word rather than a definition that helps us focus on the centrality of the concept of knowledge. However, the Encyclopedia of Philosophy is helpful in offering a definition that is not only useful, but can be easily comprehended by the average person. According the Encyclopedia of Philosophy knowledge is "justified true belief." This is a wonderful definition in that it is concise, comprehensible and useful. But before we can use it, we must examine it's constituent parts. We must ask what we mean by justification, truth, and belief. This is a minor digression, but one that will prove useful in gaining an understanding of the essence of the concept of knowledge.

Belief is a conviction or assertion about something in the world. You may believe that all grass is purple and that is a legitimate belief. It isn't true. But it doesn't have to be true in order to be a legitimate belief. It only has to be an assertion about the world. You may

also believe that all unicorns are purple. If you do not believe that unicorns exist, then this is not a legitimate belief because it is not an assertion about something in the world. However, if you believe that unicorns exist, then the belief that they are all purple is a legitimate belief.

Belief alone does not constitute knowledge. The belief must be true and it must be justified. In the sections above, we introduced some methods for testing knowledge and that is part of the justification. We will return to that shortly. But first we must digress again and address the question of what is truth.

Again we are taking on a deep philosophical problem, but we will not get any deeper into it than we have to. To say that something is true is not a statement of fact, as some might think. This is because absolute truth is no more attainable than absolute knowledge. To say that something is true is, instead, an assertion of likelihood. To say something is true is to say that it is likely that other people at other times or under other circumstances are likely to agree that it is true. Truth is determined socially and the social organism that determines truth extends across space and time. So when you say something is true, you are saying that a larger social organism, people from other places and other times, are likely to agree with you in this particular assertion about the world. For example, the claim that grass is purple is not true because it is very unlikely that other people across time and space with their own biases and agendas will agree with that statement. However, the statement that grass is green is true because it is likely that other people

across time and space with their own biases and own agendas will agree with that statement.

Truth and knowledge sound a somewhat similar in that they require lots of people over time and space to agree. However, the difference between truth and knowledge is belief. Something can be true but people don't believe it. For example, prior to Columbus people supposedly believed that the world was flat. It was true, at the time, that the world was round, but people did not believe that. So the assertion that the world is round would be considered true, but would not be considered knowledge.

Justification is also an important component in knowledge. In Plato's view, knowledge had to have three features. It had to be true. You had to believe it. And you had to be able to explain it. Roughly around the same time as Plato there was a group of philosophers called the Atomists who believed that matter was made up of tiny particles called atoms. As it turns out, they were right and the claim that matter is made up of tiny particle is true, at least according to atomic theory. So they believed this claim and it was true. But it was not justified. They really had no way of knowing that matter was made up of tiny particles called atoms and this assertion was at best a lucky guess. But it would not be credible as a knowledge claim.

Dreams, Beliefs, Illusions, and Opinions

Now we can turn to other things that look like knowledge but are not. First we can dispense with dreams and illusions. An illusion is a false idea or conception. Since

knowledge must be true, an illusion cannot be knowledge. Dreams are a little trickier. Suppose for example, that a person has a dream that it is going to rain the next day and, indeed, it rains. Could that dream be considered knowledge? Before answering, lets push the issue a bit. Suppose that this person dreams of rain every evening before it rains. Further they never dream of rain prior to a day when it does not rain. So their dreams agree with the weather one hundred percent. Granted that this scenario is very unlikely we can still ask - Do these dreams constitute knowledge? The answer is – No. The person believes that it will rain and it does rain so the belief is true. But it is not justified. The person cannot explain how these dreams come about so the dreams cannot be considered as knowledge. Assume, however, that after exhaustive testing somebody figures out that the person is unusually sensitive to the changes in barometric pressure that precede rainstorms. Assume further that these changes in barometric pressure cause the person to dream about rain. In this case it may be considered as knowledge because the mechanism by which it occurs is understood. And in fact the belief that it will rain is justified.

An opinion is a belief that is not based on certainty or positive knowledge but on what seems true, valid or probable in one's own mind. It would be easy to get into a tar pit of semantics as we compare knowledge with opinions. Absolute knowledge is unattainable so, to some extent, all knowledge is opinion. We get into this snare because the concept of opinion is also very fuzzy. If somebody asserts that the whole world is crazy then that is his or her opinion. And, as we say, everybody is entitled to his or her opinion. However, we also have expert opinions where knowledgeable people speculate or offer claims that are not completely verified. These opinions are a lot closer to knowledge than

off hand claims made by people who have no obligation to truth or veracity. So we can view the range of opinions as the steps toward knowledge, which is to say that opinions are claims subject to verification. If the opinion turns out to be true and we understand why it is true then it crosses the rather fuzzy line into the realm of knowledge.

Another Example

We can put all of this together, and take it a step further, with another simple example. Suppose you have guests over for dinner and upon leaving one of your guests asked you if you know where his coat is. You say that the coat is in the upstairs guest room and that is a knowledge claim. It is a belief and it may be true. In fact, to simplify things we will assume that it is true. But we have not yet verified it. You may appeal to authority by asking your spouse if he or she put the coats in the guest room. You may appeal to reason by reasoning that the downstairs coat closet is always stuffed and hence the next most logical place for the coats is the upstairs guest room. In this case an appeal to observation is probably the best method of verification. That is you might go upstairs to the guest room and see if the coats are, indeed, there. Finally, after verifying your knowledge claim you respond to your guest that you do indeed know where the coats are and to avoid any philosophical entanglements you simply retrieve the coat and let the guest go on his way.

As your guests leave your knowledge claim that you know where the coats are begins to deteriorate. When all the coats were in the guest room, you could claim that you knew where the coats were but as the guests disperse this is no longer true and your knowledge

deteriorates. This suggests that not all knowledge has equal status. Some things are true for a brief period of time and some things stay true for much longer. Clearly, in the interest of intellectual economy we should prize knowledge that endures over knowledge that deteriorates. So endurance is an important property of knowledge.

Let's go back to the guests in your house and assume, once again, that your guest asks you if you know where the coats are. You reply to your guest that it is the middle of summer and people don't wear coats in the summer. This knowledge is much more general than the specific knowledge as where the coats are on a given day. It not only remains true over time, but it is true over a wider variety of situations. Again, in the interest of intellectual economy knowledge that is more general and covers more situations has greater status than knowledge that is specific to single situation. So generality is another important property of knowledge.

This is to say that, somehow the knowledge that the earth is round has greater standing than the knowledge that your friend's coat is in the upstairs guest room. And the knowledge that all rotating planets take on a spherical shape is more valuable than the specific knowledge about the earth.

Back to Knowledge Management

So far we have defined knowledge as a true belief that we can justify. We have looked at some of means of justifying knowledge. And we have observed that knowledge should be

as enduring and as general as possible. Now we can turn to knowledge management and ask if it is, indeed knowledge that is being managed.

We need to distinguish between two types of knowledge that are the focus of knowledge management systems. The first is explicit knowledge and the second is implicit knowledge. Explicit knowledge is in the form of memos, project reports, historical data of various kinds and case studies. Implicit knowledge, sometimes called tacit or background knowledge are the things that people have learned over time about the way things operate as a result of their experiences. It is called implicit knowledge because it is frequently difficult for people to articulate what they know.

We can group memos, project reports and case studies together in that they represent the written opinions of the authors about a situation that is not repeatable. It is a snapshot of reaction to a unique incident or time. It is very unlikely that these items have been verified in any rigorous manner and even more unlikely that the verification could be repeated any number of times by independent observers. Thus, these documents will probably never achieve a status greater than opinion. However, as we have seen there is a great range of opinion from simple fancy to expert opinion. Hence, not knowing where on the scale these items lie, they are all very suspect when the goal is collecting knowledge.

More rigorously produced data such as the structured reports derived from computerized systems probably can qualify as knowledge since, theoretically, the programs can be run

any number of times on the same data and the same result will be produced. However, there are still two problems with structured historical data. First, the knowledge derived from data is only as good as the data. If data is collected based upon poorly defined concepts or poorly understood measures, then all of the results derived from that data are of questionable value. Second, even if the data is perfect, it represents a single moment in time for a single organization. It is not enduring and it is not general. Hence, its standing in realm of knowledge is not very great.

Knowledge management experts like to talk about tacit or implicit knowledge. However, given what we know about knowledge from the previous discussion, it seems like tacit knowledge is an oxymoron. That is to say that knowledge cannot possibly be tacit or implicit. In order to be knowledge a claim must be a justified true belief. If it cannot be articulated it cannot be shared. If it cannot be shared, it cannot be justified. If it cannot be justified, it cannot be knowledge.

To explore this idea of tacit knowledge a little further consider a baseball player who has a great batting average and an uncanny ability to hit home runs. You might say that he really knows how to hit home runs. That may be correct in everyday speech, but the fact is that he is able to hit home runs, but he may not know how he does it. From his perspective, the ball comes at him, he swings and the ball goes over the fence. Unless he knows the specifics of the method that he uses it would not be correct to say that he knows how to hit home runs. It would only be appropriate to say that he is able to hit home runs.

Why Call It Knowledge Management?

It seems like most of the stuff contained in a knowledge management systems would not pass even the loosest tests of knowledge. So why call it knowledge management? Why not call it opinion management or document management. The most likely reason is that most people understand that knowledge has great value. They also understand that opinions and other things have less value. In order to justify the costs of developing knowledge management systems it is necessary to claim that they are repositories of knowledge. But, generally, they are not. But that is not to say that they are not repositories of other valuable things. Opinions have value. Historical records and historical data have value. In fact there are lots of other things such as intuition and artistic creativity that have great value but do not qualify as knowledge. In order for people to communicate effective we have to use precise words that refer to clear concepts. When we use fuzzy ideas and concepts it nearly always leads to misunderstanding and inefficient progress toward goals. Knowledge management experts should figure out exactly what it is that they are trying to achieve with knowledge management systems and justify the systems according. But calling it knowledge in order to justify it is simply misleading.

Chapter 3: Concept Analysis

When we name things correctly, we comprehend them correctly, without adding information or judgments that aren't there. – Epictetus

Suppose a person vehemently proclaims that he or she is opposed to software piracy. What do you do? You can simply ignore them and try to justify it to yourself by saying that you just don't have time to deal with this issue. You can accept their claim on face value and adopt the position that you are opposed to software piracy also. You may decide that you are ambivalent about this issue but don't wish to explore your ambivalences. Or you might decide to take the bull by the horns, as they say, and explore the issue more fully. Realistically, if you are to have any kind of a legitimate opinion, you must to the latter. You must dig into the issue and decide what you really think about it. But how do you do this?

The first step is to realize that the issue as presented already contains built in biases. The phrase 'software piracy' is a bad thing on the face of it. The word 'piracy' is not a neutral word. It refers to pirates and people who take things that don't belong to them. So nobody could be in favor of software piracy and the question becomes - what do you think of the activity that software piracy refers to? The difficulty, of course, is that software piracy refers to a collection of activities some of which you might support and some of which you might be against. So it may be prudent to turn

back to the person who made the claim for some clarification since piracy is such a vague and value laden word.

When asked for clarification, the person is likely to respond with something like "I am opposed to software piracy because it is stealing!" There may be any number of other valid responses but we need one to work with so we will take this one. We have made some progress because we now have the allegation that the activities that comprise software piracy fit neatly within the set of activities that we call stealing. If this is really the case, then software piracy is a bad thing because stealing is a bad thing.

On the other hand, a person who copies a software program might justify their action by saying that what they did was not really stealing. They may justify or even rationalize their activity by saying that they were just borrowing the software for evaluation, stretching the license, or simply making a backup copy. What they are saying is that the concept of stealing defines certain behaviors that do not include the action they just took. In order to figure out who is right we need to refine our understanding of the concept of stealing and concept analysis is the technique we use to do this.

Think of a concept as a sphere. Things that are clear instances of that concept are inside the sphere and the closer to the center they are the closer they are to the essence of the concept. Things that are clearly outside of the sphere are clearly not

instances of the concept. The boundaries of the sphere are fuzzy. While many things are clearly in or out, many other things are on the boundary some more in than others and some more out than others. Finally, we can see the concept sphere reflected on other surfaces where the concept is being applied to other meanings. Looking at these other meanings can be misleading because the essence of the concept has been transformed to include other things. Our goal in concept analysis is to define the concept sphere. We want to determine what is clearly on the inside and what is clearly on the outside. We want to know what is on the boundaries and in what way the essence of the concept has been transformed to include other meanings that are not part of the essence of the concept. We can begin our concept analysis by following a few simple steps.

Look for Exemplars: Look for clear examples that are as close to the essence of the concept as possible. These examples should be undisputed by any reasonable person. If it is not possible to find exemplars, then, possibly, the concept is too poorly defined to be useful. In the case of stealing we can imagine a burglar breaking into an appliance store in the middle of the night and leaving with a television set. It is hard to imagine any circumstances in which this would not be stealing. Although there should be agreement on exemplars, in a group setting a variety of exemplars may reflect a variety of different understandings of the concept. If people begin with different understandings of a concept, then any conclusions based on those different understanding might be quite different also.

Look for Counterexamples: Look for clear examples of things that are similar to the concept but clearly not contained within the concept. These counterexamples should also be undisputed by any reasonable person. In the case of stealing, find examples of actions that are like stealing but are clearly not stealing. For example, a customer comes to the store during working hours and purchases a television set. It is hard to imagine any circumstances in which this would be considered stealing (even if the customer got a really good deal and declared that is was a steal). As with exemplars, there should be agreement. However, in a group discussion a wide variety of irresolvable counterexamples may indicate widely varying understandings of the concept.

Look for Borderline Cases: Look for examples that are not clear. Borderline cases help to focus on key attributes and make up the fuzzy area between the exemplars and the counterexamples. For example, what if the thief in the exemplar was also the owner of the store? Can a person steal from his or her self? What if the person in the counterexample had changed the price tag? What if a person purchases an appliance and pays for it. However, upon leaving the store, they notice they got the wrong color. So they put back the one they paid for and took one they didn't pay for. Is it stealing? Borderline cases help us clarify our understanding of the concept.

Identify Attributes: List the essential characteristics of instances that are included in the concept in question. These attributes should apply to all exemplars and no counterexamples. The borderline cases should provide some refinement on these

attributes. With the identification of attributes we begin to get some insight into the essence of the concept under consideration.

Consult the Dictionary: When asked to provide a definition of a concept, some people would prefer to defer to the dictionary. This is certainly much easier than trying to figure it out for your self. However, the dictionary provides definitions that reflect common usage that is usually not sufficiently clear for concept analysis. Nonetheless, consulting the dictionary can be a useful data point in clarifying a concept. The dictionary will usually list several definitions that reflect the variety of ways in which a word it used. These definitions can be used to further discussion of the concept. In addition, the dictionary will usually provide an etymological derivation of the word, which is frequently useful to see what the heart of the concept was at one time or will provide insight into the essence of the concept. Stealing, for example, is defined as taking another's property without permission. Still, it is important to not take dictionary definitions too seriously. If a person buys a car but fails to make the payments the car may be repossessed. The repossessor almost certainly is taking the car without permission, but it is clearly not a case of stealing.

Consult Legal Definitions: Sometimes there are legal definitions that may be useful. However, like dictionary definitions, legal definitions cannot be taken as the final word. Sometimes legal definitions are constructed to fit within the legal system and do not agree with the intuition of the average person. Nonetheless, they are the

usually the result of serious deliberation and consideration and hence may provide insight into the essence of the concept.

Continuum Construction: Come up with a list of all instances mentioned so far and and any more that you can think of that belong, to some degree, to the concept. Rank all examples in a continuum from clear cases to counterexamples. Try to determine, especially for the borderline cases, if there is a single attribute or a collection of attributes upon which they differ.

It should be mentioned that concept analysis is not a mechanical process and the preceding steps are merely a guide. Practice is important here. The more concepts one attempts to analyze the better one gets at it. While the steps just presented help to develop an understanding of the essential concept, the following suggestions will help to test and refine that understanding.

Construct an Essential Definition: Attempt to come up with as simple a definition as possible that covers the essential attributes of the concept. This may vary greatly from the dictionary definition that reflects common usage or the legal definition, which is used to establish facts in court. The essential definition is intended to determine whether or not specific instances should be included within the concept and thus serves a conceptual membership role rather than the role filled by other definitions.

Inclusion of Particulars: Search for specific examples and determine whether or not they are examples of the concept. This tests the strength and merit of the essential definition.

Socratic Shuffle: Socrates would often keep his interlocutors off balance in the following way. He would ask them to explain a give concept. If they provided instances of the concept he would ask for a general principle that covered all instances. If they provide a general description of a concept he would provide an instance that was not covered in the general description. In some sense this is a little unfair since there probably is no general description of a concept that includes all instances. Nonetheless, it is a useful technique for exploring subtleties of a concept.

Invent Cases: Sometimes it is necessary to construct cases to test essential attributes. These contrived cases may be unlikely to occur, or they may even be impossible, but they help clarify key attributes of the concept in question. For example, what if the shop owner and the thief are the same person. He came to the store in the middle of the night because he needed a present and forgot his key. But, he did not record the fact that he removed a television set. Is he stealing?

Embed Within a Social Context: The meaning of a concept is often defined by its social context. For example, if the owner of the store was your favorite uncle you might be less inclined to call the action stealing. If the owner was you next-door neighbor who repeatedly dumped his lawn clippings in your back yard, you might be

a little less lenient. We often define concepts based on who or what we like or dislike rather than based on search for precise and useful definitions.

Explore Metaphorical Uses: Many concepts have essential meanings as well as metaphorical meanings. The metaphorical meanings must be examined in two ways. First, the metaphorical meaning is an instance of the concept applied to elucidate another situation and as such may contribute to a deeper understanding of the concept. Second, since the metaphorical meanings do apply to a different situation they should be identified and excluded from the original concept. For example, if a person says – "Officer arrest that woman. She stole my heart!" we are using a metaphorical meaning of stole which does not belong in the original concept.

Concept analysis does two things. First, it helps individuals clarify their thinking. And second, when used in a group discussion, it reveals the different understandings that individuals may have and helps develop a consensus view of the meaning of a concept. However, in order to be productive in getting to the heart of a concept, a group discussion must have several important features.

Members Must Take Turns: It is important to get everybody's perspective on a concept under scrutiny. If one or two people dominate the discussion then the activity become biased by their opinions. Concepts are usually quite rich and it is important to get as many perspectives as possible in order to recover the richness of a concept.

Members Must Be Respectful of Other's Opinions: In order for members to contribute freely to a discussion, other members must be respectful of their opinions, regardless of how different they may seem initially. In a group setting, concept analysis is a group search for understanding and that understanding cannot be achieved if some members of the group are reluctant to offer opinions.

Members Must Be Skeptical and Critical: While it is important to be respectful of the opinions of others, it is also important not to accept the opinions of others at face value. Members of the group must be convinced of the validity of the contributions of others, otherwise the concept under investigation cannot be clarified.

Member Must Share a Common Goal of Achieving an Understanding of the Concept Under Investigation: The purpose of a group discussion in concept analysis is to achieve a greater understanding of the concept in question. Thus, group members should be more committed to the progress of the analysis than they are to their individual opinions. Effort to derail the ultimate goal of the discussion should be avoided.

In the Platonic dialogs, Socrates was the de facto discussion leader. In this role his goal was to elucidate philosophical principles though the process of a question and answer discussion. For the average group leader, leading a question and answer discussion can be difficult especially when discussion participants may be reluctant

to contribute for fear of saying something wrong or stupid. Follow are some standard types of questions to keep the discussion going in a productive direction.

The first type of question is a process facilitation question. These questions are used to make the group discussion process productive. The discussion leader may ask questions such as "Is this topic even worth discussing?", or "Mr. Smith, what do you think of Mr. Jones' claim ?" Process facilitation questions are important for two reasons. First, they help to get group members to buy into the discussion format. And second, they keep the discussion going in a productive direction.

The second type of question is a concept refinement question. These questions were discussed earlier and serve to help the group in achieving a greater understanding of the topic under consideration. Clearly, this type of question assumes that the group members have accepted the group discussion format and hence rely on the success of the first type of questions.

The third type of question attempts to get at underlying justifications or inconsistencies in a group member's position. In this category there are three types of questions: rational, emotive, and disposition to act. Rational questions are of the form: "What do you think about X ?" and calls the person's reasoning processes into play. Emotive questions are of the form "How do you feel about X ?". Emotive questions call for an affective response and may not agree with rational claims. Questions that inquire about a disposition to act are of the form: "What would you

do in case X ?" and seek to determine what the person would do if called to act in a situation. Asking questions of these forms identifies inconsistencies in the participant's moral concepts and sets the stage for further inquiry into the nature of the concept.

When you make a claim such as “I am opposed to software piracy,” it is important that you know what you are saying in order to for your claim to have any merit. In order to know what you are saying it is important to understand the essence of the concept involved in the claim. This becomes even more important if the claim is a moral claim or a claim that demands action of some kind. Concept analysis, as presented in the this chapter, is a useful technique for getting to the essence a concept and thus helping us to clarify the content of our claims. The remaining chapters in this book will apply techniques of concept analysis to philosophical issues in information systems with the goal of shedding light and providing greater insight into some of these complex issues.

Chapter 4: What is a Right to Privacy?

Liberty is the right to do whatever the law permits. – Montesquieu

We can continue our exploration of philosophical issues in information systems by examining an issue that has moved, very rapidly, to the foreground of our concerns in the past couple of decades – the right to privacy. At first people were simply concerned about the existence of electronic databases containing personal information. Having this information available electronically made people feel vulnerable to invasions of privacy. Even though much of this information was available anyway, having it in electronic form allowed others to piece together information about a person and create a richer picture of that person's habits and personal history.

Recently that concern has grown as the volume of electronic information collected about the average person has grown, particularly information collected on the World Wide Web. But the Web not only collects information, it makes it more easily available, at least potentially, to anyone in the world. So larger and larger collections of personal information are becoming more and more easily accessible and concerns about personal privacy are increasing.

Unfortunately, most of those concerns are visceral reactions to a poorly formed concept called invasion of privacy. And the advocates of protecting privacy in cyberspace exploit those visceral reactions to great advantage. Privacy advocates would have you believe that privacy in cyberspace is just like privacy in your own home. And having someone view electronic information about you is the same as having someone break into your home and catch you in the shower with a video camera. Certainly, nobody would like to be caught in the shower on videotape and consequently nobody would like to have electronic information about him or her viewed in cyberspace.

Most people would agree with the following statements – Privacy is a right. We must protect personal privacy. And, restricting access to personal information in cyberspace can help protect personal privacy.

But if you were to ask – How did privacy become a right? Or, exactly what does that right entitle you to? You might very well get a blank look or even an agitated look. If you ask – What is the downside to privacy? You would probably get a shrug. And if you asked – In general, are you in favor of keeping information hidden from public view? Or, do you believe in censorship? You would probably get firm negatives in both cases.

All of this suggests that the issue of privacy in cyberspace may not have been as well thought out, as it should have been. This chapter will look at the concept of rights and the concept of privacy. We will ask – What are rights and where to they come from? Then we will explore this issue of privacy and attempt to determine what the phrase 'right to privacy' really means. Then we will carry this discussion a step further by looking at the issue of protecting privacy in cyberspace. Current thinking on the issue of privacy is rather heavily one sided, and one of the goals of this chapter is to provide some balance in our thinking. But before we get too far ahead of ourselves, let's go back to the issue of rights and where they come from.

What Are Rights and Where to They Come From?

A right is a privilege afforded to an individual by a social group or organization that allows that individual to exercise discretion in a particular set of circumstances simply by virtue of participating in the social group or interacting with that organization. That sounds a little lofty, so let us consider a few examples.

The most trivial example would be rights afforded to a consumer who interacts with a particular grocery store. The grocery store may proclaim that you have the 'right' to return any product with which you are unsatisfied. Let's examine this situation a little

more closely. First of all, this 'right' means that you have an opportunity to exercise your discretion in a certain area. In this case, you are allowed to decide whether or not your money should be refunded for a particular product. If you feel that a given product is unsatisfactory, you can request a refund for it and the grocery store must honor that request.

Assume that the grocery store begins loosing money due to this policy. Is this 'right' inalienable? That is to say, can this right be taken away? Of course it can. The grocery store can simply change its policy. Of course, the policy cannot just be changed haphazardly. The store may have to announce the change. Or they may be required to put up signs and provide some sort of a grace period for returns from the prior policy. But the right can be revoked. It is not an inalienable right.

Next, assume that you are shopping at Store A, which has a return policy as described above. Then, one day, you purchase a product at Store B. Once you get home and sample the product, you find it is completely unsatisfactory and take it back to Store B. Store B claims that it does not have the same refund policy as Store A and consequently refuses to provide the refund. You may protest that being able to return an unsatisfactory product is your right and they would probably respond that this is your right only when shopping at Store A. The point is that a right is granted by a social group or organization and a different group or organization may or may not grant that same right.

This raises the question of whether or not there is such a thing as an inalienable right, since one group may grant a right while another may not. Worded differently, are there any circumstances in which an individual has discretion that cannot be taken away. The answer to this will be addressed shortly.

Next consider the right to free speech that is granted by the Bill of Rights in the U.S. Constitution. Citizens of the United States generally enjoy the freedom to speak their minds about political issues. But let's say that a citizen of the United States takes a vacation in the Despotic Islands, a fictional group of islands in the Pacific somewhere.

The Despotic Islands are ruled by an insecure tribal king who has proclaimed that there is a $100 fine for anyone saying anything against the king. So the vacationing U.S. citizen gets off the plane and remarks, “the king does a terrible job maintaining the airport.” Immediately, the citizen is fined $100. The citizen may protest that they are a U.S. citizen and consequently have the right to criticize the government. Where upon the King of the Despotic Islands may point out that the right to free speech exists within the United States, but does not extend to the Despotic Islands. So even rights that we consider fundamental are still a function of the social group or organization with which we are dealing.

Even rights that we consider fundamental are not absolute. For example, the right to free speech has many boundaries. The most widely sited example is the fact that one does not have the right to yell “Fire” in a crowded theatre. Here the right to free speech is in conflict with public safety and in this case public safety wins.

Why do we have certain rights? The answer is simply that lawmakers and statesmen felt that society would be better off if people had those rights. The framers of the U.S. Constitution could have decided that an individual has the right to shoot another person if they don’t like their looks. However, in their wisdom they chose not to make that a right. Just like a grocery store may grant rights to make people better consumers, governments grant rights to make people better citizens.

We hear a lot of talk of basic human rights and inalienable rights. It bears looking at these two expressions for just a moment. First of all, an inalienable right is a right that cannot be taken away. So, is there such thing as an inalienable right? Is there a situation in which an individual exercises his or her discretion and that choice cannot be taken away?

Since rights are socially conferred, they can be taken away if the society decides that the right is no longer to the benefit of society. Even rights that are agreed upon can be taken away from certain individuals if they have transgressed in some way. The only right that cannot be taken away from a person is the right to decide whether or not to participate in

a rights conferring social unit or organization. If I choose not to shop at Grocery Store A, there is nothing they can do to make me shop there. Similarly, if I choose not to be a member of Society A, then there is nothing Society A can do to force me to be a member. Granted nonparticipation may ultimately be costly. If Grocery Store A is the only store in a two hundred mile radius, I may have to drive a long way to buy groceries or I may have to grow my own food. But I cannot be forced to shop there. Similarly, if I choose not to participate in Society A, I may have to leave the country, or I may be jailed if I don't. But I cannot be forced to participate. So, other than the option to participate, there are no inalienable rights.

What about basic human rights? If rights are a matter of social convention does that mean that there is no such thing as basic human rights? The answer to this question is a little tricky. Certainly, basic human rights are not somehow part of the fabric of reality. They are, however, a part of our social reality. That is to say that there are certain rights which most people in a given type of society would choose to have. The problem with basic human rights is that they are basic to a particular social organization and not all possible societies are the same.

For example, one might feel that the right to own property is a basic human right. Certainly it is fundamental to our modern view of life. However, it is possible to imagine a healthy functioning society where all property is shared. Further, it is possible to imagine that anyone in that society who wanted property for him or her selves might look selfish to the others. So the right to own property is basic to one social structure and worldview, but not to another.

It should be mentioned, before concluding this exploration of the concept of rights, that rights cannot be determined in a haphazard fashion. Rights must be agreed upon by the members of the social unit and must be somewhat consistent with the rights of surrounding social units. Further, they must be exercises of discretion that are likely to continue to be of value in the future.

When people refer, with great reverence, to their right to privacy, they are assuming that an individual discretion with regard to their privacy has been worked out and agreed upon. Further they assume that this agreed upon right is exactly consistent with their own rather poorly thought out and often visceral concepts of privacy. While there are aspects of our lives in which our right to privacy is fairly well understood and agreed upon, privacy in cyberspace is still up for grabs. Consequently we need to think through the issue and decide what is good for both individuals and society in terms of our rights to privacy in cyberspace.

What is Privacy and Why Do We Have a Right to It?

As we have seen, rights are not absolute. Rather they are largely a matter of social consensus. And, the value or importance of certain rights may be different at different historical times or in different social or political circumstances. So in order to examine the concept of a right to privacy we need to ask – What is privacy and why should there be a right to it?

Here it might be useful to begin with a dictionary definition of privacy which is roughly that privacy is the quality or condition of being withdrawn from public view or secluded. Privacy is derived from a Latin word meaning to separate. When something is private we are separating it from that which is public and keeping it to ourselves. So a right to privacy is a right to keep some things to our selves and out of public view.

This seems fairly reasonable. People need to feel autonomous and allowing them the right to define themselves in public and keep parts of their life private helps in providing individual autonomy. Further, the cost to society or to other individuals of allowing each person to keep some things to them selves is probably not all that great. People also need to have a place where they can let down their guard without worrying about public scrutiny so protecting privacy in our homes also seems very reasonable. In fact we have a wide variety of laws that support these rights to privacy. From laws against Peeping

Toms and illegal wiretaps, to constitutional prohibitions against illegal search and seizure and self-incrimination we feel that it is important to allow individuals to keep certain things to themselves.

It is important to note at this point that there is a benefit to this right to privacy as already noted, but there is also a cost. I might want to know your dirty little secrets or I might like to eavesdrop on a conversation you are having with your spouse or lover. But I don't get to do that. It is hard to argue that there is a great cost in frustrating my prurient interests but there are other examples where the costs are more real. For example, the police have to work a little harder to catch you if you are doing something illegal. They cannot just tap your phone or break into your house and force a confession out of you. They must get court orders for some intrusions and other intrusions are not allowed under any circumstances. On the balance, the benefits of allowing people to have some space of their own where they can feel safe, autonomous and away from public scrutiny seems to outweigh the costs substantially. Thus, we believe fairly strongly in a right to privacy. But it is important to remember that this right is a matter of social consensus and not a part of the fabric of reality. Further, this right, as with all rights, has limits.

How far does this right to privacy extend? It seems to cover my home and my thoughts, but does it go any farther than that? Here it is useful to consider the legal definition, which is based on the concept of "a reasonable expectation of privacy." Since it is not possible to describe all circumstances in which the right to privacy might apply, the legal criterion is based upon what a reasonable person might expect. If I am visiting a friend, I might expect to have my privacy protected even if it is not my house. In fact, if I am in Washington DC and I call a friend in California, I have a reasonable expectation that my telephone call will be private even though the conversation leaves my house and travels over three thousand miles of country side. The telephone call is private for several reasons. First, the telephone line is, in some sense, an extension of my house. Second, a reasonable person would expect it to be private. Third, the benefit to the public of having private telephone calls is greater than the benefit to the public of allowing anyone who

wishes to do so to eavesdrop on my conversation. But how far can we carry this right to privacy.

Here continuum construction, one of the techniques discussed in Chapter 3, comes in handy. We will construct two continua, the first dealing with extensions of private space and the second dealing with private communications.

If you were at home in the shower, you would certainly expect your right to privacy to protect you. If someone peers in your window, you can probably call the police and have them arrested. However, assume that a tornado passed through your town a few days ago and your house is missing a wall so that your shower is exposed to passersby on the street. If you were to take a shower and demand that everybody on the street be arrested, you would probably be laughed at or dismissed as a lunatic. Once the wall was destroyed, you no longer maintained your reasonable expectation of privacy. Let's take this a step further and say that in the middle of a heavy downpour, you decide to take a shower on your front lawn. Can you reasonably expect privacy in this situation? Of course not! In fact you might be arrested for public indecency. So your right to privacy is only protected in those situations in which you have a reasonable expectation of privacy and those reasonable expectations don't go very far before they bump into restrictions or conflicting rights.

Next consider your right to privacy when it comes to telephone conversations. It is reasonable to expect your right to privacy to be upheld when you are talking to someone on the telephone. Suppose, however, that you are using a public telephone in a shopping mall. You might expect your conversation to be free from wiretaps, but what if the person at the public phone next to you overhears your conversation. Does your right to privacy include the person next to you? Probably not. Suppose you wish to proclaim your undying love to your significant other so you hire a skywriting airplane to write a love message across the sky. Can you expect others to not look up lest they invade your privacy. Of course not. The inconvenience to others certainly overrides your right to

privacy in this case and since rights are always a trade off, you loose your right to privacy when communicating with your loved one via skywriting.

So a right to privacy is a right to have some space in which you can keep things to yourself away from public view. This right is important because people need to feel safe and autonomous and need to have a place where they can let down their guard without worrying about public scrutiny. However, the right to privacy is not all-inclusive nor is it absolute. Beyond a reasonable expectation, the right to privacy does not apply. Nor does it apply in some cases when it comes into conflict with other rights or when the benefits of protecting privacy do not outweigh the costs or lost benefits.

Is There a Right to Privacy in Cyberspace?

Privacy in cyber space really gets down to the issue of whether or not people have a reasonable expectation of privacy in cyber space. This in turn gets down to the question of whether or not people should have a reasonable expectation of privacy in cyberspace. And that, of course, gets down to the simple question of what do you gain and what do you loose. Privacy advocates would say that privacy must be protected at all costs. In the next sections we will consider both sides of the issue of protecting privacy in cyberspace and see that the right answer is not all that obvious.

When evaluating the right to privacy in cyberspace we have to consider two main dimensions of the issue. The first is, considering the costs and benefits of protecting privacy in cyberspace, what is an appropriate allocation of the rights involved. And second, given a fair evaluation of those trade offs, does or should a person have a reasonable expectation of privacy in cyberspace. We will begin by examining some of the typical reasons that people give for wanting to protect privacy in cyber space and evaluate the merit of these claims. Then we will look at the downside of privacy, an aspect of the issue that is rarely acknowledged. Finally, we will look at reasonable expectations given the trade offs presented.

The claims and counter claims have been given names in the following discussion simply for the sake of clarifying key arguments. These names are not by any means universal and another commenter may break up the arguments differently and name them differently. Hopefully the names and the organization of the arguments will prove useful in increasing our understanding of the issues at stake.

The Claims For Why We Should Protect Privacy in Cyberspace

There are several recurring themes regarding the need to protect privacy in cyber space. For purposes of simplicity we will group these into three categories: the nuisance claim, the protective claim and the iconic claim. Each of these claims is legitimate to some extent, but none is an absolute call to action. That is to say that the concerns are real, but must be considered in terms of alternative solutions or the cost of remedy through suppression of information.

The essence of the nuisance claim is that if a person's activity in cyber space is not kept private, solicitors of various kinds will annoy them. For example, if I go to a bookstore site on the web and buy a book on Tai Chi I don't want to receive phone calls at dinnertime from people trying to sell me Tai Chi classes. And I don't want my electronic mailbox, nor my regular mailbox for that matter, filled with Spam and junk mail from solicitors trying to sell me Tai Chi related products.

A darker instance of the nuisance claim, which is still prevalent in many forms, is the concern that if I go to an adult web site and purchase adult materials, I do not want my mailbox filled with solicitations from vendors of this type. Particularly, I may not want my family or roommates to know what kind of questionable materials I have been purchasing. Adult materials are a rather heavy-handed instance of this problem, but one can easily see how other, more benign materials such as political or religious materials might also fall into this category.

An extreme instance of the nuisance claim would cite the fact that vendors may be able to construct a consumer profile, which allows them to predict, with a high degree of accuracy, what kinds of products I am most likely to buy. Knowing my vulnerabilities, they fill my mailbox with solicitations or post banner ads the next time I appear at the web site with the result that I land up purchasing things that I would not have purchased otherwise. The extreme instance of the nuisance claim suggests that the impetus for shopping behavior is taken away from me and given to the vendor.

The protective claim suggests that protecting privacy in cyber space protects me. That is, if my privacy is not protected, then bad things can happen to me. In this category we see claims such as one might be denied a job due to information acquired about him or her from the various data banks in cyber space. For example, if I apply for a job at a conservative trade association and they find out that I have been downloading materials from a politically opposite web site, they may decide not to hire me. There are limits on what an employer can ask on a job application but there may be no limits on what they can find out about you in cyber space.

A more serious form of the protective claim suggests that things in my past might keep me from being hired. For example, if I have a criminal history, I might be denied a job. But my past does not need to contain anything as clear as a conviction for a heinous crime. What if I have tickets for jaywalk or parking illegally? Does that show sufficient lack of character to deny me a job for which I am otherwise qualified?

Taking the protective claim a step further, what if I am denied medical insurance due to a medical condition? Does my right to privacy allow me to apply for insurance without revealing all of my medical history? Maybe, maybe not. It may depend on the specific history. But what if my medical history is a little more vague. What if I appear to be "accident-prone"? Does an insurance company have a right to deny medical insurance to me because I might be a high risk for clumsiness? Taking this a step further, what if my genetic code makes me a high risk for certain genetic diseases. Can I be denied medical insurance based on a genetic predisposition?

The protective claim says that you do not always know the reasons for which you may be denied a job or medical insurance or any other thing than another person may be granted. The decision maker will simply look of the available information and make a decision. Hence, it is better to provide as little information as possible in order to avoid having unfair or detrimental decisions being made.

The iconic claim is a little trickier than the nuisance or protective claim because presumably no action is involved in the iconic claim. This claim asserts that the information available about you in cyber space is somehow a representation of your person. When somebody views information about you in cyber space they are in fact viewing you. In order to fully understand the iconic claim you have to ask "Would you want the whole world to know that you rent adult videos?" Or, "Would you want the whole world to know that you collect fertility statues?" Or, "Would you want the whole world to know that you have a fascination with explosives?"

A more insidious form of the iconic claim would ask if you would mind if some stranger in cyber space were to piece together a profile of your behaviors from the time you wake up until the time you go to sleep. It is important to distinguish between the iconic claim and the previous two claims in that in the iconic claim a person may know something about you but they do not act upon that knowledge. If they act upon that knowledge, then it becomes either the nuisance claim or the protective claim. The iconic claim suggests that just having somebody know something about you can be to your detriment. Certainly, if somebody catches you in the shower on his or her video camera they have gained knowledge about you and you have been diminished in a variety of ways. So the iconic claim does have merit.

Responses to the Claims

The claims for protecting privacy in cyber space are not without merit. Neither are they absolute calls to action. This section will provide responses to the claims, not to diminish them, but to put them into perspective.

The advocates of protecting privacy in cyberspace have cleverly managed to connect privacy in cyberspace with privacy in the home. The net result of this is that most people react viscerally to the privacy issue. The above claims make people feel exposed, vulnerable and violated. And it is difficult to overcome those feelings with critical arguments. This issue will be taken up in greater detail in the final chapter.

There are three general responses to the nuisance and protective claims. They are: the surrogate response; the faulty assumptions response; and the irrational clumping response. The essence of the surrogate response is that suppressing access to personal information in cyberspace is being used as a surrogate solution to other problems. The essence of the faulty assumptions response is that if you examine the assumptions behind the nuisance and protective claims they appear to be somewhat faulty and perhaps even contradictory. Finally, the irrational clumping response asserts that there are many different kinds of information in cyber space that are being clumped together. Some of that information may very well be appropriate for restricted access while other pieces of information should clearly be publicly available.

Another way of stating the surrogate response would be to say that what others do not know about you shouldn't hurt you. Perhaps you believe that you should never be denied a job or medical insurance based on personal information. Allowing you discretion over the release of personal information would allow you to keep any potentially harmful information out of the hands of decision makers. Unfortunately, restricting access to personal information is simply being used here as a benign surrogate for the real issue. The real issue is - what criteria are acceptable for denying a person a job or medical insurance? If we feel that a person should not be denied medical insurance based on their past medical history, then there should be legislation to that effect. If we feel that a person should not be denied a job based on things that occurred in their past then that

should be legislated. Suppressing personal information is a clumsy, heavy handed, round about and counterproductive way to achieve these goals.

Imagine the following job interview:

Interviewer: Good morning, Mr. Smith, please tell me a little about you.

Smith: Well, I can't really tell you about myself because that information is private.

Interviewer: Do you have any relevant experience for this position?

Smith: Yes, but I can't tell you about it because it is private.

Interviewer: Perhaps you can tell me about your educational background.

Smith: I could, but I prefer not to. You see it really is private.

This could go on, but the point should be clear. An employer interviewing a person for a job needs some information upon with to make his or her decision. If we feel that certain criteria should not be taken into consideration then we should pursue legislation to that effect rather than making all personal information off limits.

The faulty assumptions response asserts that if you examine the assumptions behind the nuisance and protective claims they appear to be somewhat faulty and perhaps even contradictory. For example, if a vendor obtains your past purchasing information they may call you at dinnertime trying to sell you something. Of course, this could be terribly annoying. But why would a vendor want to annoy prospective customers. If a vendor annoys its potential customer base enough, those potential customers will not want to do business with that vendor and the vendor will go out of business. So it would be foolish for the vendor to behave in that way. Further, being annoyed by vendors is not a product of making purchases in cyber space. Mailing lists and telephone lists are currently being

passed around by vendors and charitable institutions; regardless of whether or not a person has been shopping in cyber space. In the case of telephone lists, you can ask that your name be removed from the list. Perhaps similar rules need to govern email. But suppressing information is not the way to do it.

Perhaps the most extreme instance of faulty assumption response is the assumption that vendors possess both extreme intelligence and extreme stupidity. The extreme intelligence assumption lies in the fact that we believe that vendors are clever enough to create a profile of our purchasing behavior and then present products to us that we simply cannot resist purchasing. The extreme stupidity assumption lies in the fact that we believe vendors will annoy us with banner advertisements and spam email until we no longer want to do business with them. It is unlikely that most vendors possess both extreme intelligence and extreme stupidity. And the contradiction inherent in this assumption shows the fundamental irrationality of our concerns.

These extremes can be seen in the case of hiring also. If an employer is clever enough to do a thorough background check on a prospective employee then how could they be stupid enough to deny employment to a qualified candidate based on irrelevant background information.

Irrational clumping is another example of an irrational reaction to the privacy issue. There are many different kinds of information available in cyber space and different kinds of information must be treated differently. However, rather than examining the wide variety of types of information, most reactions to the privacy issue are the same. If it is personal information then access should be restricted.

If I send you an electronic message then the fact that I sent you a message is somehow a different kind of information than the contents of that message. Most people are far less concerned about having others know that they sent a message than they are concerned about having others know what was in the message. If I send you an email from home it is different from sending an email to you at work.

Certain kinds of information are more 'naturally occurring' in cyber space. For example, web logs and online purchases are part of the cyber information space. My bank records, may be in electronic form at my bank, but something would have to be done in order to make them available in cyber space. Other information, such as the shows I watch on cable TV might require some substantial effort to be recorded electronically and then additional effort might be required to make that information available in cyber space.

There is some basis to the irrational clumping claim in that every time you turn around you find another piece of information that you thought was private appearing for public view in cyber space. It appears that all information is potentially available and hence all kinds of information get lumped together.

The iconic claim is a little trickier to address because it represents an essentially visceral belief that a person is somehow diminished by having somebody simply know something about them. As was mentioned before, the iconic claim assumes that no action was taken otherwise the incident would become an instance of the nuisance or protective claims. How is it that a person is diminished by having another person simply know something personal about them? One possibility is the confusion between iconic images and real things. Primitive peoples discovered by anthropologists early in the last century were reluctant to have their pictures taken because they were afraid that the camera was stealing their soul. In some sense, constructing an electronic profile of a person is stealing their soul in the same way. But the question we have to ask is, can a person be damaged simply by having another person know something about them? This is a complex question that must be broken down into several cases.

First of all, most people would agree that a Peeping Tom looking in your window is a fairly flagrant violation of your privacy. Even though they have not done anything but look there is a real threat of physical danger. Although it is possible that all they want to do is look, most people would feel the risk of the person taking further action is great. So having somebody simply look in your window represents a real physical threat and thus

makes one feel unnecessarily vulnerable. Other invasions of privacy such as planting recording devices in your home or tapping your telephone also represent real threats, although perhaps not physical threats. One would not bug your house or tap your phone simply because they were mildly interested in your activities. Such an action would be done for a specific purpose which undoubtedly would be to the detriment of the person whose privacy was violated. So these are not really iconic threats. They represent a real danger to the person.

Second, if your neighbors or coworkers find out something about you, you may loose social status or even opportunities to advance in your career or social standing. So in these ways you can also be diminished.

However, in the case of someone who will never meet you and never have contact with you knowing something about you, it is hard to see how you are diminished. If somebody finds something out about you and tracks you down then it becomes a physical threat. If one of your neighbors or coworkers finds information about you in cyberspace then it may turn into case two. However, if somebody has no connection with you and no impact on your life it is hard to see how they can diminish you in any way. Thus, it is hard to see how the pure iconic threat is a threat at all.

Now we need to turn back to the first two cases where information discovered about you in cyberspace leads to a physical threat or a diminution in your social status. That is, somebody finds something so compellingly interesting about you that they wish to make physical contact or something so distasteful is discovered that you suffer a loss of status. The argument for protecting privacy in these cases is inherently contradictory. If there is a compelling fact about you in cyberspace, there is probably that same compelling fact about a million other people. This piece of information is only compelling in your case, if you are the only one that it is known about. If the same piece of information is known about everybody for whom it is true, it is no longer such a compelling piece of information. Similarly, for information that would result in diminished social status. If that piece of information is known for everyone for whom it is true, it is no longer as

important as it would be if it were leaked out about one person. For example, if someone were to say to me that they know that I watch cable TV, I would respond that nearly everybody watches cable TV. The point here is that information about a person usually attracts attention because of its novelty. Novelty is due to a belief that this piece of information is rarely true and knowing it provides additional insight into the person whom the information is about. Since certain types of information are not widely available, we believe that certain activities are more novel than perhaps they really are. Thus, suppression of personal information leads to greater interest in those specific cases when it is released.

For example, lets say that the IRS audits your neighbor and the audit reveals that he cheated on his taxes. This is a really interesting piece of information about your neighbor because it reveals that he is a tax evader. However, this is only interesting because you know this piece of information about your neighbor and not about others. A full audit of all taxpayers may very well reveal that a large percentage of taxpayers cheat on their taxes in some way. Some may be small or trivial violations or even errors while others may be flagrant attempts to avoid paying large tax bills. But the availability of this information on everyone makes that availability of this information on one individual less novel, less interesting, and certainly less damaging.

Perhaps the real risk might be that if all this personal information were readily available we might find out that there is nothing at all interesting about our lives. Certainly the mystery of hidden information creates interest and allows the possibility of fantasizing that others are dying to find out about us. But the benefits of pretending to be interesting must, as with the other perceived benefits be weighed against the corresponding costs. And it is those cases that we turn to next.

The Downside to Privacy in Cyberspace

Perhaps one of the reasons that people feel so strongly about protecting privacy in cyberspace is that there is no reason not to. They feel that privacy protection comes without social and personal costs while the economic costs are so vague and systemic that they don't seem to affect anyone directly. But there are some downsides to protecting privacy and one does not have to look very far to find them. This section will consider the economic, social, and personal costs associated with protecting privacy in cyber space against which the benefits can be weighed.

It is difficult to say much about the economic costs of privacy without gathering data and proclaiming specific costs such as privacy in medical records cost so many million dollars per year in copied medical records, duplicated labs tests and conflicting or redundant treatments. Perhaps inventory costs could be reduced by a few percent if vendors had a better idea of what purchasers were likely to buy. Or advertising costs could be reduced if advertisements could be shown to only those consumers who are likely to interested in a given product.

From the perspective of the consumer there are also economic costs associated with privacy. For example, if I choose to buy a few mystery stories to find out whether or not I like mysteries, I would have to buy a crate of books in order to find out what I might like. It is likely that in that crate I would find a few books that I really like and a lot that I don't care for at all. It would be nice if one of the bookstores on the web would make some recommendations but since they don't know anything about me, they can't say whether I am more likely to appreciate British or American mysteries. Nor can they tell which of the sub genres I might be disposed to like. The economic cost to consumers can be seen in bookshelves filled with unread or half read books, and even more dramatically in basements filled with purchases that seemed, at the time, to be a good idea.

All of this is to say that a lack of information leads to uncertainty and uncertainty is costly. Consider the following examples. Assume you are having a party and you hire a person to plan the party. The person asks you if the party is going to be for five, fifty, or five hundred. You don't know. The party planner would just give up at that point because

trying to plan a party that could be for anywhere from five to five hundred people would be extremely difficult and outrageously expensive if you were to go ahead with the plans trying to cover all possibilities. Alternatively, assume that you are buying clothes for a new job, but you don't know what the new job is. It could be anything from a construction job to a job in the executive suite. Trying to buy a wardrobe to cover all those possibilities would be prohibitively expensive.

This is the same problem that businesses, insurance companies, governments and other social, political or economic institutions face when trying to operate on limited, imperfect or uncertain information. These costs are then passed back through to the consumer in the form of higher prices and less effective service.

The social costs of privacy are the social effects of suppressed information. Suppressed information leads to gossip, unsubstantiated supposition, rumors, innuendo and outright defamation as the information vacuum is filled with false or manufactured information. The desire to suppress private information in cyber space stems from a lack of trust that we have in each other and our institutions. We believe that our neighbors will think less of us or won't accept us if they know certain things about us. We are afraid that vendors will trick us into buying things that we don't need or want. We think that the government will spend its time building cases against citizens by piecing together disparate pieces of information. We think that employers and insurance companies will act against their own best interests just to frustrate our interests. This is not to say that none of these fears are real. However, if these fears are real we should deal with it through social policies. If they are not real then we should not have to carry the burden of those fears. It seems a little funny that we create rumors by producing information vacuum but don't want anybody to act on those rumors. On the other hand we don't want to provide exactly the information that would fill that vacuum and make decision making at the social, political and economic levels more rational and more effective.

If one is viscerally attached to the concept of protecting privacy, then it is easy to ignore the economic and social costs. After all, those costs are apportioned over a very large

number of people and the cost to a given individual is not that great given the perceived cost of invaded privacy. However, when one considers the cost to one individual of protecting the privacy of another individual that begins to change. What we have here is a clash between one person's right to privacy and another person's right to know. Actually, the economic and social costs also involve a clash better a right to privacy and a right to know, but nowhere is it more poignant than in a clash between individuals.

One of the classic, and rather heavy handed examples, of the clash between one person's right to privacy and another person's right to know is when a sex offender moves into a neighborhood or when a sex offender applies for a job at a day care center. In these cases the need that parent's have to protect their children seems to make their need to know outweigh the privacy rights of the sex offender. This example is heavy handed because having a sex offender move into the neighborhood conjures up images, for many people, of children being molested and thus eliminates any opportunity to explore the issue rationally. Most people would say, "I am all for protecting personal privacy, but not in the case of sex offenders." So this case doesn't really get you very far.

It doesn't get you very far because on one hand you have a visceral reaction in favor of protecting privacy. Having your personal information viewed in cyberspace makes you feel like you have been caught in the shower with a video camera. You certainly don't want that so you feel that protecting privacy is of paramount importance. On the other hand, the idea of a sex offender moving into your neighborhood and molesting children results in such strongly negative visceral reactions your visceral reactions in favor of protecting privacy are overwhelmed. Hence the claim, "I am all for protecting personal privacy, but not in the case of sex offenders," accommodates both visceral reactions but hardly provides a rational basic for policies regarding privacy.

So we need to examine some cases that are a little less emotionally loaded. We will begin with a very benign case. Let's assume that I have gotten on to my university library's search system and located a book that seems to be crucial to some research that I am doing. Unfortunately, the book is currently checked out to another faculty member. This

is disappointing but it also represents an opportunity. Perhaps another faculty member is interested in the same area of research and we might have an opportunity to collaborate or share ideas. After all academia is suppose to be an arena for exploring, developing and sharing new ideas. So I contact the librarian to find out who checked out the book. And guess what? The librarian can't tell me because it violates the privacy of the borrower. So much for collaborating and sharing new ideas.

This policy of protecting the privacy of the borrower is certainly not without merit. There have been some rather egregious instances of federal investigators trying to obtain the library records of individual who are suspected of seditious activities. A person should certainly be able to read whatever he or she wants without having the government come after them for reading the wrong books. However, the policy is not with cost either. And that is the heart of the privacy dilemma. We are often so gung ho to protect privacy that we run rough shod over other rights such as the right to know without ever considering whether or not that right has any claim in the situation being considered. The right to know is an important right also and should not just be pushed aside in our rush to protect privacy.

The right to privacy makes individuals feel safe and autonomous. The right to know also makes individuals feel safe and autonomous in that it allows them to know things they may need to know about others to increase their safety or allow them to oriented themselves with respect to others.

In the next example we will take the conflict between the right to privacy and the right to know a step further. Let's say that my son gets a "C" in math on his report card. Having no idea how he stands in relation to the rest of the class makes it difficult for me to know how to interpret this information and hence makes it difficult for me to know what to do. If everyone in the class got A's and B's, then I would interpret it one way. If there was one A, one B, one C and the rest failed, I would see it differently. If one group in the class did markedly better than another group in the class - say girls versus boys, well behaved versus restless, neatly groomed versus the latest rebellious style – that would tell

me something else. If this particular teacher routinely received poor evaluations and routinely graded harder than the norm then that would tell me something else. However, I am not privy to the information I would need to know these things. The privacy of the other students and the privacy of teacher protects them from having this information made readily available. My right to know has come into direct conflict with the right to privacy of others.

Most people are concerned about their right to privacy when it comes to medical records. So, in the next example, we will consider medical records. Let's say that I go to my doctor for some malady that is not immediately obvious. My doctor runs some tests, can't figure out what is wrong and refers me to a specialist. The specialist in turn runs some additional tests in which he or she tracks some lab results over time. Then in the middle of this medical investigation, I change jobs. My new insurance requires me to change doctors. When I go to my new doctor she asks me to have copies of my medical records sent to her, which I do. At some point I get transferred out of town and have to find another new doctor. After several such job changes and moves, I have been to a half a dozen doctors, an equal number of specialists and have had dozens of lab tests done. In addition, getting copies of all my medical records is becoming increasingly more difficult because some of my doctors may have left the practices in which they worked when I saw them. Further, some of the practices may have been bought out by larger corporate practices. The end result is that it has become extraordinarily difficult for me to assemble a complete set of records that reflect my medical history. If would be nice if all my medical records would be in one central location accessible via the web so that any doctor I visit can have access to my full medical history. But this is impossible, of course, because making all this information accessible from a central location via the web could potentially have a negative impact on my right to privacy.

One of the concerns in the protective claim is that information about a person may be used against them in some way. This is particularly insidious if a decision is made on a particular piece of information and the individual does not know which piece of information led to the decision. But what if a decision is made about a person on

incorrect information? What if I have a bad credit report due to incorrect or false information? What if an employer chooses not to hire me because credit worthiness is considered to be an important trait? I might not think of checking my credit records and might go for quite a while having incorrect information used against me. Privacy protects my credit records from public scrutiny, but it also protects incorrect information. If I can easily check my credit over the web, I can be sure that it does not contain incorrect information. If others check my credit record, they may turn up errors and tell me about them. In short, the more people who check information, regardless of what their reasons may be for looking at that information the more likely it is that the information will be correct. To a large extent, privacy protects incorrect information.

These examples have been somewhat concrete. Next we move to some conceptual concerns in which the right to privacy must be evaluated, from a more humanitarian perspective, against a lesser recognized right – the right to know.

People often think of their right to privacy as an issue of individual rights versus the rights of society. And to a large extent this is true. However, right to privacy also involves the rights of one individual against the rights of another individual and it is on this level that we will turn next. A person's right to privacy is important because people need to feel that there is a part of the world where they are safe from public scrutiny. And people certainly should have a right to feel safe and autonomous. At the same time people also should have the right to know things about the world around them. The heavy handed case of the sex offender brought this issue into clear focus and the examples library loans and school records showed how one person's right to know comes into conflict with the privacy rights of others. But while withholding information from me regarding library loans or the grades of other students may restrict my intellectual curiosity or my ability to make decisions regarding my son's report card, neither instance diminished me as a human being. So at a more esoteric level we must ask whether or not I, as an autonomous human being, have the right to orient myself with respect to others in the world and whether or not I have the right to make sense out of my life. These two issues will be addressed next.

As an autonomous individual in a free society, I have a wide variety of choices to make about my life and how I spend my resources, especially time and money. I make choices that will affect the availability of my resources. Some choices such as where to go on vacation or what car to buy are less important in the larger scheme of thing but are still important to me as an individual. Other choices such as where to live and whether or not to accept a job offer are more serious. Even more critical are decisions about what school to go to, what curriculum to pursue and what career path to embark upon. We can make these decisions on purely objective data such as comparative costs and expected outcomes. But often we turn to our friends and ask them what they have done in similar circumstances or what they might do given the choices we are facing. We do this because purely objective data sometimes misses things that are important to us as individuals. Hence, information from and about others is often key to our individual decision-making. But what if we were to turn to our friends for advice and they clammed up saying that the information we had asked for was private? We might feel betrayed or alienated. We feel that people somehow owe us information that we need to make good personal decisions. We feel that we have a 'right to know'.

Let's say that you have decided to accept a job with a new company that requires a transfer out of town. You consult with your friends, all of whom tell you that it sounds like a good deal, so you sell your house buy a new house in the new town and move in order to take the new job. Then after three months the new company lays you off. Since you did not stay for six months they claim no responsibility in helping you with you moving expenses or the expenses associated with your real estate transactions. You sell the new house at a loss and move back home having to rent because you don't have enough to buy back you house. Telling your tale of woe to your friends, two of them admit that the same thing happened to them. "It was horrible for at least three years trying to recover financially," one admits. You look at them in horror and say "Why didn't you tell me that? Don't you think I had a right to know that information?" No matter how strongly one might feel about privacy, there are cases in which a right to know is a conflicting and strongly competing claim.

Not only is information about others needed to make good decisions, it is also needed to help people make sense out of their circumstances and even their lives. People need to know about other people in order to understand themselves. Perhaps this is the reason we have seen such growth in the number of talk shows and reality shows in the past decade. People want to know what other people are doing and how they compare with others. The freedoms of the late twentieth century have give rise to a wide variety of alternative lifestyles and have allowed people to explore a wide variety of interests that they may have previously ignored. Unfortunately, too many of these are viewed as nasty little secrets that must be protected by privacy. But this protection leads to alienation and isolation. People seek out others with similar interests so that they can feel less alienated. We see this in a wide variety of clubs and interest groups. When somebody faces a medical or psychological crisis they may seek out a support group to help them deal with the problems and make sense out of their circumstances. People certainly have their right to privacy but they also have a right to know how they fit in with the rest of society and they certainly have a right to make sense out of their lives and feel comfortable about themselves. Know about others helps them make sense out of what they are doing. Letting people suffer thinking that they are abnormal or different is perhaps more immoral than any invasion of privacy could ever be.

Making Some Important Distinctions

When we talk about privacy in cyber space we are really addressing a wide variety of forms of information and no blanket policy will address everything. For example, the fact that you sent an email to such and such a person on such and such a date is less personal than the contents of that email. Some information, such as visiting web sites naturally occurs in cyberspace but other information such as criminal history or medical records does not naturally occur but may be placed in cyber space. Attitudes towards protecting privacy in cyber space often group all information together and have one solution which is to suppress all personal information. Any reasonable policy must be a little more

refined than that. For the purposes of this discussion we can make a few distinctions that may require different privacy policies. First, we can divide cyber space into three domains: public, restricted, and personal. Public cyber space is any web site that you can freely access over public Internet lines. Most web surfing, shopping and browsing would fall into this category. Restricted cyber space would include information that is available electronically but not available to anybody who wishes to look at it. Restricted cyber space would include data with access that is restricted in some way such as privileged users, corporate privacy policies, or paid subscriptions. Personal cyber space would include information that is stored in electronic form, but is not intended to be viewed by anyone who is not a user of the specific computer or a member of a trusted workgroup. This distinction is important because many people view privacy in cyber space as an all or nothing at all policy. That is, if you don't protect privacy then you have no privacy at all. If your foot print in cyberspace is not protected then the contents of your hard drive is not protected either. This would be a foolish policy and it would make are great deal more sense to create categories of information and decide privacy policies as appropriate for each category.

A second important distinction is between facts and content. For example, the fact that you sent an email to a specific person at a specific address on a specific date is a fact. The substance of the message is content. This distinction is important because policies may vary on facts and content. The fact that you sent a message to a given person reveals far less about you than the contents of the message might. Hence, making the fact public is potentially far less damaging than making the content public. Further, the benefit of the information derived from facts, which are generally highly structured information, is far greater than the benefit derived from slogging through content which is generally unstructured and possibly highly contextual. Hence, the cost benefit balance of facts versus content is quite different.

So when somebody talks about protecting privacy in cyber space you really have to ask them what they mean by cyber space and which categories of information they believe should be restricted.

A Reasonable Expectation of Privacy

Privacy in cyber space really gets down to the issue of whether or not people have (or should have) a reasonable expectation of privacy in cyber space. We have a reasonable expectation of privacy in our homes and that reasonable expectation extends to technologies such as the telephone that can be viewed as extensions of our homes. However, once we go to work or go shopping at the mall, we give up a lot of that privacy. The question then becomes do we view cyber space as an extension of our homes or do we view cyber space as a public space where it is not reasonable to expect to have our privacy protected. It would be nice if there was a simple answer to this question but ultimately we must answer this question by making distinctions about the public or private nature of cyber space, distinctions regarding different kinds of information, and then decide for each strata of cyber space and each category of information where the right to privacy or right to know prevails given the costs and benefits of each alternative.

Chapter 5: Computers and the Quality of Life

It is quality rather than quantity that matters. Seneca

You have to wonder, sometimes, where all this technology is going. In general, information technology helps us manage information more effectively. We can store and retrieve great quantities of information at incomprehensible speeds. We can search for just the right pieces of information; organize or summarize them if we choose; display them in a wide variety of pleasing formats; and then respond to the information we just received all within a few seconds. The rate at which the capacities of information technology have increased has grown so rapidly that it is just simply beyond the comprehension of the average person. An example will help put this into perspective. Imagine a twenty-volume encyclopedia on your bookshelf. They take up an entire shelf and together the twenty volumes are much more than the average person could lift. There are approximately 20 million words in these twenty volumes of the encyclopedia[1]. In 1980 it would have taken 232 days to download the encyclopedia. In 1990 that time would be reduced to a little over a minute and a half. By the year 2000 that time would drop even further to a couple of seconds. If you needed something from the encyclopedia and it took nearly eight months to download, you would not even consider it as an option. However, at a download time of a few seconds, you would easily consider several downloads a day to get the most up to date information.

With these enormous capacities at our disposal, our traditional information processing tasks are easily handled and we have to come up with new ways in which to use all this capacity. And we do think of new ways. We have electronic mail, web sites, accessing bank records, paying bills on line, chat rooms, teleconferencing and on and on. But, at some point, you have to sit back and ask, “Where are we going with all this technology?” And that question can be a real stumper. Ultimately we would like to make life better somehow, but that insight doesn’t really get us very far.

[1] 20 volumes x 1000 pages per volume x 1000 words per page.

A Story

Imagine a town in the middle of the dessert hundreds of miles from its nearest neighbor. In this arid land, the town gets its entire water supply by taping into an underground spring that flows under the town. The tap is a pipe about the size of a thick pencil and all the water that the town uses has to come from that little pipe. As you can well imagine, this shortage of water would not only make the resource precious, but it would affect the behaviors and the values of the town. There would be no showers, no lawn watering, nor any backyard gardens. There would be no car washes, or swimming pools; and water balloon fights would be unheard of. The values of the town would be conservative and conservation of the precious water would be at the head of the list of shared values.

If you were to go to this town and ask the inhabitants what they would do if they had more water, you would get some very conservative suggestions if any at all. One person might want a few more ounces to drink each day. Another person may want to bathe a couple more times per year. But, in general, the ideas that the towns people would come up with would not include imaginative leaps such as a water theme park. They would be restricted to ideas that would make tomorrow just a little bit better than today.

The point here is that having a scarce resource affects behaviors and values. And when people are asked what they would do if the resource were not so scarce, their imaginations are limited by their experiences and what they think is possible.

Now imagine that a dam is built many, many miles away with the end result that much more water is now diverted to the underground stream. In fact, due to an increase in water volume, the stream comes to the surface. Now, where the tap used to be, flows a steam yards across and several feet deep. It is not hard to imagine that the increase in available water will affect both the behaviors and the values in the town. There would be more water for bathing and cleaning, which may increase social interactions. There would be water for gardens and perhaps car washing, which would give rise to a demand for new

products. And the value of conservation might be replaced by an increased value of innovation as people come up with new ways to use the water.

Let's assume next that the dam caused some fundamental shift in the way that water from the river flowed into its various tributaries. The net result was that the stream flowing through the middle of town turned into a river a full kilometer wide. The days when water was in short supply are now a distant memory. The innovative uses for water now seem somewhat pedestrian. Instead of thinking in terms of gardens and car washes, the entrepreneurs of the town are thinking of swimming pools, a lake resort, and possibly even a water theme park. And there are new problems being faced. There are drownings, and floods and leaky basements that the people of the town have to contend with.

The point here is that the scarcity or abundance of a resource can have a major impact on the behaviors and values of a society. In the case above, it was the availability of water. In our case it is the availability of information. In the last few decades we have experienced an increase in information bandwidth comparable to the increased supply of water in the town above. At first, there was information to meet basic needs. Then there was ample information and we had to figure out what to do with it. Now, with the rise of the World Wide Web, we are facing a problem of information overload.

Sometimes we get so caught up in this raging torrent of information that we don't even question where it is all going. But ultimately we need to take charge of this rapidly evolving information technology and ask where should we be going with all this development. When you ask people – What should we be doing with all this information technology? They will often respond that this technology makes it easier to manage information. Why do you want to manage information more effectively, you might ask next. At this point most people will shrug and walk away. A few more reflective individuals might suggest that the point of all this technology is to improve our quality of life. They will then relax feeling that they have satisfactorily answered your question. They will relax until you ask the next question – What is quality of life and how does information technology improve it. At this point, even they will walk away.

Most people would agree that improving quality of life is a noble goal for information technology and many people actually attempt to pursue this goal. But if you do not know what quality of life is, how do you know if you are improving it? This chapter will analyze the concept of quality of life and explore ways in which information technology may help to improve it.

What Is Quality of Life?

The few people, who think beyond the productivity gains and greater access to information that we achieve using information technology, see improving the overall quality of life as a laudable goal. But, how to improve quality of life is a little more perplexing than one might imagine. We always get back to the question – What is quality of life? – which needs to be addressed before we can go any further with the issue of how information technology can improve quality of life. Surrogate problems abound. Surrogate problems are problems that one addresses when they feel the real problem cannot be addressed. For example, one way to improve quality of life for a limited group of people would be to use information technology to aid the physically handicapped. This is certainly an important goal, but it does little to address the over all quality of life issue. Other surrogate responses include providing computers or access to computers for those in lower socioeconomic groups. Providing universal access to the Internet is a version of this response. In the past there were allegations that computers were dehumanizing workers. If this is true then computers should be taken away rather than given to people. How do you know what to do? Does the World Wide Web contribute positively to quality of life or does it detract from it. On one hand you have access to enormous amounts of information and a wide variety of products. On the other hand you may feel overwhelmed by all that information and all those choices. You may feel what ever you have in your life is insignificant compared to all that is out there and in that way the Web might just make you miserable.

Examining quality of life from the perspective of concept analysis is still somewhat problematic. If you ask somebody what would improve their quality of life you might get a wide variety of answers. One person might say – "At the moment a cold beer might do a great deal to improve my quality of life," and indeed, at that given moment and for a few succeeding moments that may very well be true. However, it is hard to see how that beer may result in any lasting improvements in the person's quality of life. Another person might say that having a satellite dish on their roof so they could access all possible television channels might improve their quality of life. And, again, it might well provide some improvements in the short term. Once they get bored of all those channels they may no longer see any net improvement in their quality of life. Somebody else might say that a better job would improve their quality of life. Certainly, a new job that paid more money and provided greater satisfaction would seem to be a lasting improvement in quality of life. But what if that job had an extremely high stress level that more than offset the additional money and greater satisfaction?

What if that wonderful new job required a move to a terrible new town? We sometimes see rankings of different cities in terms of quality of life indicators that tell us which cities offer the best quality of life and which offer the worse. But this raises further questions. Is it better to have a great job in a city with terrible quality of life or a terrible job in a city with great quality of life? Or is everybody in a city with a great quality of life happier than the people who live in a city with poor quality of life? All of these questions reflect the fact that we are unsure about what quality of life really entails. Certainly everybody wants to have great quality of life and if technology can improve quality of life then we certainly should do that. But until we understand what quality of life really means it is difficult to do anything to improve it.

Objective Quality of Life

To begin unraveling this tricky problem we need to go right to heart of the quality of life concept. Quality of life can be viewed from two perspectives: sociological and psychological. It is this distinction that we often find at the heart of much of the

confusion. Sociological quality of life is based on objectively measurable factors that sociologists use to compare one city to another. Those objectively measurable factors are things that one might think would contribute to the quality of life in a city. For example, cleanliness of the air and water, cost of basic necessities, commute time, or availability of recreational activities might go into the computation. The benefit of using objectively measurable factors is that cities can be evaluated by any number of different researchers and, presumably, the same quality of life index would be derived. So the quality of life in the city does not rely on the subjective opinions of a particular individual. Further, these indices can be compared over time to determine if the quality of life in a given city is changing. And indices can be compared from one city to the next to determine which city offers the best quality of life. Ultimately, the cities can be ranked based upon these indices and that is how we get rankings like the ten best cities in which to live.

There are two problems with using objectively measurable factors. The first is that they must focus on measurable quantities and the second is that while they may measure the quality of life that the city has to offer, they do not measure the quality of life of the individuals in that city. Things that affect one's quality of life such as stress levels or overall feeling of well-being are very difficult to measure. For example, you cannot go up to a person and ask, on a given day, how would they rate their stress level on a one to ten scale. Each person may apply the scale differently. One person may feel that anything over a six is stressful while another person might rate the same feeling of stress as a nine. Even the same person might not be consistent in rating their stress levels. On two different days, in which they were experiencing the same levels of stress they might assign a six one day and a seven the other. So, in order to eliminate the subjectivity and interpersonal variations in subjective scales, researchers choose objective scales that are arguably related to the factor they are trying to measure. After all, it seems reasonable that the longer one sits in traffic the higher their stress levels, so commute time can be an objective component of stress. Similarly, the greater the percentage of one's take home pay that goes to basic necessities, the more stressed they might be from paycheck to paycheck. This percentage is objectively measurable and thus can be another objective

component contributing to stress. This operationalization of quality of life is far from perfect but in order to obtain an objective measure this is usually what needs to be done.

The second problem is that a city may offer great quality of life and the residents of the city may be completely miserable. For example, assume that a city offers a great quality of life: clean air and water, inexpensive living, great recreation, etc. But the main industry of the town closed down for some reason leaving a large percentage of the population unemployed and the city in a deep unending recession. It would be hard to imagine how the residents of this town would be enjoying a great quality of life.

For another example, imagine a town that is upwind from an industry that produces foul smelling air that is extremely repugnant but not harmful. As long as the wind blow from the town to the industry, the town enjoys terrific quality of life. But a couple of day each month the wind shifts making the town nearly impossible to live in. In fact, assume that the wind shifts at random intervals so the townspeople never know when they will be trapped indoors nor would they know for how long. Quality of life indices may not catch anomalies such as this and indeed a town or city may rank highly in quality of life indicators when the residents may be truly miserable. So objective measures of quality of life do not always capture the phenomenon that they purport to capture.

Subjective Quality of Life

Subjective quality of life is a qualitative individual assessment of the general positive or negative quality of the sum of that individual's life experiences. It can be thought of a psychological quality of life as opposed to sociological quality of life because it is a state that is experienced by an individual rather than an a measurable external phenomenon. Socrates called this state eudaimonia, which translates roughly as good demons or perhaps good spirits. It has something to do with feeling good. Perhaps it is better to simply think of it as happiness. Aristotle believed that happiness is the only goal that is an end in itself. Of all the things that a person may desire - wealth, social status, and political power - people only desire them for the happiness they bring. Happiness,

according to Aristotle, is the ultimate good and the purpose of being. Sociological quality of life fails to provide a meaningful framework for understanding individual quality of life because it measures those concepts that are assumed to contribute to happiness, not happiness itself.

The reason that this difficulty exists is quite simple. In order to measure a phenomenon an operational definition must be developed. This operational definition must have certain characteristics, specifically reliability and validity. Reliability means that must it be objectively measurable and repeated measurements should produce the same results. Validity means that the phenomenon being measured actually represents the phenomenon being researched. Often times validity will give way to reliability for purely pragmatic reasons. This is simply because the construction of an objectively measurable concept nearly always moves the researcher away from the concept being studied.

In the case of quality of life, the objective measurements are generally valid within a limited range. If a person has a little more money and a little more free time he will probably be a little happier. In the work place, if a person has a slightly more challenging job and a little more say in how things are done, she will also be a little happier. However, these improvements cannot be pushed too far. If employees have a lot of extra money and nothing to do they may just get themselves into trouble. In the work place a job that is too challenging puts too much strain on the employee. Having too much voice in how things are done may give a person excessive and unwanted responsibility. So, the objective view of quality of life, whether in the workplace or in society in general, is valid only within a limited range of improvement. However, the validity of this concept does not hold up when computer technology is involved. Computers introduce the problem of scale in that they allow things to happen on a previously unprecedented scale. Since current measures of quality of life are valid within a limited range they cannot be used to evaluate the impact of computer technology. Like the desert town that is now suffering leaky basements, excessive improvements or increases just move everything off of the previous scales. Either new objective measures must be developed, or we must more fully understand the subjective experience. Since objective measures of the future

impacts of a phenomenon are even more elusive that the current subjective factors, this leaves subjective quality of life as the only realistic basic for evaluating the impact of computer technology.

Detractors versus Contributors

There is an old story about a guy who kept hitting himself on the head with a hammer because it felt so good when he stopped. Indeed there are many things that detract from our happiness or our quality of life and removing or eliminating them will increase our level of happiness. Robert M. Pirsig the author of Lila, a popular, but nonetheless excellent book on metaphysics and moral philosophy, observed the following; "Any person of any philosophic persuasion who sits on a hot stove will verify without any intellectual argument whatsoever that he is in an undeniably low-quality situation." (pg. 74) Indeed, hoping off of that hot stove will, most likely, improve that person's quality of life. But once he has left the hot stove, the next step to achieving the next incremental improvement in quality of life may not be quite so obvious.

We understand a lot about detractors from quality of life and we often define quality of life in terms of the removal of detractors. But what happens when all detractors are removed. Are there no further opportunities for improving quality of life? Is not being hit on the head with a hammer the best we have to look forward to? Certainly not! The removal of detractors is often where we focus when looking for ways to improve our quality of life, but once those detractors are removed we must look at things that positively contribute to an improved quality of life. And it is those contributors that we must ultimately identify as appropriate goals for information technology.

Another point that must be made before meeting quality of life head on is that many pleasures which appear to be contributors to quality of life are really just a disguised form of detractor removal. For example, having a beer or a glass of wine before dinner may be a very pleasant experience. And one might think that the luxury of being able to have a drink before dinner improves their quality of life. However, the drink before dinner

simply removes some of the accumulated effect of detractors and does not actually improve the person's quality of life. When the effects of the drink wear off they are right back where they started.

Flow: The Psychology of Optimal Experience

Mihaly Csikzentmihayli from the Psychology Department at the University of Chicago has developed a concept of happiness, which he refers to technically as optimal experience and informally as "Flow". Flow is a state in which people experience an immediate sense of happiness and, at the same time, develop their capacity for sustained long-term happiness. This section will describe the psychological concept of flow, show how it fits in with some philosophical concepts, and discuss some of its applications to other areas of quality of life.

> "Whenever the goal is to improve the quality of life, the flow theory can point the way. It has inspired the creation of experimental school curricula, the training of business executives, the design of leisure products and services. Flow is being used to generate ideas and practices in clinical psychotherapy, the rehabilitation of juvenile delinquents, the organization or activities in old people's homes, the design of museum exhibits, and occupational therapy with the handicapped." [Csikzentmihayli, 1990, pg. 5]

Csikzentmihayli has done extensive empirical research to show that what makes people happy is fairly consistent whether a person is at work or at home, and independent of culture, gender, economic class or any other significant attributes. People are happiest when they are pursuing challenging activities that bring them closer to their goals in life and leave them better suited for the challenges they face. When pursuing these kinds of activities, respondents reported being totally focused on the task at hand and frequently loosing track of the passage of time. They often reported the experience as being "in the

flow" and hence the term "Flow" was adopted to refer to this state of optimal experience. The goals could be either personal or professional, and examples of activities include rock climbing and learning a musical instrument as well as a variety of work related skills such as computer programming. These tasks, from rock climbing to computer programming all have eight characteristics in common.

> "First, the experience occurs when we confront tasks we have a chance of completing. Second, we must be able to concentrate on what we are doing. Third and fourth, the concentration is usually possible because the task undertaken has clear goals and provides immediate feedback. Fifth, one acts with a deep but effortless involvement that removes from awareness the worries and frustrations of everyday life. Sixth, enjoyable experiences allow people to exercise a sense of control over their actions. Seventh, concern for the self disappears, yet paradoxically the sense of self emerges stronger after the flow experience is over. Finally, the sense of the duration of time is altered; hours pass in minutes, and minutes can stretch out to seem like hours. The combination of all these elements causes a sense of deep enjoyment that is so rewarding people feel that expending a great deal of energy is worthwhile simply to be able to feel it." [Csikzentmihayli, 1990,pg. 49]

Most people would agree with these findings based on their own experience. However, Csikzentmihayli offers some theoretical insight into the reasons why flow activities produce happiness. First, the key to optimal experience is control over the content of one's consciousness. Second, optimal experience occurs when one pursues activities that result in a high degree of order in one's consciousness. And third, these activities must contribute to one's goals in life. That is, the order in one's consciousness must contribute to being better adapted to the challenges that one faces.

These principles are consistent with some philosophical views on happiness. Bertrand Russell defined happiness in terms of things that contribute to happiness - work, family

and affection - and things that detract from happiness - boredom, envy, and fatigue. His recommendation was to focus on contributors and avoid detractors, in essence agreeing that one must exert some control over one's experiences in order to be happy. This is consistent with the concept of flow. Experiences that contribute to the order and complexity of conscious energy contribute to a person's happiness. Experiences that introduce randomness or unpleasant distractions contribute to unhappiness.

The second tenet of optimal experience is that it occurs when there is a high degree of order in one's consciousness. Csikzentmihayli writes,

> "Following a flow experience, the organization of the self is more complex than it had been before. It is by becoming increasingly complex that the self might be said to grow. Complexity is the result of two broad psychological processes: differentiation and integration. Differentiation implies a movement toward uniqueness, toward separating oneself from others. Integration refers to its opposite: a union with other people, with ideas and entities beyond the self. A complex self is one that succeeds in combining these opposite tendencies." [Csikzentmihayli, 1990, pg. 41]

Optimal experiences result in personal growth. This is consistent with Aristotle's view of happiness. He believed that happiness was achieved as one pursued the development of personal excellences or virtues. Aristotle also used the term ‘eudemonia’ for happiness. This term is sometimes translated as prosperity, again suggesting the connection between growth and happiness. However, just getting better at something is not enough. You must get better at something you want to or need to be better at. This leads to the third principle of optimal experience. "This ordered consciousness comes about when "psychic energy - or attention - is invested in realistic goals, and when skills match the opportunities for action." This concept of goal directed growth is consistent with the observations of Ludwig Wittgenstein; a student of Russell’s who said that one of the principles of happiness is that you must have a purpose.

So we are generally happier when we are performing tasks that are challenging and thus allowing us to grow and develop. This growth and development makes us become better at doing something that we want to or need to be better at and in doing so creates a more complex self. This more complex self absorbs our conscious energy in and orderly and stable manner thus produces more long-term happiness.

To understand this concept a little better it is useful to understand a distinction between pleasure and enjoyment. Pleasure is hedonistic and does not produce happiness beyond the immediate moment. Pleasure may help us avoid the detractors that cause unhappiness, but it does not produce the stability we associate with happiness. The pursue of one's goals results in personal growth and development. This growth focuses one's conscious energy on the creation of a more developed self and hence produces more stability and a more lasting happiness. Further, enjoyment does not have to occur only in free time activities. Work can be a tremendous source of enjoyment.

Thus, in order to improve the quality of life, people must have opportunities to pursue tasks that challenge them, and provide them with opportunities for personal growth toward a set of personal goals. At the same time they should not be bombarded with random bits of meaningless information that cause disorder in their conscious energy and potentially cause them to unnecessarily lose sight of their goals. Pleasurable activities may distract people from unhappiness, but fail to produce more complex selves. Hence, purely pleasurable activities do not contribute to long-term happiness.

Television and The Quality of Life

A recent work on television and the quality of life by Csizszentmihalyi and his colleague R. Kubey used the flow concept to evaluate the impact of television and provides a model for evaluating information technology. Television is a much more pervasive medium than information technology, but as the information superhighway penetrates further into the homes of individuals this may change.

Television, through both programming and advertising creates goals that cannot be pursued. Instead of encouraging and helping people to identify their own personal goals based on their individual strengths and weaknesses, it encourages people to live vicariously through the actors and actresses. Their goals become to live through others who are physically attractive and who have problems that are all satisfactorily solved in a relatively brief period of time. If we believe the advertising, our goals are further distorted to believe that happiness results from the acquisition of material goods rather from the development of individual talents that make one better suited for functioning effectively in the world. So television interferes with our pursuit of quality of life, first, by interfering with our goal structure and then by providing goals which are not conducive to providing optimal experiences.

A second problem with television is that it provides pleasure without enjoyment. That is, a person viewing television is pleasurably distracted from the causes of unhappiness for some period of time, but this distraction does not result in a more complex self. Hence, any happiness produced by television exists only when the set is turned on. When the set is turned off, the viewer is no more equipped to deal with the world than they were when they started their viewing activity.

> "Just as a drug that masks pain but does not heal may be of limited value in the long run or may cause addiction, so can viewing encourage a false sense of well-being in some people who might be better off taking active steps to change the conditions of their 'real' lives." [Kubey & Csikzentmihayli, 1990, pg. 190]

Many people would argue that they enjoy television immensely and that having a T.V. in every home improves sociological quality of life. Yet, when evaluated from the perspective of its contribution to flow, television has a very dubious contribution to the quality of life. One might argue that there are television programs that do indeed enrich

people and by doing so improve their quality of life. Programs on public broadcasting channels and some cable channels that are dedicated to quality viewing, do indeed achieve this. However, the people who view these programs are probably those who are least in need of personal development.

> "The irony is that television may benefit most those who least need it. People who are already reasonably happy and in control of their lives will be more inclined to find useful information on television and will be less inclined to become dependent on the medium. Those who are less happy and less able or skilled in creating order in their experience are more likely to become dependent, and yet derive less enjoyment from their viewing." [Kubey & Csikzentmihayli, 1990, pg. 187]

Thus the concept of flow provides a useful means of evaluating the impact of technology on the quality of life that can usefully be applied to computers and information technology in the same way it has been applied to television.

Computers and the Quality of Life

In this section the concept of flow will be used to evaluate the impact of computers on the quality of life at three levels: social, organizational and user computer interfaces. At the social level we will examine the potential impact of the information superhighway as a medium, which will compete with (and possible merge with) television. This analysis will follow the reasoning used earlier to evaluate the impact of television. At the organizational level we will examine the impact of information technology on the quality of work life. In this analysis, the concept of flow will be applied to the design of organizational information systems and provide some design guidelines that will improve the quality of work life. Finally, at the user computer interface level we will see how the flow concept can be used to provide improved user interfaces.

Social Impacts

Computers and information technologies detract from the quality of life to the extent that they provide a means of distracting the individual with seemingly random pieces of information that preclude one from organizing his or her conscious energy into a more stable and complex patterns of adapted behavior. They also detract from happiness to the extent to which they force individuals to question their goals and thus interfere with growth activities that may be directed toward those goals. They contribute to happiness to the extent that they provide opportunities for personal growth and development and provide the individual with greater stability and greater control over his or her environment.

Thus, at a social level, the impact of computers on the quality of life can be assessed in terms of three factors: 1) how information technology impacts personal goals, 2) the extent to which information technology contributes to order or disorder in conscious energy, and 3) the extent to which information technology provides opportunities for personal growth and development.

Organization Impacts

Organizational information systems should be designed to support clearly articulated organizational goals. Workers using these systems should be able to perform tasks that clearly contribute to those goals. Further, the system should be designed in such a way that workers can develop greater competence in their ability to contribute to those goals and be rewarded for there increased contributions.

Information systems in the organization detract from the quality of work life when they create unclear goals or when they provide random bits of information that the users of the systems must resolve. For example, "technology for the sake of technology" detracts from flow in an organization. Technology is not a useful goal because, by itself, technology contributes little to the quality of life. Further, with the rapid pace of

technological evolution technological currency is a moving target, which results in a constantly changing goal structure. The effects of the pursuit of technology for the sake of technology have become so pervasive that they have even been given a name “techno stress”.

> "Technostress is a modern disease of adaptation caused by an inability to cope with the new computer technologies in a healthy manner. It manifests itself in two distinct but related ways: in the struggle to accept computer technology, and in the more specialized form of over identification with computer technology. The primary symptoms of those who are ambivalent, reluctant, or fearful of computers is anxiety." [Brod pg. 16]

The struggle to accept computer technology results from the fluid goal structure it creates and from the seemingly random information that one is bombarded with as new technologies evolve. The second component of technostress is a little more complex. It was mentioned earlier that computer programming is a flow activity. It has clear goals. It provides immediate feedback. It is challenging. And one can get better at it in incremental steps. Video games are a similar phenomenon. Since these technologies provide goals, it is easy for a person to over identify with these activities and become oblivious to higher, more long term goals. Thus, while these have the potential to be flow producing activities, the pursuit of them may have a longer term detriment to the quality of one's life.

Flow in Human Computer Interfaces

At the user interface level, information systems should provide meaningful challenges to users that allow them to develop greater proficiency at the work that they do. One of Murphy's Laws of Technology states that if you build a system that any dummy can use, only a dummy will want to use it. Simple user interfaces are great for first time or casual users. However, more sophisticated users become bored and frustrated. A user interface

designed with the concept of flow in mind will provide progressive and stepwise levels of difficulty that correspond to a user's level of sophistication and adaptation to the task being performed. It will allow users to get better at their jobs as they learn more sophisticated capabilities of the system. Thus, user interfaces cannot be designed or evaluated outside of the goal structure that defines their usage. They can only be evaluated with respect to the goals they support and the extent to which they allow users experience meaningful challenges and stepwise improvements toward achieving those goals.

Conclusions

Examining the impact of computers and information technology on the quality of life from the perspective of optimal experience provides new insights into current problems at all levels of computer technology. Computers contribute to the quality of life to the extent that they help individuals establish meaningful goals and provide a means for personal growth and development toward the achievement of those goals. They detract from the quality of life to the extent that they interfere with the development of personal goals and preclude users from becoming better adapted for the challenges that they face in their lives. Further, computer and information technology can detract from the quality of one's life by bombarding the individual with random pieces of information that preclude him or her from organizing his or her conscious energy into highly organized and stable patterns.

Thus in order to improve the quality of life computer technology should be developed with purposes in mind. These purposes should be grounded in basic human needs such as the need to survive and feel secure, the need to belong to a social group, and the need to differentiate oneself as an individual. They should provide challenging opportunities for users of the technology to make progress toward these goals. Ultimately, the concept of optimal experience may provide a non-technical answer to the perpetually vexing question – just what is all this technology good for?

Brod, C. (1984) TechnoStress: The Human Cost of the Computer Revolution. Addison Wesley Publishing.

Csikszentmihalyi, M. (1991) Flow: The Psychology of Optimal Experience. Harper Perennial.

Kubey, R. and Csikszentmihalyi, M. (1990) Television and the Quality of Life: How Viewing Shapes Everyday Experience. Lea.

Chapter 6: Can Machines Really Think?

He who seeks intelligence lacks intelligence. Friedrich Nietzsche

The question of machine intelligence has been with us since the very early days of computers. In the Smithsonian National Museum of American History there is an exhibit called the Information Age in which there is a dramatization of Alan Turning discussing the issue of whether or not machines could ever be intelligent. At the time this was considered to be such a realistic possibility that there were concerns about machines taking over and subjugating humans. Hubert Dreyfus took much of the wind out of the sails of proponents of machine intelligence by saying that machines could not develop intelligence because they lacked physical experience. This was a serious conceptual blow to the project of artificial intelligence, but not fatal. Interest in this issue rises and falls as new technologies provide new promises and eventually new disappointments. In the 1980's it appeared as though expert systems would save the artificial intelligence project. In the 1990's it was neural nets. Currently interest in machine intelligence is at a low due to some typical set backs and the issue of machine intelligence has receded to the background again. However, it will likely come back again as new technologies are developed. In fact, it is likely that we will come back to this topic time and time again.

This chapter will address the question - Can machines really think? In keeping with the approach that we are using in the book we must address the problem by first asking what is thinking and what is intelligence. Once we have answered that we can ask if that is something that machines can do. Finally, in order to understand how we got into this discussion in the first place we have to examine why anybody would pose these questions to begin with. But before we get to far ahead of ourselves we need to consider the magnitude of the problem.

An Example for Perspective

From the time that the first human was born until the modern day there have been less than 20 billion people on the planet, and each of these people lived less than an average of 25 years. Both of these estimates are generous but for an outside estimate we can say that the human race has contributed less (probably substantially less) than 500 billion person years to our presence on earth.

The average person is capable of approximately one addition instruction per second if you use addition or counting as the activity to gauge this. This assumption can be questioned on many grounds. For example, a person can recognize a face in one second and that seems to represent a much more complicated mental activity than simply adding. However, it could also be argued that anything a person could do in one second could be hard coded on a chip that performs the activity in a single instruction cycle. So that question is not as damaging as it may seem at first. To keep things simple, and for purposes of this example, we will place human cognition at one instruction per second.

A person, who lives to be twenty-five, will live less than 1 billion seconds. This number has been rounded up substantially just to make the arithmetic a little easier. So the aggregate life spans of all the humans that have ever been on earth is less than 500 billion billion, or 5×10^{20} seconds. This in turn says that the cognitive capacity of the entire human race, thus far, is 5×10^{20} instructions. This estimate is very generous in the number of people who have lived and the average life span over all of human history. It is also very generous in that it assumes one instruction per second sixty seconds per minute, sixty minutes per hour, twenty four hours per day, and such. Clearly people must rest and sleep and this assumes that they could just keep counting. It is conservative in the estimate of one instruction per second, but if one assumes ten, one hundred or one thousand instructions per second that results can easily be adjusted accordingly.

Now let's take a look at computers. Assume, for the sake of argument, that there are currently 100 million computers in the world [this is an old estimate but does not affect the point] each with a processing capacity of ten million instructions per second. Further

assume that within a decade or two those numbers will increase to 1 billion computers with an average processing capacity of fifty million instructions per seconds.

Using the first estimate of 100 million computers each producing ten million instructions, the computers of the world provide 1000 million million or 10^{15} instructions per second. At that rate it would take the computers of the world 5 x 10^5 or 500,000 seconds to duplicate the cognitive capacity of the human race. 500,000 seconds is just under six days. So the computers of the world at this estimate can reproduce the cognitive capacity of the human race every six days. In a decade or two, using the assumptions above, the computers of the world would produce 5 X 10^{16} instructions every second. At this rate it would take 10,000 seconds for the computers to duplicate the cognitive capacity of the entire human race. Ten thousand seconds is less than three hours, which is to say that within a decade the computers of the world will duplicate the entire cognitive capacity of the human race every three hours. Are we in trouble or what?

What is Machine Intelligence?

If machines become intelligent then we humans are in deep weeds because our rather primitive intellectual capabilities are no match for intelligent machines. Machines are faster, much faster! They are focused. And they don't need sleep, rest, or even coffee breaks. So the question of whether or not machine intelligence is possible is certainly a question of great importance. In the past we have seen some overzealous advocates claiming that machine intelligence would be achieved by the end of the last century. That, of course, did not happen but it reflects the optimism that was feeding the artificial intelligence movement. Other claims have been made that machine intelligence will never be possible, thus representing the polar opposite. How could people vary so widely on this issue? The answer lies in how you define intelligence.

Intelligence in humans is not a particularly well-defined concept. Certain IQ tests provide an operationalized view of intelligence, which allows us to measure individuals and then order them according to intelligence. However, most people would agree that while IQ

test do measure certain factors that are considered to be components of intelligence, they do not measure intelligence in the full sense of the word. There are many instances of what most people would agree is intelligent behavior that is simply not captured in IQ tests. (This is the same problem discussed in objectively measuring quality of life.) But for all their weaknesses, IQ tests are a fairly standard way of gauging intelligence in humans. So does this mean that we can measure machine intelligence by giving IQ tests to machines? Probably not. As poorly defined as the concept of intelligence is in humans, it is even more poorly defined in the case of machines.

The most famous test for machine intelligence is the Turing Test, named after Alan Turing a mathematician who worked with computers in the early days of computing. The Turing Test poses a situation in which there are two rooms, one inhabited by a computer and one inhabited by a person. Another person, the tester, communicates with both the computer and the human via a keyboard. If the tester cannot determine, through any series of questions, which room contains the human and which room contains the computer then the computer can be considered to be intelligent.

There are several problems with the Turing Test as a measure of machine intelligence. The first and most damaging is that the Turing Test does not provide any definition of the concept of intelligence. It only shows how to determine whether or not it occurs in a given machine. Second, it is unlikely that the test would produce the same results in a reliable fashion. For example, would different testers arrive at the same conclusions? Would a machine be determined to be intelligence against all possible human subjects? If the tester were to conduct the test repeatedly would the machine be evaluated the same every time or would the tester become more clever at ferreting out the machine over repeated trials?

This all points to the fact that the concept of intelligence is simply too poorly understood and too poorly defined for any claims to be made at all about machine intelligence. However, claims and counter claims have been made and it is useful to examine some of these positions to see where some merit may lie before just tossing out the whole concept

on the basis of a lack of conceptual clarity. The following sections will examine some of these claims and counter claims. After wading through some of these arguments we will revisit the concept of intelligence and see where the future of machine intelligence is likely to lie.

Fallacies in the Arguments for Machine Intelligence

There are so many fallacies in the arguments for machine intelligence that it is a little unfair to point them all out. Nonetheless, the following sections will examine a multitude of arguments relating to the possibility of artificial intelligence. However, it is important to remember that just because the optimism behind artificial intelligence is based on fallacious arguments does not mean that it is not achievable. It only means that the current arguments do not support it. So after we have put aside the faulty thinking about machine intelligence we will see what can be salvaged form the concept and see where the future may actually lie.

The most damaging fallacy in the concept of machine intelligence is the conceptual weakness described above. If you do not know what intelligence is then how will you ever know if a machine possesses it? However, putting that aside we can see a vast array of other instances of faulty reasoning.

The Galileo Fallacy

The Galileo Fallacy goes something like this. They laughed at Galileo who made scientific advances that were ahead of his time. And they are laughing at me. Hence, I must be making scientific advances that are ahead of my time. The Galileo Fallacy is used frequently to provide plausibility for new ideas that do not fit neatly within current views of the world. The claim is that in Galileo's day, the prevailing worldview was that the planets all orbited around the earth. Galileo challenged this idea, was scoffed at and eventually imprisoned. In hindsight we know that Galileo was right and we use him as an example of modern science overcoming prevailing beliefs in spite of the fact that the

scientific results were unpopular at the time. This fallacy is used to make people feel like they might be scoffing at then next major breakthrough in our understanding of the world.

The Galileo Fallacy is important because people should remain open minded to new ideas even if they are uncomfortable with them. And the idea of intelligent machines is certainly an idea that many people are uncomfortable with. However, there are two major flaws in the Galileo Fallacy.

The first is that there is quite a bit of spin in the story of Galileo as a harbinger of modern scientific thinking. One could even say it is a bit of public relations for modern science. The story of Galileo as scientific hero does not reflect the full situation of what was known or believed at the time nor does it reflect the extent to which Galileo created some of his own problems. It is certainly a romanticized version.

But even if the story of Galileo is completely true and representative of the situation, the technique of using this story still contains a rather large fallacy in its reasoning. To restate the essence of the claim - They laughed at Galileo who made scientific advances that were ahead of his time. And they are laughing at me. Hence, I must be making scientific advances that are ahead of my time. Does this mean that everybody who is laughed at is making scientific breakthroughs? Of course not. If this were the case, every standup comedian would be a future Galileo. For every Galileo there are thousands, if not tens of thousands of crackpots with completely faulty ideas who also get laughed at. If somebody gets laughed at, the odds are very much against them and to suggest that they are more likely to be a Galileo than a crackpot simply does not hold up. If somebody gets laughed at for having bizarre ideas the odds are overwhelmingly in favor of the position that they are, indeed, a crackpot.

The Extension Fallacy

The extension fallacy points to advances that have been made so far, in a field of research and development, and suggests that further advances are just an extension of what has been done. In the case of computers the extension fallacy points to the amazing things that computers have been able to do so far and suggests that intelligent behavior is just an extension of the progress made so far. This is plausible - particularly if you allow extension far into the future - until you look at the extension fallacy applied to other technologies. Actually, any number of technologies could be used to demonstrate this but let us consider the progress from bicycles to motorcycles. Early bicycles could probably achieve speeds limited to a few miles per hour. Modern motorcycles can achieve speeds well in excess of a hundred miles per hours. Using the extension fallacy we might conclude that motorcycles will eventually fly or possibly go fast enough to fly into orbit around the earth or even fly through space.

Of course it requires a completely different technology to fly or travel through space and even though two wheeled vehicles have made impressive increases in speed, one cannot extend these increases indefinitely. Similarly with computers, processors have made impressive achievements in adding, comparing, and moving data. But to go from data processing to intelligent behavior is a little like the motorcycle going form impressive ground speeds to space travel. The extension simply does not hold.

The Reaction Fallacy

The reaction fallacy is, in some ways, the flip side of the extension fallacy in that it suggests that a past event is an indicator of future events whether that relationship exists or not. For example, if one never believed that a computer could accurately predict an election result but found out that the computer did just that, then their skepticism would be called into question. So when they are skeptical about machine intelligence it would look like just another example of skeptical conservatism. The argument goes like this – You did not believe that computers could do task A and they did; Now you do not believe they can do task B, so they probably can since you were wrong before. However, the fallacy in this reasoning can be seen in the following example. A person may have been

surprised when telephony allowed them to talk to a person many, many miles away. However, if told that they could talk to the dead via a telephone they would likely be skeptical again. The reaction fallacy would say you were skeptical about talking over large distances and you were wrong. Now you are skeptical about talking to the dead. Hence you must be wrong again. Clearly, talking to the dead is quite a different problem from talking over large distances. And machine intelligence is quite a different problem than statistical prediction.

In fairness it should be mentioned that these fallacies are used to support the possibility of machine intelligence. Knocking down these fallacies simply says that they do not provide logical justification of the concept of machine intelligence. Which is to say that nobody has shown that machine can think. Neither has anybody shown that machines cannot think. So the question is still open.

Arguments Against Machine Intelligence

The arguments against machine intelligence seem to hold up a little better than the argument for machine intelligence. The following sections will present a variety of claims that have been made suggesting why machines cannot achieve the status of intelligence, at least in the fullness of the word as we use it to describe humans.

The Simulation Argument

The Simulation Argument suggests that while a machine might be able to simulate intelligent behavior that is not the same as actually being intelligent. This suggests that, at best, machine intelligence may just be a cleverly constructed façade. After all, the Turing Test only tests the appearance of intelligence. It does not test intelligence directly. If a machine can produce convincing intelligent behavior then it can pass the Turing Test and may be considered intelligence. But appearing intelligence and being intelligent are not necessarily the same thing. If a person was to be hypnotized and under hypnotic suggestions they demonstrated intelligence behavior would we still say that they are

behaving intelligently. Perhaps not. When a person demonstrates intelligent behavior there is usually some sort of an internal subjective experience associated with that activity. That internal subjective experience has something to do with being goal directed, self-willed, and well adapted to the problem being faced. The hypnotized person probably would not feel that internal subjective experience and the machine would certainly not. This issue of machines not being able to feel will be addressed later in the biological argument. But for now we can say that that appearing intelligent and being intelligent are not the same thing.

The Moral Compass Argument

The Moral Compass Argument uses your moral intuition to create an unpleasant moral dilemma if you allow the possibility of intelligent machines. It is a variation of reductio ad absurdum in which a premise (machines can be intelligent) leads to an absurd conclusion (machines should have the same rights as human).

Consider the following situation. A person is walking a dog across a street and a truck is speeding out of control towards the pair. Reacting quickly you can save one, but only one, of the pair so you have to decide whether to save the dog and let the person get run over or save the person and sacrifice the dog. It is an unpleasant situation either way, but most people would choose to save the person. The justifications for this are many but most reduce to the claim that humans are a higher form of life. When pressed to justify this position most people would cite consciousness or self-awareness or something like that.

Next consider a similar situation in which a person is wheeling a computer across the street. The truck is barreling towards the two and you can only save one. For most people this dilemma is much easier than the first. You would save the person and let the computer be smashed to bits. Most likely you would not even loose any sleep over it.

But how would this situation change if the computer in question were intelligent. In fact, how would it change it the computer were intelligent in the full sense in which people are intelligent. The computer would be aware of its impending death. It would feel pain from the accident and react to the pain. It may survive in a disabled state, perhaps in constant pain. Is the decision still so easy?

Creating machines that are intelligent in the full sense in which people are intelligent allows for possibilities such as this. But it gets much worse. No only would we be creating intelligent entities that could suffer the way humans do, but these intelligent entities might not even have the escape of death. Machines have replaceable parts and consequently could be immoral. Immortality could be pretty miserable if an entity can suffer.

It doesn't take too long to get to some pretty dire consequences of machine intelligence and before long one is inclined to admit that when we talk about machine intelligence we don't really mean intelligent in the same way that people are intelligent.

The Logical Argument

If the moral dilemmas above aren't enough, the Logical Argument suggests that the assumption of machine intelligence leads to logical paradoxes, which are equally as bothersome. Years ago, the mathematician Bertrand Russell introduced a paradoxical situation aimed at the foundations of set theory which became known as Russell's Paradox. In simple terms the paradox can be understood using the following example. There is a town in which the barber shaves everybody who does not shave himself. So far so good, it seems pretty straightforward. Then we ask the question – Who shaves the barber? Now we are in trouble. If the barber shaves himself then he does not shave himself. If the barber does shave himself, then he does not shave himself. It is a logical paradox. This paradox is brought on by the fact that the barber is a member of the set upon which the barber is acting. Crossing this level creates a self-reference problem that leads to a paradox.

A basic expert system contains data, rules, and a control structure. The control structure contains the strategies employed in interpreting the rules. If alternate strategies are available, an additional control structure is needed to select between the strategy control structures. If the system is to learn, then a control structure is needed to control the adaptation of lower level control structures. And so on. Control structures can be thought of as levels or types. If a system contains a finite number of levels then the proof of its lack of intelligence can be found by posing problems that would require a meta control structure one level above its highest level control structure. This problem is often joked about by citing an expert system that can diagnose cancer but can't comprehend a joke or a threat. Thus if a system is to be intelligence it cannot have a finite number of levels in its control structure. Since an infinite number of levels is clearly not possible the only other alternative is a self-referencing control structure at some level. But now we have Russell's paradox to contend with. Consider an adaptive control structure, which maintains all control structures, which do not maintain themselves. Who maintains the adaptive control structure? Gotcha!

We don't really know how humans handle the problem of self-reference but they apparently do and they seem to do it quite well. Russell's Paradox suggests that creating a logical structure to handle self-reference won't work. It is hard to imagine a machine that could see the humor in Abbot and Costello's comedy routine "Who's on First". In this routine, it is posited that the baseball player's have funny names. The first baseman's name is Who, the second baseman is What and the third baseman is named I don't know. These implausible names lead to a hysterically funny dialog in which the names are constantly being confused with their other possible meanings. For example, when Costello asks "Who's on first," he is asking for the first baseman's name. Abbot interprets this as a claim regarding the first baseman's name and responds, "That's right."

Crossing these levels of meaning would be difficult for an intelligent machine, and according to Russell's paradox, ultimately impossible. But the simplest human being can immediately see the humor in this. How can that be? One possible explanation is that the

person listening to the comedy routine does not follow all the logical twists but can relate to the frustration and confusion that Costello is experiencing because they have experienced something similar themselves. But this says that the routine is funny not because of its logical structure. Rather it is funny because of the experiences it recalls. In order to see the humor in this a machine would have to be able to relate to the experience which suggests in turn that a machine would have to have experience. This leads us to the next, and probably the most damaging argument against machine intelligence.

The Biological Argument

We are conscious beings. We are aware of others around us. We are aware of ourselves. We are aware of being aware. We are aware of our awareness of awareness. And so on. Notice that self-reference across levels is not a problem in human intelligence. But where does our intelligence come from? Clearly we have some sort of innate learning ability, but our intelligence appears to come from our interactions with the world around us. These interactions produce experiences and these experiences are evaluated in terms of whether or not they were fruitful and productive and favorable or unfavorable. Our goal in learning is to produce more productive experiences the next time. Sometimes we abstract principles from these experiences that guide future behaviors, again hopefully in a more productive way. But the key to intelligent behavior seems to lie in the desire to produce favorable experiences and avoid unfavorable ones. If this is the case, then a key component in intelligence is bodily experience.

Computer hardware running a software program may be able to produce an external behavior that appears to be intelligent as was argued in the simulation argument, but it is not truly intelligent behavior. Perhaps the reason for this is that the machine is not driven to perform in an intelligence manner by bodily experiences. It is simply running software. Certainly sensors could be added to simulate sensory behavior, but the task of formalizing the responses to a nearly infinite variety of potential experiences would be daunting.

John Searle from Stanford University provides an example that makes this issue very concrete. In this example, a person is in a room with a list of rules for manipulating Chinese symbols in such a way as to appear to understand Chinese. Symbols are passed into the room, a rulebook in English is consulted, and the resulting Chinese symbols are passed back out of the room. To an outsider, the person in the room would appear to understand Chinese and would appear to be producing intelligent behavior. But the person does not understand Chinese and the behavior only appears to be intelligent when in reality the person is just laboriously following instructions. The person may very well be intelligent, but their behavior in the Chinese room does not demonstrate this.

The reason for this is that the person in the room is just processing symbols without understanding them. They see a symbol, look up an instruction, and then select another symbol. The symbols do not connect with anything in their bodily experience. And the decision of what symbols to choose is driven by instructions, not by anything in their bodily experience. It is this lack of connection between the behaviors that they are producing and their bodily experience that makes the behavior appear to lack true intelligence.

In the same way a computer that receives a string of symbols and returns a string of symbols is not producing truly intelligent behavior regardless of how intelligent it may seem externally. The computer is not making any connection between the symbols and its bodily experience because the computer does not have any bodily experience. Nor, as far as we understand computer technology today, will it ever have bodily experience. The reason for this is that the person in the room is just processing symbols and those symbols have no interpretation in terms of his or her experience. Experience and the emotional interpretation of experience play a larger role in intelligence than we realize.

Why do I bother to learn anything? Is it because it makes me feel good? Or more secure in my job, or more integrated in my environment? Perhaps I faced a threatening situation once that I could have avoided by being smarter. Maybe I was embarrassed

once for having stated with confidence something that was blatantly false. Even the pursuit of knowledge for its own sake, results in a feeling of well-being.

Regardless how formally intelligent we can make a computer; it can never really 'feel' anything. If it cannot feel, it cannot develop a notion of self-awareness and self-interest. Without these driving forces it will have to rely on its human partner and will merely carry out its instructions as an extension of the person.

The Sociological Argument

The development of a truly autonomous machine of any kind is a radical departure from the traditional role of tool and machines. Most inventions and new technologies, in some way, extend the inherent capabilities of people. It is easy to see, at a conceptual level, that a club is an extension of the hand. With a little imagination we can extend this idea to see that telescopes, microscopes and possibly even television are extensions of the eyes, while radio, telephones, and microphones are extensions of the ears. Less obvious are the notions that all inventions in transportation are extensions of the feet while weaponry and tools (in the traditional sense) are extensions of the hands. How far can this concept be carried? Crossing from the physical to the metaphysical we can easily see that the computer is an extension of the mind. Marshall McLuhan takes the concept to its completion by stating that modern technology is, in effect, an extension of consciousness. The key word here is extension, not replacement. Technology extends human capabilities.

We delegate cognitive functions to computers but do not provide the computer with its own cognitive functions. Through delegation of these functions we can extend our own intellects. This relationship is consistent with the history of technology. The development of a fully autonomous machine is unprecedented and represents a divergence from the sociological tradition.

The Iconic Argument

There is so much momentum behind the drive to create intelligent machines that one has to ask: Why would we strive for such a goal if it were unattainable? This goes against our intuition. We cannot accept failure as a possible outcome because we have seen so much success in the short history of science and technology and we cannot accept the fact that legitimate scientists in the twentieth century would be in pursuit of a golden fleece. In answer to the latter, the case of the alchemists searching for a way to turn base metals into gold seems very similar to the pursuit of true artificial intelligence. You can't turn lead into gold unless you go down to the molecular level and actually make gold. Similarly you can't make a machine intelligent unless you actually construct a human being. Both are seductive goals, which violate basic laws of nature.

Why then, one might ask, do we strive for this unattainable goal? Is it possible that we just want to imitate ourselves? There are examples from Pygmalion, to Rabbi Loeb's Golem, to Mary Shelly's Frankenstein in which we try over and over again to create something in our image.

It is our narcissistic urge that drives us to search for the golden fleece of artificial intelligence. It is natural, no more or less than the greed that drove the alchemists in search of the secret formula that would turn lead into gold. It is our urge to look ourselves in the face and validate our existence. It is our desire for opposition; the need to believe we are not alone in the universe; and the desire to know other sentient beings (albeit artificial). Like a subject in a sensory deprivation experiment who fills the vacuous reality with hallucinations, we attempt to fill the void with our own creations to give us meaning and put us in touch with an all to fleeting reality. Unfortunately, I think we are searching in the wrong place for validation and purpose.

Evaluating the Pros and Cons

The arguments in favor of artificial intelligence are essentially plausibility arguments and are, generally, rather weak. None demonstrate convincingly that machine intelligence is

possible; they merely show that it cannot be ruled out automatically. Of the arguments against artificial intelligence, the sociological and iconic arguments are of roughly the same status. The iconic argument explains why anybody would believe in artificial intelligence even though it is probably not possible. It is a counter to the plausibility argument. The sociological argument essentially claims that machine intelligence is not in keeping with the way in which we understand technology. We haven't done anything like this yet so we are not likely to do it in the future. This is related to the extension argument, which says we have come so far we can go further. Both attempt to use past behaviors to extrapolate future behaviors.

The moral compass argument and the logical argument are related. The moral compass argument uses emotional reasoning to suggest that machine intelligence leads to some fairly serious moral dilemmas. The logical argument uses logical reasoning to suggest that machine intelligence leads to some fairly serious logical paradoxes. Neither of them is fully destructive. The moral compass argument says that with our moral framework of today, machine intelligence would lead to serious problems. It says nothing about potential moral frameworks in the future. Similarly the logical argument says that with our current understanding of logical and the technologies we use to produce machine intelligence we are not likely to be successful. It does not address possibilities in the future for other logical frameworks or other technologies.

The simulation argument is pretty hard to refute. It may be possible to produce behavior in a machine that simulates intelligent behavior, but to have the machine actually be intelligent is much more than an extrapolation of today's technology. In fact there is nothing on the horizon that even suggests that the transition from simulation to reality is even remotely possible. This is made even clearer by the biological argument that states bodily experience is necessary for intelligent behavior. This is by far the most damaging argument. Subjective experience is not particularly well understood in humans. Producing subjective experience in machines seems extremely remote.

What is Intelligence?

We still cannot dismiss machine intelligence without clearly defining what we mean by intelligence. Consider the following dialog between a machine intelligence advocate and a machine intelligence skeptic, modeled after a famous dialog in Joseph Heller's Catch-22.

Advocate: Machines can think!

Skeptic: No they can't!

Advocate: Do you know what thinking is?

Skeptic: No.

Advocate: Then how do you know that machines can't think?

So there we are. Unless we can get a clearer understanding of intelligence, we cannot proceed any further with this discussion, pro or con. And in keeping with the concept analysis approach that we are using, we are ultimately left with the conceptual questions – what is thinking and what is intelligence? Until we can answer those questions, we are probably just wasting our time discussing the issue.

Where Is This Likely to Go?

Avoiding the currently unanswerable questions about the conceptual definitions of thinking and intelligence, we can ask where this technology called machine intelligence is likely to go. This is much easier to address. Computers are very powerful and will continue to get more powerful. They can perform routine well-defined computational tasks at incomprehensible speeds. If we view them in their role as leverage for human intellectual activities there is no foreseeable limit on what they can do. But, if we are

looking to create some super intelligent artificial being that we can turn to an ask how to solve our problems we are probably just dreaming. Our problems are human problems and it will always take humans to solve them.

Dreyfus, H. (1963) What Computers Can't Do: A Critique of Artificial Reason. Harper Colophon Books.

McLuhan, M. (1964) Understanding Media : The Extensions of Man. McGraw-Hill.

Searle, J. (1984) Minds, Brains, and Science. Harvard University Press.

Chapter 7: Information Systems Technology and the Distribution of Social Power

To exercise power costs effort and demands courage. Friedrich Nietzsche

You hear a lot about some individuals being averse to adopting new technologies while others are eager to adopt every new technology that comes along. You hear about people in organizations being resistant to change, when new technologies are introduced, while other organizations seem to employ cutting edge technologies to great advantage. Less prevalent today are claims that information technology dehumanizes people by making their work routine and mechanical. Finally, you hear a lot about a society of information haves and have nots or a technological elite dominating a technological underclass. Are some people really resistant to change while others are attracted to the excitement of the new? Will the future be characterized by those who use technology dominating those who do not? It is likely that some people are more willing to accept new things. It is also likely that people who use technology will have an advantage over those who do not. However, those answers provide little insight into this phenomenon. Perhaps something more subtle is going on. Perhaps information technology affects the distribution of social power, and it is the potential redistribution of social power that is the more subtle underlying phenomenon. Some people fear that redistribution of power and resist new technologies. Some people see benefits in that redistribution and welcome it. And dividing society into two classes where one class benefits from information technology while the other is exploited by it is one possible distribution of social power, but not the only possible outcome. This chapter will examine the concept of social power and how information technology affects its allocation. It will look at ways in which the current balance of power may be affected by information technology and make some recommendations for how this perspective might be useful in deciding what to do with information technology. As we will see later in this chapter, the concept of social power is not very well understood. However, in order to make the following discussion comprehensible we need a working definition of social power. We can view social power as the influence that individuals have to make things happen or not happen. And we can

extend this a bit to include the influence that individuals have over things that happen to them, in which case they can make things happen or to keep them from happening.

A Story to Make a Point

Walter Wizard works for the Ace Appliance Company in customer service. Appliance repairmen call in for help in appliance repair and a customer service representative helps them diagnose the problem or repair the appliance. Much of Ace Appliance's market appeal rests on it reputation for being repairable rather than throw away.

Walter is an expert at diagnosing and repairing Ace products. Ace Appliance wants to build an expert system to leverage Walter's expertise. The plan is to capture this expertise in an expert system and provide copies of the system to all customer service representatives, thus improving the level of customer service across the board.

Walter does not wish to participate in this effort because he feels that his expertise belongs to him and not to the company. Many people, he points out, have been with the company as long as he has and have taken the same training programs and shared the same experiences. Walter feels that his expertise is special because he studied more on his own time and took his work more seriously than others. Ace Appliance feels that they provided training, manuals, and paid him a salary while he developed the expertise and thus they have some claim to it. At the same time they recognize Walter's point and are willing to provide Walter with a substantial financial incentive.

Walter refused the incentive and refused to participate in the development effort. Walter's boss, who has known Walter for over fifteen years, suggested the company was being very generous and if Walter did not cooperate his job may be in jeopardy. Walter still refused.

Ace Appliance's legal council was called in for advice and offered the following possibility. If Walter refuses to provide his expertise to build the system he may not only

lose his job, but may find himself in court facing a suit for stealing intellectual property. Walter suggests that if forced to comply he may not be as diligent at correct diagnosis and repair as he would have otherwise been leaving the status of the resulting system unsure.

Is this story about resistance to change? Is it about intellectual property? Is it about the issue of who owns knowledge? In some aspects it is about all of those questions. But most fundamentally it is about a redistribution of social power. If Walter provides his expertise in order for Ace Appliance to build an expert system, Walter looses a lot of his personal power in the organization. Ace Appliance, on the other hand, becomes more adept at providing customer service so the organization gains power. Independent repairmen who use Ace Appliance's customer service to get advice on repairing appliances also gain power in that they can service their customers more effectively. So their individual power increases. There are many stakeholders in the resolution of this problem. There is Walter Wizard and Ace Appliance. There are other customer service representatives. There are independent repairmen. There are customers and stockholders of Ace Appliance. And there are a variety of individuals such as Walter's manager and other ace repairmen who compete with Walter for top billing. But the entire situation can be analyzed in terms of the redistribution of social power. Who gains power and who looses power? For those who gain power, what are they likely to do with it? And for those who loose power what are they not going to be able to do that they could do before? But before we can do that we need to understand a bit more about the concept of social power.

How Does Information Technology Affect Social Power?

It may not be immediately obvious how information technology can affect the distribution of social power so some examples will help in revealing this phenomenon. The following examples are far from being exhaustive. They are merely intended to provide some instances to support the credibility of the notion that information systems technology can affect the distribution of social power.

Early Adopters vs. Late Adopters

Early adopters, of new information technologies will tend to gain power, or maintain their pre-existing power, as a result of adoption. This is because information technologies normally provide better information, which means better decision-making and greater control. In some cases information technology means automation in which case the adoption of the technology can lead to greater leveraging of resources and greater productivity. There is a risk, however. Early adoption of a poorly understood technology can lead to failure. Resources are expended but the corresponding benefits are not achieved, resulting in an overall loss of power.

Centralization vs. Distribution

The more centralized an information system and its operations, the more likely the power is to be centralized within a single group, rather than scattered across multiple groups and the functions or departments that they serve. This leads to a concentration of power usually at the top of a formal or even informal hierarchy of some kind. Centralized reporting systems, centralized development groups or centralized control of information resources leads to a centralization of social power within an organization. In society at large, centralization of control over information systems such as the World Wide Web leads to concentration of any social power that may be derived from this implementation of information technology.

Competency Enhancement vs. Competency Destruction

People working together in an organization have a wide variety of skills and specialized knowledge that they contribute to the tasks that the organization performs. These competencies can be magnified or reduced by the introduction of information technology. For example, one used to have to go to the purchasing department to find the best price for a desktop computer, now you can go to the World Wide Web and find the best price

for yourself. Similarly, you used to have to go to a travel agent to get the best price for airline tickets. Now you can often find a better price yourself. In both of these cases, information technology has, in effect, eliminated the competencies of specialists and has allowed non-specialists to do the work for themselves. In other cases the introduction of specialized tools such as software development environments has enhanced the competency of certain groups by allowing them to further leverage their capabilities with tools that require even further specialized knowledge.

Content of the System

Once an information system is implemented, it requires that people work within that system's model, based on standardization of names, procedures, codes, and are limited to that model's content. The things in an organization that are named, counted, and tracked are the things that become important. Further, once the content of the system is solidified through implementation, it becomes the reality for the organization and making conceptual changes at a later date becomes very difficult. Hence, content of the information system frequently becomes the workplace reality. Therefore, the design of the system's content can determine power distribution by specifying what is important to the organization, and the act of designing that system is in itself an exertion of power over other members in the organization.

Change versus the Status Quo

Information technology does not always change the distribution of social power. Sometime it reinforces it. For example, in a highly centralized, hierarchical organization the implementation of a centralized reporting system by a centralized group of developers would serve to reinforce the status quo. Alternatively, introduction of local area networks and distributed processing and development would serve to redistribute social power within the organization. So information technology may not change the power structure within an organization. It may enhance it.

What is Social Power?

Bertrand Russell, who we called upon in an earlier chapter, viewed power as the most fundamental concept in social analysis. In his words, "the fundamental concept in social science is power, in the same sense in which energy is the fundamental concept in physics." If this is indeed the case then it makes sense to view the impacts of information technology from a distribution of power perspective, and see how information technology may result in a redistribution of power. But, what is social power? In this case, existing definitions are not as useful as they normally are. Sometimes power is viewed as synonymous with influence or control. Social psychologists and political scientists view the concept of power differently. Frequently political scientists view power distribution as the allocation of power between the government and the governed, or between the governments of competing countries. Social psychologists often view the distribution of power as the power that individuals or groups within society have to pursue their interests. Unfortunately, with the wide variety of thinkers and writers on the concept of power there is little agreement on the essence of the concept. So we have to take a stab at a useful definition. Intuitively, power is the ability to influence one's environment or to avoid being unduly influenced by one's environment. This can be reduced to Russell's definition of power, as the ability to produce desired outcomes. So the distribution of social power is the extent to which specific individuals or groups in an organization or within society can produce outcomes that they would view as desirable. And a redistribution of power means that the ability to influence desired outcomes shifts from one individual to another, from individuals to a group, from a group to individuals, or from one group of individuals to another group of individuals. If social power is redistributed the shift may be from few to many, many to few, or just from one group to another group. The impacts of the distribution or redistribution of social power issue lie on two dimensions. Those who used to hold power are not going to be able to influence outcomes. Hence, we are going to loose something. Is it something we can afford to loose? Is it something that we would like to loose? Or is it something that doesn't matter either way? Those who gain power are going to influence outcomes in a new way. Is that

something we want to have happen? Is it something that we don't want to have happen? Or is it just a different way of doing things that is neither better nor worse? In order to gain a better understanding of these questions we need to dig a little deeper into the concept of power and the options for distribution of power.

Sources of Social Power

In a now classic article written in 1959, social psychologists French and Raven identified five bases of social power: reward, coercive, legitimate, expert, and referent. Reward power, as the name suggests, is power that is derived from a person's ability to reward. The strength of this power is increased with the magnitude of the reward that the person wielding reward power is perceived to control. Coercive power is power that is derived from a person's ability to punish. The strength of coercive power increases with the extremity of the punishment that one can inflict. These two bases of social power are often characterized as the carrot versus the stick view of influencing outcomes. Legitimate power is derived from values which people have internalized regarding one person's right to influence another. For example, in a management hierarchy, the boss has legitimate power over his or her employees because everyone has internalized the validity of the organizational structure. So Smith has legitimate power over Jones by virtue of Smith's position in the organization.

Expert power is derived from the unique knowledge that a person has in a specific area in which specific knowledge is needed and valued. An example of expert power is the power that professionals (such as medical doctors or lawyers) have by virtue of their profession. However, expert power also covers unique technical skills. The strength of the expert power that Jones has over Smith is a function of Smith's perception of the extent to which Smith perceives Jones as being an expert in that area. For example, if Jones is a medical doctor, Jones will have substantial expert power in assertions made regarding the health of Smith. However, if Smith decides to provide Jones with financial advice, the expert power of Jones over Smith would diminish significantly, especially if Jones is in need of financial advice.

Referent power is the ability of Smith to influence Jones by virtue of the fact that Jones identifies with Smith and wishes to emulate Smith. The strength of the referent power of Smith over Jones is a function of the extent to which Jones identifies with Smith and the attraction that Jones has for Smith. If Smith is an upwardly mobile employee who is well liked and has strong values, Jones' attraction to Smith will be strong as will Smith's referent power over Jones. If neither Smith nor Jones has been promoted in the last five years, Jones may identify with Smith but find very little attractive about Smith. Hence, the referent power of Smith over Jones would be limited. In simple terms Smith's referent power over Jones is derived from the fact that Jones desires to be like Smith.

It is difficult to see how reward and coercive power are immediately affected by information technology. There are certainly small ways. For example, in an organization, more productive workers might be rewarded with the latest technology. Or, at home, parents might reward a good report card with access to the Internet. But in these cases, information technology is just viewed as something desirable, which can be used in reward or punishment situations, but probably does not affect the distribution of power in a significant way. Similarly with legitimate and referent power. Information technology may affect who is in positions of legitimate or referent power, but a more subtle examination would be required to see the connection.

Of these five choices, expert power seems to be the area where information technology most affects the distribution of power. You hear about the rise of a technological elite or the technological haves and have-nots. Knowing how to use information technology is beginning to affect the power that people have over the options they have in life. Even people who are experts in areas other than information technology look silly if they sit down at a keyboard and cannot perform the simplest of tasks. But even expert power does not seem to account for the shifts in social power that have been occurring as the result of the increased use of and reliance upon information technology.

Unfortunately, this characterization of social power does not seem to get us very far and thus raises the question – why go down this path? The answer is that concepts can be

characterized or refined in a variety of ways. Some characterizations are more useful than others for specific analyses. We have seen in earlier chapters how some definitions are more useful than other definitions. Here we see that some characterizations or refinements are more useful than others. The French and Raven view of social power may be useful in analyzing relationships between individuals in an organization, but it doesn't seem to help much in understanding how information technology affects the distribution social power. So when we look at the ways in which various writers and researchers have characterized their concepts we must do so with an eye for which characterizations might have the greatest potential for the concept that we are analyzing.

Several years after the original article, Raven added a sixth base of power called information power. This sounds like it should be a good fit for understanding how information technology affects the distribution of social power. But it still needs a little work in order to advance our analysis. Information power, according to Raven, is based on knowing the internal workings of the organization and how to get things done. In this view, secretaries and administrative assistants have a great deal of information power because you frequently have to go to them in order to find out how things are done in an organization.

This still stops short of what we need but can be made workable with a little extension. We can generalize this concept to the social power that is derived from simply knowing things. This in turn can be refined into knowing what and knowing how. Know what is having accurate information about a given situation that can lead to better decisions and better choices. Knowing how is having an understanding of how to accomplish certain things that need to be done in order to achieve desired goals. Knowing what is supported by the information in an information system. Knowing how is supported by the procedures embedded in the software and the operating scenarios provided by the user interface.

So this more general view of information power (knowing what and knowing how) seems like a reasonable source for the social power allocated by information systems technology. Given that we know where the power comes from we have to ask next where should the power go. And for some insight into that question we can look at what social philosophers

have said about the distribution of social power.

Social Philosophy and the Distribution of Social Power

Social philosophy gives us three dimensions of power distribution to be considered for the introduction of information systems: 1) vertical distribution of power, 2) horizontal distribution of power, and 3) the conflict between what is and what ought to be. The vertical distribution looks at how broad or how narrowly power should be distributed. Some social philosophers suggest a shift of social power upward to the hands of a few; others suggest a shift of social power downward to the hands of many. Horizontal redistribution looks at the impact of shifting power from one group to another. What is vs. what ought to be looks at the question of impact on the status quo. The following sections will examine each one of these dimensions and show how they are affected by the introduction of information systems technology.

Vertical Distribution of Social Power

Vertical distribution of social power addresses whether power should be held tightly at the top of the social pyramid by a few rulers or held widely by larger numbers of the populace. In terms of organizations, this is the argument of whether power should be centralized or decentralized, and where it should reside within the organization. Plato, Aristotle, and Marx, as we shall see, disagreed on the appropriate vertical distribution of power, and Machiavelli provides some insight into the trade-offs.

Plato put forth a vision of the ideal government in his book *The Republic*, which can be used as a standard against which to compare other views. In Plato's Republic the populace was divided into three categories. The philosopher king was the single authority and had that authority because, through understanding philosophy, he understood how things should be run. Below the philosopher king were the guardians - a select group of warriors who carried out the wishes of the philosopher king and from whose ranks the next philosopher king

would be chosen. Below the guardians were the artisans - the rank and file citizens who were incapable of understanding philosophy and consequently lived for the pleasures in life.

In this structure we see the three possibilities of power distribution: rule by one, rule by a few, and rule by the many. Aristotle called these options a monarchy, a republic, and democracy. In Plato's view, the philosopher king should rule supported by the guardians, but Aristotle did not agree. Aristotle followed his rule of the mean and viewed both monarchy and democracy as extremes. Aristotle favored rule by an aristocracy because he saw it as the moderate course between two unacceptable extremes.

Plato's guardians were not supposed to own property because the ownership of property, according to Plato, would cause jealousy and unhealthy competition between the guardians. This would eventually lead to disorder and decline of the state. Instead the guardians just shared the available resources. More than two millennia later, Karl Marx would promote Plato's artisans to leadership, and place the same restrictions on them with respect to the ownership of private property. So we see in these three philosophers the three possibilities of the vertical distribution of power: rule by one (Plato), rule by a few (Aristotle) and rule by many (Marx). Of course, we can carry rule by many to rule by everybody (democracy), which Aristotle condemned, interestingly enough, as mob rule.

The vertical distribution of power is still a problem today and in today's organization. Information systems can centralize or distribute power, which is also to say that they can affect the vertical distribution of power. Centralized power in the hands of a few does not necessarily mean the few at the top of the organization. It may simple mean in the hands of a few in the information systems department. Regardless of where the power is concentrated or distributed, the issues are the same.

While Plato, Aristotle and Marx all had different views of the best way to distribute power, Machiavelli gives us some insight into the major trade-off in vertical power distribution. In his book, *The Prince*, Machiavelli provides strategies for the prince to employ to maintain his effectiveness as a ruler. Machiavelli does not question whether this approach was an

appropriate distribution of social power. He simply provided strategies to help maintain princely power. Many later commentators have viewed Machiavelli as a pragmatist rather than an idealist because he provided expedient approaches for the Prince to maintain his power rather than questioning whether the monarchy was the best way to distribute power. However, in later works, Machiavelli stepped back from his support of the monarchy and favored a wider distribution of power. His reason for this is that it is too easy for a prince to become corrupt and hence it is too risky for a political body to place all of its power in a corruptible prince. This transition captures the most important trade-off in vertical power distribution. Closely held power can operate more efficiently. However, there is a risk associated with closely held power - its effective use relies heavily on the capabilities and motive of the people holding the power. Widely held power is much less efficient, but the risks associated with it are much less.

Information systems can affect centralized power or distributed power. If they reinforce centralized power, then the successful functioning of the organization is determined by the quality of those who hold the central power. If they are capable and effective, the organization will benefit. If they were incapable or ineffective, the organization would benefit more from a more distributed power base.

Horizontal Distribution of Power

Prior to Karl Marx, social philosophers limited their analysis to the power that the government had over the governed. Marx restated this as a class struggle between the working class and the capitalists. Russell broadened the concept even further to include the power of one group over another group. Modern organizations demonstrate this kind of power redistribution with power struggles between departments, between labor and management, or even between innovators and protectors of the status quo.

Information systems can redistribute power across groups or across individuals in an organization. For example, the implementation of an accounting information system may reduce the power of the accounting department and increase the power of the information

systems department. It may also cause power shifts across individuals. For example, by embodying expertise in an expert system, an individual's expert power may be reduced. However, the largest effects come from the redistribution of information power, as the know what and know how is reallocated from one individual or group to another.

Is vs. Ought or Change vs. The Status Quo

Social philosophers have always been concerned with the difference between the way things are and the way things ought to be. Plato's Republic was an idealized view of government where a king led by an understanding of philosophy. The king was supported by the guardian class, who went without personal comforts in order to support the king's philosophical vision. Plato viewed this as the quintessentially rational form of government. However, Plato's view was an idealized view, not just a slight improvement over the current form of government. When asked what would happen if the guardians questioned their role in the government, Socrates said you have to lie to them and tell them that it was the natural order. He called it the "noble lie," which was a legitimate untruth used to maintain the ideal order.

This incident in The Republic points out the problems in idealized power structures. Since they are not in place today, we may have trouble justifying them, and if we were to put them into place as happened with Marx's idealized view much later, we may not understand all of the implications well enough to make them work.

Taking "the way things are" as a basis for social structures creates other problems. Aristotle was opposed to Plato's view because it was too idealized. He felt that the current order in the Greek cities was the natural order and that should be the basis of all political structures. However, Aristotle did not believe that women were capable of holding political office and slavery was socially acceptable in his time. Thus, conforming to the way things are does not allow us to solve social problems, and projecting an idealized order creates a risk in successful implementation.

Interpreting this issue in the context of information systems requires us to look at how information systems impact the status quo. In other words, we look at whether power will be shifted or if the information system will simply reinforce existing power structures. Then we must look at what a change from the status quo might mean.

Information systems designers always confront the problem of change vs. the status quo. Should a new system be developed based on the way things are done today, or should a new system be developed based on the way things ought to be ? If the latter, how do we know how things ought to be ? If we build an information system that reinforces the status quo, we reinforce any problems or inequities in the current system. If we build a system based on an ideal, we have to justify the ideal and run the risk of failure if we fail to fully understand the ideal.

Information Systems and the Distribution of Social Power

In terms of the vertical distribution of power, we look at whether power should be held in the hands of many or in the hands of a few. An information system can result in a centralization of power in the hands of a few or can distribute the power among many, depending on its design. It is, therefore, important to understand the implications of each of these power structures. Closely held power is typically more efficient, however it has an element of risk associated with it - all the eggs are in one basket. The effectiveness of the organization then depends on the capabilities of the few who hold the power. Widely held power tends to be less efficient, but diffuses the risk across more parties.

Information systems can also redistribute power horizontally across groups or across individuals. The two most important factors to be considered with respect to this horizontal power redistribution are that a shift may meet with resistance, either by those who fear the loss of power or by those who fear gaining power (hence responsibility) and not knowing how to deal with it. This resistance should not be misinterpreted as a simple resistance to change. Second, one must consider how the shift in power changes the capabilities of the

organization - certain things will now be possible, while others may be not longer possible within the new power structure.

Finally, the development of information systems provides the opportunity to do things differently in an organization and thereby correct problems in the way things are done. However, change is risky because it is not possible to see all the implications of an organizational change, and radical changes are difficult to justify and implement. New systems should introduce only the amount of change in the power structure necessary to correct significant problems but should not introduce excessive change in the power structure and thereby risk the stability of the organization or social unit.

The Open Question

As we have seen, information systems technology affects the social power that is derived from information power. Information power is power that results from knowing what (information) and knowing how (mechanisms for achieving goals). Information technology can affect the vertical distribution of power (centralization versus decentralization). It can also affect the horizontal distribution of power (power shifts from one group to another). Finally, information technology can be used to reinforce the status quo or bring about dramatic changes for the better. Hence, information technology can and does have a major impact on the distribution or redistribution of social power. Ultimately, this leads the question - How should organizations in particular and society in general be structured? This question is as old as philosophy itself and is well beyond the scope of this chapter.

French, J.R.P. Jr, and Raven, B. (1959) "The Bases of Social Power", in Studies in Social Power. Dorwin P. Cartwright (ed) Institute for Social Research. University of Michigan. pp. 150-167.

Raven, B. (1993) The Bases of Power: Origins and Recent Developments. *Journal of Social Issues.* 49(4). pp. 227-251.

Russell, B. (1960) *Power: A New Social Analysis.* London: Unwin Books.

Chapter 8 – Science vs. Technology

Science and technology multiply around us. To an increasing extent they dictate the languages in which we speak and think. Either we use those languages, or we remain mute. J.G. Ballard

We so often use the phrase "science and technology" that we do not see the two as separable at all. In fact we often see technology as some sort of an application of science, the former dependent upon the later for its very existence. When we think of the modern world we view our progress, in many ways, as the progress of science and view technology as a by-product of that scientific progress. Yet it could be argued (and will be argued below) that technology is much more far-reaching than science and that science is dependent upon technology for its very existence. Further, if we view the past few hundred years as the age of science, we might well view the next few hundred years as the age of technology. This chapter will compare these two highly synergistic concepts of science and technology, show how they are really quite different in their essences, and will use that comparison to put the phenomenon of technology into a different perspective.

An Illuminating Example

Imagine a shepherd sitting on a hillside a few thousand years ago, watching over a large flock of sheep that are grazing peacefully in a lush valley. This morning, as the sheep left the corral, the shepherd carefully counted each one. He remembers the number all day because when the sheep return to the corral for the night he must count them again to make sure that all of the sheep are accounted for. It is important to count correctly because if the number returning is less than the number that left in the morning, then something may have happened to one of the sheep and the shepherd must spend time looking for the missing sheep. However, this process is tedious and error prone. If a distraction occurs when the sheep are leaving or returning to the corral, the shepherd may loose count. Any number of events that occur during the day may cause him to loose count. And confusing the count from one day to the next is always very likely. Any of

these mistakes is likely to cause a great deal of extra work and a much longer day. So the shepherd comes up with an ingenious solution.

He collects a large bunch of pebbles and puts them in a pile near the entrance to the corral. In the morning as the sheep leave the corral, he takes one pebble for each sheep and puts it in a second pile. In the evening when the sheep return, he takes one pebble off of the pile for each returning sheep. If the pile is exhausted when all of the sheep have entered the corral, then the shepherd knows that all of the sheep are accounted for. This method has several benefits over the previous approach. First, the pile of pebbles remembers the count at any point. So if the shepherd is distracted by something while counting the sheep, he does not have to remember where he was. When the distraction is over, he can simply return to moving the pebbles from one pile to the next. Second, the pile of pebbles remembers the count for the day and nothing that happens during the day, short of somebody fiddling with the pile, will cause the count to be confused with the count from another day. Third, the cognitive process of counting and remembering is tedious and error prone, while the physical act of moving pebbles requires very little attention or effort on the part of the shepherd.

This pile of stones, it could be argued, represents the first computer. In this case, a person has delegated cognitive functions, which are tedious and error prone, to an external object and in essence has used an external object to leverage cognition. However, whether or not one is willing to view this pile of stones as an early computer, it is certainly an early example of technology. In this example, the shepherd came up with a solution to problem. This solution does not exist in nature. It is something that the shepherd constructed. All of the scientists in the world, studying atoms, compounds, stars or galaxies would never have found this solution. This solution does not exist in the natural world. It is something constructed in the mind of the shepherd to achieve his goal of having a more reliable way in which to keep track of his sheep. And therein lies the difference between science and technology. Science is a collection of methods that are used to produce reliable knowledge about the natural world. Technology is the means by which we construct a solution to a problem or satisfy a set of objectives. But that is a

pretty big step. So we need to step back and build up to that distinction a little more slowly.

Analysis and Design

In the early days of information systems development there was a stage in the development process, preceding the programming phase, called systems analysis and design. These two activities were lumped together as though it was a single step involving a single process. Later these two steps would be separated and it would be recognized that analysis and design are quite different activities requiring different cognitive skills and producing quite different products.

Analysis is the process of decomposing an object or phenomenon in order to understand it better. In the case of software development, systems analysis is an attempt to understand the existing manual system in order to understand what the automated system must do. Design is the process of constructing a solution, usually but not always out of existing components, in order to meet a set of objectives.

These two processes require very different cognitive skills. In order to be a good analyst a person must be able to look at a complex phenomenon and break it down into its constituent parts. In order to be a good designer, a person must be knowledgeable in the techniques or components that are available and know how to combine them in order to solve a problem or serve a purpose. Although it is not always true, most people tend to be better at taking things apart or putting things together, but not both. Hence, the person who is good at analysis and design is somewhat rare.

Analysis always begins with an existing phenomenon and ends up with a better understanding of the existing phenomenon. Design begins with a problem to be solved and ends up with a constructed, but not necessarily implemented, solution.

This distinction is not unique to systems analysis. Other forms of analysis such as market analysis, financial analysis, even psychoanalysis all attempt to understand a complex phenomenon by decomposing it into constituent parts. And, if a problem is discovered in the process of analysis, a design activity is required to come up with a solution. If the problem is small or fairly routine the solution may appear to arise automatically from the analysis, but the larger or more unique the problem the more it will rely on a creative constructive effort for solution.

A non-technical example will help to further illustrate this distinction. Suppose there are two towns, which we will call Appleseed and Citrusville, on opposite sides of Raging River. Both local economies are in trouble. Appleseed produces more apples than it can sell locally keeping the prices low and the apple farmers barely breaking even. When the apple farmers take their apples to Citrusville, they can sell them at a premium but getting the apples across the Raging River is difficult because the river is very wide and the current is very strong. The orange farmers in Citrusville have a similar problem. Excessive local supply holds prices down and exporting organges to Appleseed is difficult because of the river. Both sides believe that if they could export a quarter of their crop to the other side of the river then local prices would rise and they would benefit greatly from the export revenues. The analytical activity has examined the facts of the case and identified the problem. Now the question is what to do about it and this becomes a design activity.

The design activity constructs a solution to the problem usually, but not always, using solution components. While the facts in the analysis are a part of the world at the time of the analysis, the design solution is not part of the world and must be constructed. Perhaps it would be fair to say that the analysis is a process of discovery whereas the design is a process of invention.

The design solution that comes to mind most readily is to build a bridge across the river thus allowing traffic to flow easily between the two cities. This is certainly reasonable to consider but may not be feasible. Building a bridge may be too costly or too difficult for

any number of reasons. So other solutions need to be considered also. One possibility is to build more powerful boats that can cross the river more easily. Maybe a cable across the river would help keep less powerful boats from being carried downstream by the strong current. Maybe it would be possible for small aircraft to carry produce from one side to the other thus avoiding the river entirely.

More creative solutions are possible. Perhaps the farmers can exchange seeds so that the farmers on each side could grow both apples and oranges in the proper proportions to maximize profits without having to export or import at all. All of these design solutions exist as possibilities in the future and no amount of studying the current situation will lead to these solutions without a creative leap of imagination at some point. Further, this distinction between what is true in the current situation and solutions that can be constructed to have things be different is going to be a key concept in understanding the difference between science and technology.

This discussion of analysis and design may seem like a digression from the point of this chapter but it is important in understanding the difference between science and technology. But before we get into all that let's look at some common understandings of science and technology and see what weaknesses there may be in those common understandings. We will then look at science and technology as broader instances of analysis and design, and will, finally, look at some of the relationships between science and technology.

What is Science?

It is not too difficult to come up with examples of sciences. For example, most people would feel pretty comfortable saying that physics, chemistry and biology are sciences. After all, when you have to take science electives in school, physics, chemistry and biology are among the most common choices. So they are pretty clearly sciences. But what about sociology and economics - are they sciences? Well, we call them 'social' sciences. Does that word 'social' modify science in some way so that these fields are not

really sciences? What makes something science? And while we are at it, what about mathematics? Mathematics seems pretty important, but is it a science? When we move outside of the realm of traditional science and into areas like astrology and alchemy we are pretty sure that they are not sciences, but what makes us so sure?

Sometimes we add the word 'science' to a field to make it sound more robust. So areas like management science, computer science, and information science need to remind us that they are sciences by putting the word *science* in their names. But, if that is all it takes, then creation science should be a science also. Some people say that anything that calls itself a science is not. That is a little harsh, but you have to wonder why some fields have to put science in their names lest we believe that they are not sciences.

While we are asking what makes a field a science we can also ask what makes somebody a scientist? Is it the fact that they work in a field that we call a science? Was Aristotle a scientist? Was Isaac Newton a scientist? Most people would sputter over these questions and say something like Aristotle is the father of science and Isaac Newton was the greatest scientist that every lived. But is this true? Well, in order answer all these vexing questions we can go back to the beginning with Aristotle and see how the concept of science has evolved.

The intellectual project that eventually led to modern science began nearly three thousand years ago when early Greek philosophers began to realize that natural phenomenon may have natural rather than supernatural explanations and that by investigating these natural phenomena we might come to understand them. These early philosophers, sometimes called Ionians, sometimes called pre Socratics, dismissed myths that explained natural phenomena in terms of the behaviors of gods and looked to the natural world for explanations.

Aristotle, who came along a couple of centuries later was the first to rigorously, systematically, and for the time exhaustively, explore natural phenomenon. Looking back, we might label him as the first scientist but he would have thought of himself as a

philosopher. One of his areas of interest was natural philosophy, but he wrote on politics, aesthetics, and a wide variety of other topics that we would not consider either natural philosophy nor science.

The giants of early science such as Galileo and Newton would not have called themselves scientists either. They would have referred to themselves as natural philosophers. It wasn't until the middle of the 19th century that the word science came into use and its applications was not as complimentary as it is today. The natural philosophers of the 19th century found increasing more of the work of natural philosophy was technical, detailed data collection and analysis rather than the development of grand ideas. They wanted to distance themselves from the technicians so the term scientist was coined to refer to these more technical people. Over the past century and a half the term scientist has become more prestigious than natural philosopher, but at the time the reverse was the case.

The word *Science* comes from the Latin word *scire* which means to know. Hence, science means knowledge in a broad sense. Scientific knowledge is redundant unless we see science as a very specific kind of knowledge. That very specific kind of knowledge is knowledge that has been tested, verified, and articulated in a form that allows others to re-verify it and build upon it. So when we talk about scientific knowledge we are talking about knowledge that has been rigorously acquired, scrupulously verified, and organized in such a manner that it can be unambiguously shared with other people so that they may verify it for themselves.

Another important feature of science is that scientific knowledge is acquired largely through observation. That is to say that science is an empirical enterprise. The role of theory in science is to organize observations and predict new things that might be observed to further verify the knowledge acquired thus far. Consequently, even though mathematics is frequently referred to as the Queen of the Sciences, it is not really a science. The properties of triangles require no reference to any real triangles. Mathematics is organized knowledge of a different kind.

The project of science has produced impressive results in our attempts to understand the natural world. But is all science natural? What about social science, political science, or management science. These disciplines do not study the natural world. They study the social world. So are they really sciences? Well, if we use the criteria that science is based on observation, then it is not clear what we are observing. Social, political and organizational structures are different from atoms, molecules and living beings in that they do not have any presence in the physical world. But to the extent that observations are rigorous, scrupulously verified, and organized in such a manner that it can be unambiguously shared with other people so that they may verify it for themselves, these disciplines to seem to be sciences.

But some sciences do not meet these criteria. Creation science and Astrology fall short of science in that they cannot be verified. The key to verification is that the claims of the science must hold up under scrutiny. This idea goes all the way back to Socrates whose primary method of acquiring knowledge was refutation. When an idea is repeatedly challenged and yet one cannot refute it, it gains credibility. A very similar idea was offered by Karl Popper in the last century. Scientific statements must be falsifiable. So if a body of knowledge claims to be a science, it must contain statements that can be proven to no be true. Creation science and Astrology have lots of claims, but few if any of them are falsifiable.

So to sum things up, science is a body of knowledge about a specific phenomenon that is reliable because it is pursued via a rigorous methodology that attempts to refute it claims. Claims that hold up over time under repeated attempts at refutation are more likely to be valid in the long term.

What is Technology?

It is not too difficult to come up with examples of technology. We have computers, television sets, video players, cars, home security systems, etc. etc. In science we have telescopes, microscopes, particle accelerators, incubators, etc. etc. When we go to work

we hop in a car, drive over a bridge to a huge building with central air conditioning and heating. These are all technologies. When we come home we may take some food out the refrigerator, put it in a microwave to heat it up and then put the dishes in a dishwasher to clean while we plop down in front of the television set to catch up on news from around the world brought to us by satellite. Technology is so pervasive in our lives that it is difficult to imagine life without it. Even if we need to get away we hop in a car, go through toll booths to get out of the city, sit at red lights that control traffic at our destination and then call home to let others know we arrived safely. But is technology merely this collection of gadgets or is their more to it than that?

Technology comes from two ancient Greek words – *techne* and *logos*. *Techne* is, roughly, a reliable means of achieving an end, and *logos* is, roughly, a body of knowledge that explains a phenomena. Logos is where we get all of our 'ologies' from. Psychology is knowledge about the psyche. Sociology is knowledge about social systems. And technology is knowledge about *techne*.

Any craft is *techne*. So cooking, rhetoric, medicine, drama and engineering are all *techne*. They are all crafts and are all bodies of knowledge that apply reliable methods to producing desired outcomes. So when we think of televisions, computers and cell phones as technology we are shortchanging the concept. Certainly engineering went into creating these produces but to suggest that these engineered products are the full scope of technology fails to recognize the full role of technology. Technology is the science of design. It is a body of knowledge that allows us to construct reliable results or solutions to well defined problems. It is not merely applied science as many would assume. In fact the Merriam Webster dictionary defined technology as applied science. But does that definition give full justice to technology?

What is the Relationship Between Science and Technology

When Galileo made his startling discovery that the sun and not the earth was the center of the solar system, he did so not as the result of new ideas. He did so as a result of new

technology, the telescope to be specific. The telescope allowed Galileo to do two important things. First, it allowed him to make more accurate observations on planetary orbits. And, second, it allowed him to see that the moon was pock marked with craters. The craters on the moon may not seem like such a big deal but Aristotle had claimed that the heavenly spheres were perfect. Seeing the craters on the moon told Galileo that Aristotle was wrong about the perfections of spheres and allowed him to question other ideas that had been around since the time of Aristotle. But, again, it was not scientific ideas that led to this challenge. It was a new technology.

Technology plays a more important role in the advance of science than one might think. Often times, new scientific data is acquired based on new technologies for observation and data collection. This new data may provide new insights or overturn previous ideas about the natural world. Other technologies such as the computer have aided the scientific project by allowing us to organize and process volumes of data that would have been impossible to process by hand. But the use of technology does not stop there. Scientific method is a process for producing reliable knowledge. So scientific method is a form of *techne* and hence a technology. In fact, the more closely we look at it, the more scientific progress appear to be the result of applying technology to the production of reliable knowledge about the natural world. So rather than denigrating technology as merely applied science, one could claim that science is merely applied technology. To push this idea a little further, we could observe that knowledge about the natural world is only one kind of knowledge that we derive from the application of *techne*. For example, drama and literature are techniques for exploring social dynamics and the human condition. These technes also help us to acquire knowledge, yet of a very different kind. So now we have reduced science, not only to the application of technology toward the production of reliable knowledge, but one of many such applied technologies.

The point here is not to demean the scientific enterprise. Rather, it is to take technology out of the shadow of science. Both science and technology are important endeavors and both must be fully appreciated in their own right. Science attempts to address the question – what is true about the world? It is primarily analytical. It does not solve

problem nor provide solutions. It provides reliable knowledge about the world. Technology on the other hand answers the question - What could be? It is primarily synthetic. It is a design activity. It solves problems and changes the world, hopefully, for the better. It provides reliable methods for providing desired results.

Science and technology simply could not get along without each other. Not only are technologies necessary to collect and process data, but the organization of scientific knowledge is again a design activity and hence a technology. On the other hand, how can we know what possibilities are likely to produce results we desire? We have to understand the possibilities and limitations of our tools. Once we have created something we need to study it to see if it does indeed produce the desired results. And the study of artifacts is once again a scientific activity. So science and technology are at once quite separate and hopelessly intertwined. But, until we pull the project of technology out from under the moniker of applied science it would be difficult to appreciate the full range of possibilities.

Implications for Information Systems Technology

We tend to think of information technology or information systems technology as computers, networks, hard drives and the like. And at some level this is true. This is technology that is used in information systems. But this is not the only *techne* in information systems. There are also a range of methodologies that are used to develop information systems and a collection of models used to represent information. When we create an information system we are, in effect, modeling the world in information. We do this because it is easier to interact with information than it is to interact with the world. The information models that we interact with have become so pervasive that we fail to see them anymore as information models. When a sales manager looks at a sales forecast and decides to put a few more dollars into advertising and promotion, he is interacting with an information model of the world. When a person uses software to balance her checkbook and then uses a spreadsheet to go over the monthly budget she is interacting with information models of the world. When somebody goes to a web site to find a book

and then orders the book from the website, he or she is interacting with an information model. It is a bookstore experience modeled in information. So information technology is not only the hardware that we use, but the variety of methods that we employ when modeling the world in information. This perspective is important because when we model the world in information we are not limited to the shortcomings of our world today. We can create a new world that is, perhaps, more orderly, more rational and more effectively meets our needs.

We have learned a lot about our world from the past few centuries of scientific explorations. But, for all we have learned, it has been limited to what is and not what could be. It may be that we are in a transition point from an age of science to an age of technology in which we focus more on possibilities than actualities. Information systems technology is an important part of that age of technology just as physics and chemistry were important parts of the age of science. But, it should also be emphasized that it is only a part and there are many more possibilities that we have yet to dream of.

Chapter 9: What is Teleology and What Does It Have to do with Information Systems Development?

To everything there is a season, and a time to every purpose under the heaven. Ecclesiastes

Traditional information systems development is based on a lifecycle development model in which the requirements for an information system are gathered and documented, and then an information system is designed to meet those requirements. The design is implemented and the system, after testing, is turned over to the users of the system. This all sounds very rational until we ask the question – where do the requirements come from? In order to gain some insight into the difficulty of answering this question, we need to turn to a fictional analogy of information systems development.

The Railroad Paradox

The Railroad Paradox is an anecdote about a railway company introduced by Gerald Weinberg [Weinberg, 1988, Gause and Weinberg, 1989] to illustrate an important point about information systems development. In this story a railroad train stops in Homeville each morning and takes commuters to work in Workville. In the evening the train picks up the commuters in Workville and returns them to Homeville. During the course of the day, the train passes through Homeville on the way to Workville but does not stop. Residents of Homeville would like to catch a mid afternoon train to Workville for shopping and other reasons so they write the railroad company and ask that the train stop on the 2:30 pm run. The railroad company sends representatives to Homeville to observe the station in question at 2:30 pm. When they are there they see no people waiting for the train, they concluded that there is no demand, and declined the request. The point of the story, of course, is that demand for the train does not exist until the train starts making a regular stop, at which time the demand for the train will develop.

It is often the case that demand for a product does not exist until that product becomes available. This is certainly true for information technology. Most major IT products, such

as the microcomputer, electronic spreadsheet, word processor, or personal database, were developed first and then users had to find ways to use them. The availability of the products and the ingenuity of the users in finding ways to use the products created the demand. In most cases the potential users did not understand what the products could be used for when they first came out so there could not have been any significant demand for the products.

This is not just true for IT products. It is also true for traditional information systems. When we develop accounting or financial systems we inquire about the demand for information that may or may not lie in the future. We ask users what information they need (ostensibly to make better decisions) and how they would use that information if they had it. We then attempt to build an information system based on those speculations. More often than not the requirements of the information system will change during development as users become more familiar with the information it may produce and become better at imagining how they might use it.

The point of the Railroad Paradox is that demand for IT products and systems does not exist, typically, until the products and systems exist to create the demand. This runs completely counter to traditional information systems development models and reveals an important flaw in our understanding of information systems. Traditional systems development models begin with requirements analysis, a phase in which the requirements of the system are documented prior to building the system. But, if the requirements do not exist until the system is actually developed, then where do these requirements come from? If the demand for the product does not exist until the product is developed and understood, how can we expect the potential users of a currently nonexistent system to provide any insight into the potential requirements of a software artifact that exists primarily in their future?

Most information systems developers ignore this paradox and go right ahead with requirements analysis as though the requirements for a future artifact exist in the world today waiting to be discovered by the diligent requirements analyst. And it works,

sometimes. As long as the requirements for the future systems are based on limited extrapolations of things that exist in the world today, it is possible for the analysts to put together a list of requirements for the future artifact. So, when a payroll system or accounting system is being automated, this process may work even with its flawed assumptions. But when there is a discontinuity between things that exist in the world today and the things that will exist in the future, this process of looking in today's world for requirements is seriously flawed.

The problem is exaggerated even further by the continuing growth of technological capabilities. As we get improved graphics and user interfaces, better database products and more sophisticated development environments, we find ourselves trying to find ways to use these capabilities and apply them to artifacts that lie in the future, the purpose of which is unknown. But this problem of knowing how to do things but not knowing what to do is not new. In fact, it is as old as technique itself. Although the components are different, it is a problem that was recognized by Plato twenty five centuries ago.

The Dialog Of Gorgias

Plato often showcased the wisdom of Socrates by putting him in opposition to other characters that represented what Plato thought of as lesser virtues or talents. One of those lesser talents was skill in sophistry. Sophists were masters of the art of persuasion and could construct clever arguments in favor of any side of an argument regardless of the real merits of any particular position. One of the most famous sophists was a teacher of rhetoric named Gogias who confronts Socrates in the dialog that bears his name. In this very revealing dialog Socrates faces the sophist Gorgias who is seen in the eyes of the Athenians as a well-respected and highly regarded public speaker. Often Gorgias will give a speech in a public place and the Athenians will shower him with praise and money for his efforts. Socrates does not accept money for his teachings, because his only goal is the pursuit of truth, and the scene is set for a showdown between the two. Socrates begins by asking Gorgias who he is. Gorgias responds that he is a teacher of persuasion. The conflict in the dialog is that Socrates is a seeker of truth. He wishes to understand the true

nature of concepts such as justice and virtue. Gorgias is a teacher of rhetoric. He believes that there is no truth - that you can convince anybody of anything. Socrates asks Gorgias to give him an example in which persuasion has value. Gorgias says that his brother is a physician and often has to convince his patients to undergo unpleasant treatments for the sake of their health. Socrates asks Gorgias, "What do you persuade the patient to do?" Gorgias responds "Whatever my brother wants the patient to do." There is a moment of realization as Gorgias sees that his art is all technique. While Gorgias is a master of rhetoric and can convince anyone of anything, he is lacking in the knowledge of what he should be persuading people of. He is diminished next to Socrates who is attempting to find out how things should be. While the rhetorical techniques of Gorgias are dazzling, they get him nowhere unless somebody else tells him what to do.

Telos Versus Techne

This dialog emphasizes an important distinction between two important concepts: telos (how things should be) and techne or technique (how to make things that way). Telos is the end state, purpose, distant goal or thing off in the distance. It is the root for words such as telescope which is a device for seeing distant things. It is also the root for teleology which is, roughly, the study of end states or purposes. Techne is a craft, methods or means for producing a desired end. It is the root for technique and technology. In the dialog of Gorgias we see a conflict between telos (Socrates' attempts to discover truth) and techne (Gorgias' glib rhetoric).

It is exactly the problem that we have with technology today. There are many Gorgian software developers who can dazzle you with technique. They can make multicolored objects dance around on the screen. They can recite version numbers and acronyms. They can talk about capacities and capabilities or object models and plug-ins. But when you say "What is the point of what are you trying to do?" they balk just as the sophists balked at Socrates. I do not mean to condemn modern technology nor the people who have worked very hard to master it. I merely want to point out that technology is just embodied technique. And if you cannot answer the question - What should you be doing ? -

technique and technology are of quite limited value. The purpose of this chapter is to suggest a teleological approach to information systems development, which changes the focus from the things we could do to the things we should do.

Teleology

Teleological thinking began with Aristotle who believed that an adequate understanding of a phenomenon required an understanding of four causes: formal, material, efficient, and final. The formal cause is the shape that a thing takes on. A boat and a picnic table, for example, can be made out of the same material, but take on very different shapes. The material cause is the stuff out of which the thing is made. A boat may be made out of wood or steel. The efficient cause is the procedure by which the thing is made. A boat can be manufactured, constructed from lumber or carved out of a log. The final cause is the ultimate purpose of the thing. The purpose of a boat is to float on water. It is easy to see that a ship builder who does not understand the final cause of a boat may not be very successful in building them.

The final cause is ultimately a teleological explanation, which Aristotle required of all scientific explanations including inanimate or physical phenomenon. This led to some problematic interpretations such as – fire rises because it wishes to return to the sun, or objects fall because they wish to return to the center of the earth. This attribution of purpose to physical objects does not sit well with one's modern sense and indeed Galileo dismissed teleological explanations from his view of astronomy leading eventually to the modern view that physical objects simply follow the laws of nature and do not have any ultimate purpose.

It is fairly easy and appropriate to dismiss teleology from physics, but that does not necessary dismiss this perspective from all scientific and intellectual endeavors. For example, in biology, there is an ongoing debate regarding teleological explanations of biological systems. Can one fully understand the functioning of the kidney if one does not know that the purpose of the kidney is to remove waste materials from the

bloodstream? It would seem that one could not. However, does that purpose exist in nature or does the observer ascribe that purpose in order to improve our understanding of the functioning of the kidney. Clearly, this debate can be easily distracted by the philosophical question of whether or not purposes exist in nature or if purposes are simply superimposed on nature in order to make things easier to understand. I would gladly accept the second and lessor of these claims by saying that it doesn't matter if purposes exist in nature. Superimposing purposes on nature in order to make things easier to understand is a good enough justification for a teleological perspective.

As we move from biological systems to human systems such as government or economic systems teleological perspectives become increasingly more important. These systems evolved or were constructed to satisfy some sort of human need. To discuss them without discussing the purpose that they serve is to miss the point completely.

The extreme of teleological thinking can be found in engineering where objectives are clearly defined and solutions are constructed to satisfy those objectives. Although it might be theoretically possible to describe a construct such as a bridge over a river in purely physical terms, you probably would not want to drive over a bridge that was built by someone who did not know that the purpose of the bridge was to get cars safely from one side of the river to the other.

One would expect software engineering to follow this model but rarely are development objectives defined for a software development project. Instead requirements are defined which are problematic for three reasons. First, it is not at all clear where requirements come from. Conventional wisdom says to ask the users what they want. But asking users to imagine a system that exists in the future puts an unrealistic strain in their imaginations. Second, as shown in the Railroad Paradox the requirements for the information system may not exist until the system is implemented. Third, requirements imply that a solution has been constructed and the requirements are the requirements of that solution. But where did that solution come from and how does one know that it is the

correct solution? Attempting to develop information systems using a functional rather than teleological approach is fraught with problems.

A Teleological Approach To Information Systems Development

Information systems development should begin with a problem to be solved. If there is no problem, then why would somebody invest all that time and money to build an information system? That problem should be clearly articulated and agreed upon. When the project is completed, its success can be determined to the extent than the problem is solved. Further, upon completion of the project, developers can reflect on how well they solved the problem and see if they can do anything to make their next problem solving effort more effective. This allows developers to improve the development process.

The problem should be decomposed into objectives, which are really just subproblems. These subproblems should be articulated and agreed upon just as the problem statement was. And just as was the case with the problem statement, the developers should reflect on the objectives when the project is finished, and then figure out how to improve their articulation of objectives.

The process of stating a problem and then decomposing it into objectives is really just another decomposition technique. Structured analysis decomposes a system into processes. Information modeling decomposes a system into entities. And object oriented analysis decomposes a system into reusable components. Why is decomposition into objectives any better?

There are several answers to this question. First, decomposition into objectives is a higher level abstraction and thus more general. Once the purpose of the solution system is understood, parts of the system may be analyzed further using any of the above techniques. However, the above techniques are all more specific and consequently don't serve to organize the overall purpose of the system being developed. Second, objectives are easier for people to understand. The average person can grasp the concept of what the

developers are trying to achieve much more easier that they can grasp the implications of a data flow diagram or object model. Finally, and because objectives are easier to understand, it is easier to gain consensus and develop a shared vision of what the system under development is suppose to do when it is complete. This shared vision is critically important because thousands of design decisions, large and small, will be made at various points by developers based on their understanding of the proposed system. If they do not have a shared vision, the resulting system is likely to be a hodgepodge of conflicting design decisions. Finally, it is much less expensive to correct design flaws earlier in the development process. If there is no agreement on what the proposed system is supposed to do, it is better to find that out when working on the problem statement than during development.

If a problem-oriented approach is so beneficial, why doesn't everybody do things this way? The answer is that defining a problem and decomposing it into objectives is difficult to do correctly. It is difficult for three reasons. First, objectives are less concrete than processes, entities, or reusable components. However, as systems analysis techniques have evolved from flowcharts and decision tables to data flows and entities, the progress has been toward higher levels of abstraction in order to control the complexity of large systems. So decomposition into objectives, although difficult, is necessary as we develop increasingly more complex information systems. Second, it takes a fair amount of experience in defining objectives before one gets very good at it. Approaching information systems development from a process improvement perspective allows you to learn from your experience get any better a developing successful information systems. Finally, problems and objectives are amorphous terms to most people. In order to be successful in a problem solving approach we need more structure in the concept and a better understanding of problems and objectives. The remainder of this chapter will address this issue.

Stating The Problem

Some times it is difficult to succinctly articulate the ultimate purpose of an information system. Usually there are multiple purposes and multiple levels. And there are many things that need to be done that are not contributing to an ultimate purpose in any obvious way. Information system requirements are usually very difficult to understand and comprehend in their entirety for just this reason. Since developers cannot reliably build a system that they cannot understand, it makes sense to reduce this complexity to something that they can understand. On this issue, Gause and Weinberg offer the following advice: "A possible solution is to regard every design project as an attempt to solve some problem."

Some developers are skeptical of viewing software development as solving a problem. Frequently, software is developed to exploit an opportunity. However, the difficulty here lies in the common semantics of the word 'problem' and a more precise technical definition. A problem is not necessarily something 'wrong' with the organization. Again, Gause and Weinberg offer some direction: "A problem can be defined as a difference between things as perceived and things as desired." A problem here is a gap between the way things are and the way we would like things to be. Since nobody's life or organization is perfect there are many instances of a gap between the current and the desired state. So every person and every organization has problems that do not necessarily reflect negatively upon them.

But not all problems are solvable problems. If a person thinks that the desired state for their organization is that everybody should be a millionaire then there is a gap that is unlikely to be closed. However, if a person thinks that paychecks should come out on time and be accurate, then it is a gap that it is possible to close. This idea is then embodied in the concept of a 'solvable problem.' A solvable problem is problem in which the desired state can be achieve given the resources available to the problem solver. That is, there is a gap between the current and desired state. It is a gap that we

know how to close given the necessary talent and other resources. And those resources are available.

Now we must take a fairly large step and claim that in a teleological approach to information systems development we should only work on solvable problems. This claim must raise eyebrows and it certainly flies in the face of traditional wisdom that claims we often do not know what we are doing in information systems development for a wide variety of very good reasons. How then would we ever build a new system or develop software that has never been developed before.

There are two important responses. First, I believe that we grossly exaggerate the extent to which software being developed is truly new. It is much more likely that developers who are unaware of what other developers have done are simply solving the same problem again in a different way. In recent years we have seen a proliferation of object models which suggest that there are common solutions to information systems problems and that it make much more sense to learn about and utility common solutions rather than constantly reinvent everything.

Second, if the software to be developed is truly new, then it is a creative process not a software development process. The problem to be solved is the gap between the developers current understanding of the application or technology and the level of understanding that would be required to approach the application from a teleological perspective. Once that problem is solved, the developers can focus on the application problem. Certainly no one would ever board an airplane that was built so that the engineers could learn about aerodynamics. But we seem to think that it is O.K. for an organization to rely on and information that was built so that the developers could learn about the application.

It is certainly fair to observe that information systems typically solve many problems. This may very well be true. But it is also fair to observe that some problems are more important the other problems and some problems are subsumed in larger problems. We

cannot loose cite of the fact that our analytical techniques in information systems are driven by limitations in human cognition. Decomposition techniques, for example, are a result of a recognition that we cannot understand complexity. We have to break it down into pieces and then understand the pieces. In the same way, a multifaceted problem may be a more accurate description of the situation, but if nobody can understand it, then it serves little purpose in the development process. A clearly stated problem may necessarily exclude something that is otherwise deemed important. But it may also increase the likelihood of success in the development process.

Defining Objectives

The weakness in current methods of decomposition such as data flows, entities or objects is that they do not address what the system under construction is supposed to do. They may address the way things are current done, or the way they could be done, but they do not address the way they should be done in order to solve the problem at hand. Using objectives in the system development process provides a means for addressing the purpose of the system in a highly structured and systematic fashion, thus filling the gaps left by traditional methodologies. Historically, the problem with using objectives has been the lack of structure in the process of defining objectives. Objectives are a decomposition of system purpose in the same sense that modules are a decomposition of system function. Viewing objectives in this way makes it possible to distinguish well stated objectives from poorly stated objectives. Further it is possible to identify different kinds of objectives (component, competing, and constraining) and how they relate to each other.

An objective is a subproblem that arises when a problem space is decomposed into (relatively) independent components. An objective is a definition of purposeful activity that not only defines a desired end state but also carries within its statement some understanding of how that end state will or should be achieved. Viewed as a decomposition of the problem space, objectives should have the following characteristics:

1) Objectives should be mutually independent, and each objective should solve one aspect of the problem in its entirety.

2) Objectives should be collectively comprehensive so that if all objectives are met the problem will be solved.

3) Objectives should lend themselves to further hierarchical decomposition so that if an objective is too large to tackle directly, it can be broken into subcomponents that again have the characteristics being stated here.

4) Every objective should state some important aspect of the resulting solution and by implication should indicate something that is not important. Hence, objectives of this type should guide the people solving the problem at every stage as to the goals of the solution they are trying to produce.

In addition, the following aspects of objectives should also be observed:

1) There are three levels of objectives and objectives must match the level of the problem that they are addressing.

2) There are as many different kinds of objectives as there are interests in the solution.

3) Objectives can be related to each other as components, competitors, or constraints. As components they can be organized hierarchically as a means ends analysis of the problem space. As competitors they represent conflicting goals and must be examined in a trade-off analysis. As constraints they place limitations on the possible solutions to the problem under consideration.

4) Some objectives are units of work toward which action will be directed. These actions will be human and organizational actions. Hence, these objectives must be achievable units of work.

Levels of Objectives

One great source of confusion in defining objectives is that there are three levels of objectives each of which are formed differently and measured differently.

Strategic objectives address long term issues, usually five years or more, and aid in the direction of corporate efforts or the repositioning of an organization. Strategic objectives can be stated by managers who are literate with respect to technology but not immersed in the techniques of systems development. Good strategic objectives are important because they provide the visionary framework within which to organize more specific activities. Measurement of strategic objectives is interpretive by nature.

Tactical objectives address the means by which strategic objectives can be achieved. They result in the expenditure of resources toward work efforts. Tactical objectives are stated by managers who understand the functional area within which the work is being done. They are measured by determining if the desired work has been accomplished. Tactical objectives are the most problematic of objectives because they require both the visionary planning capabilities of a generalist planner, and the in-depth functional area expertise of an analyst. Tactical objectives represent the transition between a business problem and a technological solution.

Operational objectives are implementation oriented objectives that determine if measurable improvements have been made. They are the ultimate performance level validation criteria for the chain of objectives. Operational objectives do not require functional area expertise. They can be stated in terms that the user of the functional area services understands. And they can be evaluated by anyone who uses the services.

Kinds of Objectives

Critics who suggest that real-world problems are just to messy to allow for the definition of objectives are probably trying to map all objectives into a single hierarchy. Usually, there are many different kinds of objectives that are not related to each other in a hierarchical fashion. For example, in an information system there are functional and data objectives that address what the system has to do. There are also human/computer interface issues that address ease of use, training time and productivity concerns. However, there may also be process objectives such as "the system will be build in-house by junior programmers" or "the system will employ a new technology such as database or expert systems so that the organization can gain experience in an emerging technology". There may be objectives for performance, quality, maintainability, and flexibility. There may be objectives for distribution of the system, platforms that should be accommodated and turnaround time for repairing system flaws. What is important to draw from this is that many classes of objectives will need to be identified for any systems development effort. Each class of objectives should be made as explicit as possible, and the relationships between the classes should be identified.

Relationships between Objectives: Component, Competing and Constraining Objectives

While the clear definition of individual objectives is very important, the definition of the relationships between objectives is also very important. Objectives can be related to each other in three ways: as components, competitors, or constraints. As components they are related hierarchically as objective to subobjective. If objectives are related hierarchically, then the satisfaction of all subobjectives at a given level should ensure the satisfaction of the objective from which they were derived.

Component objectives can be organized in a hierarchy where the objective at the top of the hierarchy focuses more on the desired end state, while the objectives at the bottom of the hierarchy focus more on units of work needed to achieve the end state. Recall,

however, that within any problem space there may be multiple hierarchies of component objectives.

Competitor objectives represent conflicting goals and must be examined in a trade-off analysis. Pair wise analysis of competitor objectives should determine which of each pair is more important and how much more important. For example, if flexibility is an important objective of an information system and performance is also important, then one must ask "Which is more important, flexibility or performance, and how much more important is it?" Techniques such as the Analytical Hierarchy Process (AHP) can be used to turn pair wise comparisons into a ranking of competing objectives. This ranking is important to designers making trade-off decisions in systems design.

Constraining objectives place limitations on the possible solutions to the systems development problem. The difference between a constraining objective and a simple constraint is that a constraining objective represents a solvable problem whereas a constraint simply represents a static limitation on the problem space. For example, an objective stating that a new system will be built entirely by the in-house systems staff limits the possible solutions. However, the staff may be trained in new methods or new technologies. Or they may be motivated by some sort of incentive program. If these avenues are open then the objective represents a solvable problem. If the new system is to be built entirely in-house and there is no possibility of training or motivating the staff in any way, then this objective is really just a constraint.

Objectives as Units of Work

At some level, objectives become units of work toward which action will be directed. These actions will be human and organizational actions which must be feasible, understandable, motivating, acceptable, and flexible. They must be feasible because people cannot achieve the unachievable. They must also be feasible within the constraints of the available people resources. They must be understandable because people cannot accomplish tasks that they don't understand. They must represent

reasonable challenges so that people are motivated without being overwhelmed. They must be acceptable to the people responsible for achieving them and they must be flexible enough to respond to unforeseen changes in either the application domain or the development team.

A Final Observation

For well-understood problems, a teleological approach is the most effective because it directs resources toward the solution of that problem. However, poorly understood problems should be viewed a learning problems. That is, the first step is to figure out what needs to be done. Once that is figured out, you can approach the development problem.

Gause, D. and Weinberg, G. (1989) Exploring Requirements: Quality Before Design. Dorset House Publishing.

Weinberg, G. (1988) Rethinking Systems Analysis and Design. Dorset House Publishing.

Chapter 10: Thinking Clearly About Concepts

Concepts, like individuals, have their histories and are just as incapable of withstanding the ravages of time as are individuals. But in and through all this they retain a kind of homesickness for the scenes of their childhood. Soren Kierkegaard

It would be nice if all concepts were clearly defined and used in an appropriate manner. It would also be nice if all issues about which we had to make decision or form opinions were presented in a balanced and objective manner. However, both of those goals are unrealistic. Sometimes people simply don't think through what they are saying. Sometimes they have goals or values, explicit or implicit that drive their characterization of things. Often times their goal is to promote a certain worldview, agenda or issue rather than question it. But even if a person is attempting to be balanced and objective they may very well be frustrated in their efforts. Real world issues are often very complicated. If you embrace the full complexity of the issue you will like go far beyond what the average person can comprehend. So you have to simplify. But the act of simplification requires that you ignore some aspects of the issue while promoting others. And this process inevitably introduces some bias. That is to say that one can try to be fair and objective, but that is nearly always unachievable.

Consequently, when one faces an issue such as privacy or quality of life, or any issue that rests on some poorly defined concepts it is important to clarify the concepts and clarify the claims. The technique of concept analysis has been used throughout this book and does not need to be further elaborated. However, there are a few surrounding issues that need to be explored further. Before embarking on a discussion of these issues I would like to introduce a story that I used frequently when opening a discussion of issues in personal privacy.

A Story About Privacy: Is Personal Privacy an Antiquated Concept?

We often make assumptions about how the world ought to be that are so fundamental to our worldview that we cannot imagine a situation in which those assumptions would not

hold. This is the case today with personal privacy. So consider the following dramatization.

The 12th Century

Imagine a minor noble sitting around the dining hall of a medieval European castle finishing off a few after-dinner drinks. The noble takes a long draw on his draught, wipes his mouth on his sleeve and proclaims to his friends,

"You know, I think that we would all be better off if everyone had a say in how they were governed. Each person could have one vote and they could apply that vote to electing their representatives. Those representatives could then do the governing and if they did not suit the needs of the people, then they would not be elected the next time around. This would make the government responsive to the needs of the people and would improve society by making everyone a stakeholder."

A round of belly laughs and guffaws scampers through the room as everyone immediately sees the patent absurdity of this idea.

"The masses are too uneducated to select their own leaders," offers one rather sinister and effete looking noble.

"The masses are too easily swayed by demagoguery," offers another.

"Elections would become popularity contests," offers a third, "where only those who are good looking with honest faces would have a chance."

"Yes, yes," replies the harbinger of change, "all those problems may well occur. But the problems that result from my idea are far outweighed by the benefits, and the difficulties that they create will diminish with time. What I have offered, is the right thing to do and it is only a matter of time before it comes about."

The 21st Century

Now imagine an information systems specialist sitting around a birds-of-a-feather session at a privacy conference. Our specialist is staring out the window lost in thought while the other attendees chant the catechisms of personal privacy. They talk about personal information, shopping history, medical conditions, bad credit reports, and sometimes a chill comes over room when the conversation drifts unexpectedly to that modern day incarnation of evil - the social security number. There is no audible "Amen" after each allegation of violated privacy, but it is certainly implied. They are soldiers in the good war - the war against the invasion of privacy by information technology.

The discussion reaches one of those predictable lulls as our specialist returns from the outer reaches of imagination and focuses once again on the presence and the reality of the small meeting room. He begins to speak in that slow and measured way which suggests that the words are being formed one at a time in his brain and delivered directly to his mouth. He says,

"I think that personal privacy is a thing of the past. Any transaction that occurs and is recorded in cyberspace should be available to anyone who want to see it. Not only do I think that there should be no restrictions on the dissemination of personal information, but I think that society has a responsibility to make it easily accessible to everyone. There should be public workstations in every library that allow anyone to come in off of the street and examine personal information about anyone they wish to know about."

The room became very quiet.

"I want to know what you buy at the grocery store, what video tapes you rent, how much money you make, what your mortgage payments are, and whether or not you have any hideous diseases. I want to know where you buy your gasoline and how much you spend in car repairs. I want to know where you went on vacation and how you got there. I want to know how often you call your mother, and what your kids get on their report cards.

Not only do I want to know, I think it is my right to know. I have a right to orient myself in the world and know things about the people around me. Suppression of information leads to isolation, alienation, supposition, innuendo, and superstition. In short, I think that any that restricts the flow of personal information is immoral."

Although it did not seem possible, the room grew even more quiet.

Nobody wanted to speak. They were not used to opposition. The moral high ground of privacy had never been seriously challenged. To find themselves on the low ground of inhibiting the free flow of information and being pursued by a flash flood of self righteousness was not something they had expected. They were too used to nods and the implied "Amens". Nobody knew how to address the opposition. So they began to raise their familiar old arguments.

"What if somebody looks at your purchases from the grocery store and starts calling you at dinner time to sell products?"

"What if you are denied medical insurance because of a medical condition?"

"What if you are denied a job because of some event in your past?"

"What if you rent adult movies or call phone sex numbers and somebody decides to broadcast that information to everyone on the internet?"

"Yes, yes," replied the harbinger of change, "all those problems may well occur. But the problems that result from my idea are far outweighed by the benefits. And the difficulty that they create will diminish with time. What I have offered, is the right thing to do and it is only a matter of time before it comes about."

Conceptual Realignment

What was attempted, in this story, was a conceptual realignment of the personal privacy issue. Personal privacy was such a sacrosanct issue that nobody could see any downside

to protecting personal privacy. It was, as they say, motherhood and apple pie – something nobody in their right mind would ever challenge or criticize. Any yet there are always two sides to any ethical issue. Ethics always consists of competing interests and the problem with personal privacy as an ethical issue was that everybody seemed to be on one side of the issue. So this story used three conceptual realignment techniques in order to allow a more balanced discussion of the privacy issue. The first was to remove privacy from the moral high ground by recasting it as a form of censorship. Aversion to censorship and a belief in free flowing information is another apple pie issue and putting two apple pie issues in opposition to each other prevents a person from merely standing on unexamined principles. Second, most people, at the time, saw the privacy issues as competing interests between the individual and big government. Hardly anyone is going to stand up for big government against the rights of this individual. So protecting privacy became another issue on which individual rights were being assailed. Recasting the privacy issue as competing interests between one individual (one who wants privacy) and another individual (one who wants to know) removes another moral high ground of protecting the rights of the individual. Third, in equating the currently unaccepted view that not protecting privacy is somehow bad with an earlier view that democracy might be bad shows how people can have strongly held beliefs based on how the world is today, but those beliefs may become outdated as the world changes. This conceptual realignment was, to use a less flattering term, propaganda; or to use a gentler term, persuasion. The piece was written intentionally to affect the views that people might have regarding personal privacy. Concepts can be used to clarify our thinking, confuse our thinking or reorient our thinking. This example with the right to privacy shows how conceptual realignment can help us to understand concepts more clearly. But sometimes, conceptual realignment can be used to confuse or misdirect us.

Misdirection and Conceptual Fallacies

Thinking through complex issues is often complicated by the fact that we rarely encounter them in a balanced presentation. Often advocates of a position have crafted their side or their opinion of how things should be so that it seems the most reasonable. In

fact, if they are good enough at crafting their opinions, it will seem like any position other than the one the one they are proposing is silly. Hence, it is important to see beyond the packaging and context of ideas in order to evaluate their true merit. Since this is a book about concept analysis we will take a deeper look into the privacy issue and see how it has been crafted to seem like the only logical position to take. Following that we will examine some techniques for undoing the packaging. Finally, we will have a few suggestions for being a little more balanced in our thinking.

Metaphors and Implied Analogies

George Lakoff and Mark Johnson in their book Metaphors We Live By have provided important insights into the role of metaphors in our understanding of reality. More than just a literary device, metaphors reveal much about our understanding of concepts. When metaphors are used in discussing concepts, we can examine the metaphorical references in order to understand implied meanings. For example, very often see the Privacy is a Territory to be Protected Metaphor. Phrases such as protecting privacy or invasion of privacy suggest that privacy is somehow an area or territory where we are vulnerable.

Emotional pleas are another powerful technique in conceptual realignment. Privacy advocates often use emotional pleas to get their point across. They try to make a connection in your mind between privacy in your home and privacy in cyberspace. Few people would argue against privacy in your home. Home is where you expect to be out of public view. You expect to be able to let down your guard and if somebody invades the privacy of your home you feel very vulnerable. Not the least of this is due to the fact that your physical presence is at home and if somebody invades your home you could be in physical danger. But having somebody looking at a transaction in cyberspace and having somebody looking in your window at home are not at all the same thing. Yet privacy advocates often connect these two kinds of privacy making you feel at those somebody invading your privacy in cyberspace is a lot like somebody catching you in the shower with a video camera. Wanting privacy in the shower has a strong emotional component to

it and when the connection is made to cyberspace there is a strong emotional desire for privacy in cyberspace.

Role of Stories

Stories can play a large role in influencing your thinking about specific concepts. Stories provide narrative arguments which can often sneak around logical arguments or prejudices and make a person see things differently before they realize that there current conceptual perspectives are being challenged. The story about the 12th century medieval castle earlier in this chapter is a good example. One has already bought into the story before one sees the point of the story. The ethical example in the machine intelligence chapter provides another example of this. One has bought into the ethical quandary of having to save an intelligent machine before one realizes that the contradiction speaks against machine intelligence. Several stories have been used through out this book to make a variety of points. Some have been identified as stories and the point has been analyzed. Others were told along the way and may not have registered at the time as stories. Hence, stories often lull the reader into lowering their defenses so a new point can be made or a new perspective can be provided. Using stories is often useful when a head on attack of an idea is likely to generate defensiveness.

Rhetoric

In Joseph Heller's bestseller, Catch-22, Nazarian claims "They are after me." His friend disputes this claim causing Nazarian to challenge "Do you know who they are?" His friend admits that he doesn't know who 'they' are and Nazarian replies "well, how do you know that they are not after me?" This bit of fancy rhetorical footwork leaves anyone who challenges Nazarian's claim which no further challenge. This is a humorous scene in Catch-22 but is frequently used in less humorous arguments. The argument in favor of machine intelligence is on. An AI advocate claims that machines can be intelligent. You dispute that claim saying that machines cannot be intelligent. The advocates asks you if

you know what intelligence is. You reply that you do not. They retort "How do you know that machines can't be intelligent." This fancy footwork is the stock in trade of rhetoric and only one trained in rhetoric can easily counter it.

Surrogates

Sometimes we replace the concept we are examining with a surrogate concept. We saw this in the chapter on quality of life. If we do not know what quality of life actual is, we can replace it with surrogate concepts such as clean air, clean water, low cost of living, good health care and the like. We often use surrogates when we cannot address concepts directly. But when we are examining concepts in order to understand them surrogates often get in the way. If we studied surrogates in the chapter on computers and the quality of life then we might land up trying to find ways in which computers could be used to improve air quality or water quality instead of directly address the core quality of life issue. Not that improving air and water quality is all that bad. But it missed the core issue.

We see another use of a surrogate concept in the chapter on machine intelligence. The Turing Test for Artificial Intelligence says that if you have a computer in one room and a person in another room with another person taking to both via a keyboard, if the person cannot tell which the person is and which is the computer then the computer is intelligent. This misses the concept of intelligence completely and replaces it with the surrogate concept of trying to fool the interrogator. It is easy to imagine how much research could go into systems that fool the interrogator without bringing us one step closer to machine intelligence.

Thinking Clearly About Concepts

Concepts are general ideas formed from specific instances of object, behaviors, occurrences or examples and as such are a core mechanism for organizing our knowledge. A large part of our ability to think clearly comes from our ability to organize our thoughts clearly. Clarifying the concepts that we use helps us think more clearly.

Following are some rules of thumb for examining the concepts that we use so that we can ensure that they are properly defined and properly used.

Rules of Thumb

In Chapter 3 we looked at some techniques for concept analysis and tried a variety of ways in which we could clarify out thinking about a particular concept. We then employed those techniques in order to explore a variety of concepts. There are some rules of thumb that we can employ in determining if concepts are well defined, but I did not provide them earlier because I did not want to make concept analysis appear too mechanical. However, they are worth mentioning before closing. These rules of thumb have been drawn from information modeling and modified to address concepts.

Membership: The first rule is the membership rule. For each concept there must be a rule that allows us to determine if a particular event or object is a member of that concept. Ideally, a membership rule for Concept X will allow all X and nothing but X.

Consistency: The second rule is the consistency rule. All items included in a concept must have similar features. Internally they must have similar attributes and externally they must play similar roles in the domain of discourse within which the concept is being examined.

Utility: Finally, the third rule is the utility rule. Concepts must have some utility which means that a given concept definition must be useful in some way. That way could be clarity of definition, our ability to gain greater understanding of the concept, or our ability to use the concept in decision making or policy making. We may even wish to define the concept in such a way that it influences people's thinking about the concept.

The Ghost of Socrates

We have conjured up the Ghost of Socrates to explore a number of philosophical issues in information systems such as personal privacy, machine intelligent, quality of life, and the redistribution of social power to name a few. We began with a technique that Socrates used in which he would challenge an idea to see how well the concepts involved hung together in a cohesive and useful manner. We refined this technique in concept analysis and applied it in an attempt to understand key issues in information systems. Finally, in this last chapter we saw how concepts can be manipulated to influence our understanding of important issues. The point of all this is that is important to say what we mean and mean what we say. Concept analysis serves us toward this end. And every now and then we need to conjure up the Ghost of Socrates to help us figure out what we really think about important philosophical issues.

Annotated Bibliography

In the introduction to the bibliography for A World Lit Only By Fire, William Manchester observes,

> "Bibliographies are useful guides for readers who want to learn more, but they can be deceptive. Traditional bibliographical structure is sometimes misleading; the order of the works which are cited is determined by the alphabetical order of the first letter in the scholars' last names. Furthermore, every entry appears as the equal of every other, which is an affront to common sense."

I am inclined to agree with this sentiment and take it even a step further. Many times long bibliographies are used to give credence to the work in which they are cited. Authors cite long lists of works some of which may have had a profound influence on the current work, but most of which just happened to be laying around while the work was being written. I have kept my citations to a minimum in this work because I believe it stands on its own. On the other hand, when I have taken a specific intellectual contribution from another author such as a quote or a specific idea I have given due credit. Nonetheless, there were specific works that influenced my thinking and I think it only fair to give them their due. Consequently the following annotated bibliography is an attempt to give credit to those from whom I got many of my ideas.

Csikszentmihalyi, M. (1991) Flow: The Psychology of Optimal Experience. Harper Perennial.

I find myself referring to this book a lot in a lot of different circumstances. I was impressed that somebody could gain such insight into the elusive phenomenon of happiness and Csikszentmihalyi's characterization of happiness fits in so well with so many related concepts, not the least of which is the ultimate purpose of technology as discussed in the quality of life chapter.

Dreyfus, Hubert L. (1979) What Computers Can't Do: The Limits of Artificial Intelligence. New York: Harper & Row.

Research in artificial intelligence was getting pretty heady in the late 1970's and a lot of critical reasoning was cast aside. Dreyfus came along with a philosophical analysis of the AI project which explained why the project was doomed to fail. For me, it showed that

philosophy has value in information systems because a deeper understanding of things can save one from wasting a lot of time on an ill conceived idea.

McLuhan, Marshall. (1964) Understanding Media: The Extensions of Man, New York: A Mentor Book.

McLuhan is often far more useful for a clever quote, sound byte, or intriguing observation than he is for a clear explanation of anything. Nonetheless, I found his concept of technology as an extension of the person very useful in my attempts of understand the role of technology. This tied in with Simon's idea of sciences of the artificial and help lead to my teleological view of information systems.

Palmer, D. [1988] *Looking at Philosophy: The Unbearable Heaviness of Philosophy Made Lighter*. Mayfield Publishing Company.

Palmer, D. [1991] *Does the Center Hold ? An Introduction to Western Philosophy*. Mayfield Publishing Company.

When I began teaching philosophical issues in information systems, my biggest hurdle was finding something for students to read that would give them some background in philosophy without totally overwhelming them. These two books by Palmer are great accessible entry points. If any one were to tell me that they knew nothing at all about philosophy but would like to learn something, I would point them towards these two book by Palmer as a starting point.

Pirsig, Robert M. (1991) Lila: An Inquiry into Morals. Bantam Books.

Pirsig is perhaps best known for his other book Zen and the Art of Motorcycle Maintenance. But I found Lila to be a much better book. Pirsig is far from being a mainstream philosopher. Yet I found this book to be one of the most poignant and revealing of all the works I have read on metaphysics and moral philosophy. In many ways is loosened up my philosophical thinking and allowed me to reevaluate many of my, perhaps erroneous, assumptions.

Russell, Bertrand (1960) Power: A New Social Analysis.

Resistance to change is a much batted about phrase in organizational studies. Many people feel that information systems fail because of resistance to change. This book by Russell introduced me to the concept of power and a means for analyzing social situations and led to much of my thinking in the redistribution of social power.

Russell, B. (1958) The Conquest of Happiness. New York:
Liveright.

Russell did not make as much progress as Csikszentmihalyi in understanding happiness. Yet it is a good idea to have a collection of great thinkers who you respect and look to

their opinions when approach a new idea. In fairness, I would have to toss in Plato and Aristotle into this collection. I always check and see what they said about any philosophical issues that I am interested in. Even if I don't fully agree with their views, their insights are invaluable.

Searle, John. (1984) Minds, Brains and Science. Harvard University Press.

This book contains Searle's famous Chinese Room analogy which helped clarify, through a very simple example, the difference between the appearance of thinking and the actual experience of thinking.

Simon, Herbert (1984) The Sciences of the Artificial

This book is a frustrating hodgepodge of ideas. And yet it was the first place when I encountered the idea of artifacts as an area of study. The idea that natural science studies what is out there, while sciences of the artificial studies the means by which we construct things led to my teleological view of information systems and the distinctions between telos and techne.