The Philosophy of Information:
Ethical, Ontological and Epistemic Perspectives.

Thesis submitted in fulfillment of the requirements for the degree Doctor of Philosophy

By

Steven Tyrrell McKinlay

August 2013

Thesis advisors

Prof. John Weckert and Dr. Morgan Luck

Charles Sturt University, NSW, Australia

Contents

Certificate of Authorship

I hereby declare that this submission is my own work and that, to the best of my knowledge and belief, it contains no material previously published or written by another person nor material which to a substantial extent has been accepted for the award of any other degree or diploma at Charles Sturt University or any other educational institution, except where acknowledgement is made in the thesis. Any contribution made to the research by colleagues with whom I have worked with at Charles Sturt University or elsewhere during my candidature is fully acknowledged.

I agree that this thesis be accessible for the purpose of study and research in accordance with the normal conditions established by the Executive Director, Library Services or nominee, for the care, loan and reproduction of theses.

Signed Steven Tyrrell McKinlay

Date

Acknowledgements

A very warm thank you goes to my principle supervisor, Professor John Weckert. I've very much enjoyed working with John, there is no doubt at all that John's encouragement and good humour, not to mention sound advice, guidance and ideas, carried me to the end. Thank you for everything John. Thank you to Dr Morgan Luck, my co-supervisor who also provided so many inspirational and thought provoking comments to my work. The three intense days in Canberra with John and Morgan, working through a draft of this thesis were indeed memorable. I thank you both very much for the time and interest you took in my work.

I'd like to formally thank the Centre for Applied Philosophy and Public Ethics (CAPPE) for funding the first couple of years of this project. Without the initial support from CAPPE it would never have got off the ground. I also thank Charles Sturt University for the generous research scholarship that enabled me to complete the work.

Thank you to Dr Stephen Clarke who was also involved in a supervisory role in the early years. I had a memorable meeting with Stephen and Luciano Floridi at Oxford in 2007 after an E-CAP conference in Sweden. Of course I am deeply indebted to Luciano, not only for providing material which continues to stimulate my interest but also for taking an interest very early on in my work and always encouraging me to continue.

I thank my employer Wellington Institute of Technology for supporting my research, providing time and also generously funding my attendance at various international conferences over the years of my study.

A special thank you to my collegue, Mariki Farrell, who very generously proof read this thesis. They say you never stop learning and hopefully, with Mariki's help, and having now finished a PhD, I might now know where to correctly place a comma. Any remaining errors are mine alone. Thank you Mariki.

There are too many friends and relations to mention individually who were interested, supportive and encouraging in their own way. I hope this acknowledgement will suffice.

A very special thank you goes to Jo Watson, Mum to our gorgeous son Jack and my friend for longer than I can remember.

I dedicate this work to my family, Morag, David, Shona and Stuart and especially to my Mum and Dad. I couldn't have wished for better parents. Knowing all along I was a dreamer, they just let me dream and encouraged and supported me without much fuss. Knowing she was not going to be around long, Mum said to me, squeezing my hand, "Make sure you finish that PhD Steven". Well, here it is Mum, for you.

Abstract

This thesis is a philosophical examination with the aim of developing more robust accounts of Luciano Floridi's informational realism and information ethics. Floridi's novel metaphysical framework argues that reality can be analysed from an informational perspective. Using object oriented terminology and theory Floridi develops his theory as a world consisting of dynamically interacting informational objects. Furthermore, Floridi's controversial information ethics argues that informational objects can be moral agents and have at least minimal intrinsic value.

Floridi wants us to take the information object idea literally. However, I dispute this claim, arguing in chapters 3 and 4, that while the concept has some utility its ontological value is difficult to defend beyond it being a useful metaphorical or linguistic term.

A complete metaphysical theory needs to address the problem of causality and I do this in chapter 5 using the metaphysical problem of absences to show how information can be employed to address causal issues.

The position I develop in these early chapters is brought to bear upon Floridi's Epistemology (chapter 6), and his Ethics (chapter 7). Chapter 6 develops the argument that at least in some cases knowledge and information are coexstensive, that is, in some cases prior or descriminating knowledge is required in order to interpret information. Chapter 7 introduces what I consider to be a signficant issue related to the use of the term "informational entropy" within information ethics. After identifying the problem I go on to offer a possible solution.

The ultimate aim of this thesis is to present a more coherent, more robust informational metaphysics and consequently this work contributes to and strengthens Floridi's information ethics as well as any epistemic accounts that are informed by the philosophy of information.

1 Introduction

The primary inspiration for this work arrived in 2004 in a landmark article entitled *Open Problems in the Philosophy of Information*. In this paper Luciano Floridi articulates eighteen defining problems for the emerging field the Philosophy of Information, and this thesis considers several of these problems covering ontological, epistemic and ethical perspectives. It was Floridi's paper that clearly expressed many issues that I had wondered about in the years prior to writing this. The problems and the questions posed by Floridi's article have kept me busy for the duration of this research and I expect will continue to do so.

The nature of philosophy is such that to move things along a robust critique of the existing material is often required. This thesis is a critical examination with the aim of advancing Floridi's philosophy of information. I begin by identifing what I consider to be significant metaphysical issues regarding the nature of information and what Floridi calls the information object, my goal is to provide a more robust account of the concept. I then deal with the issues that arise from the analysis and consequent implications for his information ethics and his epistemology.

Floridi asks us to take his information object concept literally, arguing that it constitutes the fundamental metaphysical unit. I will argue there are methodological problems with adopting this stance. In chapters 3 and 4 I argue that while the concept has some utility, its ontological worth is difficult to defend beyond it being a useful metaphorical or linguistic term. A complete metaphysical theory needs to address causality and I do this in chapter 5 using the metaphysical problem of absences to show how information can be employed to address the causal issues.

The position I develop in these early chapters is brought to bear upon Floridi's Epistemology (chapter 6), and his Ethics (chapter 7). The ultimate aim of this thesis being a more coherent, more robust informational metaphysics and that this work contributes to and strengthens Floridi's

information ethics as well as any epistemic accounts that are informed by the philosophy of information.

The central player within this work is Floridi's metaphysical concept of the informational object, its ontological status, its relationship with knowledge, explanation, reference and ethics. Note here that Floridi uses the term "informational object" and "informational entity" interchangeably. Except when quoting other authors, I will stick to the term *informational object* for clarity's sake.

In many ways the issues at stake are inspired by technology and as such this thesis is cross-disciplinary drawing upon ideas not only from philosophy but from information and computing technologies. The changes brought about by the Information Age have been profound. They have transformed our world and created new realities that challenge our values, beliefs, and confront conventional assumptions. In response to these challenges Floridi's work over the last decade or so has focused on two areas, Information Ethics (IE) and Informational Realism (2004, 2009).

What Floridi calls informational realism is in fact an informational metaphysics with the elemental unit being the informational object. There is a connection between IE and Informational Realism in the sense that IE draws significantly upon Floridi's metaphysics and the concept of the informational object as the prime entity of moral consideration.

Reality according to this account is ultimately informational and the Floridian interpretation of the universe is described as "the totality of informational objects dynamically interacting with each other" (2008, p2). Floridi's metaphysics is as fascinating as it is controversial.

At this point I ought to declare my naturalistic bias. Following Quine's naturalism, I assume that the measure of what there is, as well as how we know such things, is science. Throughout this thesis I try to draw upon empirical examples to bolster my arguments; as such there is an element of pragmatism in my methodology. Nevertheless at times, I do deal with intrinsically metaphysical issues, that is, problems that are perhaps not so easily explained with appeals to underlying phenomena. While I recognise that there is certainly potential for conflict between metaphysics and

empiricism my stance is perhaps better characterised as a form of rapprochement. I directly deal with the problems arising from this dualism in chapter 4.

Although I generally agree with Floridi that the idea of the information object offers utility to metaphysics as well as to an onto-centric theory of ethics, I begin my argument with a criticism of Floridi's use of object oriented (OO) theory and terminology as a way of clarifying the concept. Chapter 3 shows how information objects viewed from an object oriented perspective must render them second order. If information objects are to be the prime ontological units Floridi imagines them to be, then we must move away from any talk whereby such abstractions are seen as referential in nature (as this is precisely the nature of OO objects).

To save Floridi's information object it is essential we escape the referential trap. To fail, will carry transitive repercussions for an onto-centric ethics, for any potential information-first epistemology and finally the future of an informational metaphysics would ultimately prove problematic at best. For example, from an ethics perspective one such problem relates to that of intrinsic value. How could abstract referential objects possess intrinsic value? Such things must surely collapse into tokens or referential symbols. Many, if not all ethicists, no doubt would quickly dispense of any such idea.

Proceeding with what I hope is a more ontologically robust concept of the information object idea, an improvement on Floridi's original concept, we embark upon a comparative analysis. Does our new and improved information object concept stand up against or indeed improve upon more traditionalist metaphysical schemas (or is it indeed inferior)? This section will no doubt divide metaphysicians; there will be those that hold a view similar to David Armstrong (in particular 1998 and 2010). I conclude however that more traditional views, such as those of Armstrong are ultimately reconcilable with Floridi's schema. And it won't just boil down to semantic differences, indeed I provide what I hope are good empirical examples as to why we can take the Floridian information object idea seriously, if not literally, as he does.

Later chapters (4 and 5) elaborate upon the empirical evidence to support an informational metaphysics. Chapter 5 in particular presents what I hope is a compelling argument that at a most fundamental level, scientific explanation, causation and what Schaffer (2004) terms negative causation is ultimately informational by nature. This will show that an informational ontology can be defended not only from a philosophical perspective but arguably from a scientific or empirical point of view.

It is often said that philosophy raises more questions than it ever answers. And so in my closing chapter I consider the implications for Floridian Information Ethics in light of the subsequent analysis. In keeping with the naturalistic theme running throughout this work the chapter forwards a reconciliation of Floridi's use of the term entropy with more traditional scientific uses of the term.

To conclude, this thesis, *via* a series of chapters each examining slightly different but closely related themes, contributes to the emerging discipline that is the philosophy of information. We trouble shoot some quite serious shortcomings, clarifying some ontological and epistemological problems, and this in turn helps shape the nature of what looks to be a promising new addition to metaphysics and ethics.

1.1 Chapter Summaries

Chapter 3: The Floridian Notion of the Information Object.

Published in: *Luciano Floridi's Philosophy of Technology: Critical Reflections*, Part IV, 3. Springer, 2011.

This chapter deals with Floridi's informational ontology, in particular the validity of his conception of the information object. Floridi utilises popular object oriented programming theory in order to explain the properties and behaviour of the information object. The concept is important not only in Floridi's pursuit of an informational metaphysical scheme but it also plays a significant role in his information ethics. His controversial ontology presumes the universe to consist of information objects, something similar he argues to OO objects (2004). Whilst not dealing directly with Floridi's

open problem P.15: Wiener's Problem: What is the ontological status of information? (2004, p572) the chapter does address some of the issues related to this problem. Of course Floridi wants to use the terminology and theory behind OO as an example in order to explain his ontology. However, this tactic I believe exhibits shortcomings, the most problematic perhaps being that OO objects are notoriously theory-laden, that is, any discussion regarding OO objects presupposes the theory within which they exist, this being the one encompassed by Object Oriented Programming.

Further I argue that OO objects are referents and since referents are necessarily second order entities they cannot be rightfully considered basic or prime ontological units. Because OO objects are referents they exhibit explicit and distinct identity relations between the object itself, its constituent object class and its physical implementation in some computer program. In the computer world this system is described as "levels of abstraction". This isn't the case for real world Floridian information objects where the levels of abstraction are not explicit in anyway. My argument doesn't dispatch informational realism, it does however provide some evidence that, if indeed there are such entities as information objects, they can't be much like object oriented objects.

Chapter 4: Informational Realism and Metaphysics.

Paper presented at : IA-CAP 2011, Aarhus University, Aarhus, Denmark., 2011.

This chapter is a comparative analysis of Luciano Floridi's novel and controversial metaphysical account, informational realism, with the more traditional comprehensive account forwarded by David Armstrong, a world of states of affairs.

As we mention earlier Floridi postulates the information object as the entity central to information ethics and his informational realism. In developing the concept he draws heavily upon object oriented (OO) programming theory. Informational objects are reckoned by Floridi to be, in a sense, ontologically primitive and as such naturally occurring mind independent structures dynamically interacting with one another. Floridi

employs OO like terminology such as “properties” and “relations” in order to clarify his concept of the informational object.

Armstrong on the other hand postulates that the world, all that there is, is a world of states of affairs. A state of affairs according to Armstrong consists of a particular, which has a property or alternatively a relation which holds between two or more particulars. Each state of affairs as well as constituent higher or lower order states of affairs is a contingent existent. Furthermore the properties and relations attached to states of affairs are universals.

These two theories, whilst exhibiting marked resemblances also reveal fundamental philosophical differences yet both attempt to present a unified metaphysical schema, an ontology. Of great interest is the fact that here we have two strong competing theories. The situation begs critical comparison. This paper proceeds with the aim of striking some rapprochement between the two accounts. Any rapprochement of course must attempt to at least, save the good features of the respective accounts.

Chapter 5: Explanation, Causality and Information

Paper presented at: IA-CAP 2011, Aarhus University, Aarhus, Denmark., 2011.

Wesley Salmon recently discussed the connections between explanation (primarily scientific explanation) and causality arguing that to explain something is to state its cause. Equally whilst arguably not a necessary condition to “being informed” we could argue, probably in all cases, to have something explained is to be informed. It seems then that the concept of causality is linked in some way to information.

In section 1.6 above I raise the tricky metaphysical issue of absences. In this chapter we investigate this problem in more detail. Some absences (or what Schaffer calls negative causation) it seems are informative (Schaffer 2004) while Armstrong suggests absences (including preventions and omissions) seem to have no causal efficacy. Armstrong (2010) goes on to argue that in developing a systematic metaphysics we ought to treat

absences as second class cases of causality, while they add nothing additional they do supervene, "Consider for instance, a billiard table with balls moving around it sometimes hitting each other. You have all the motion of the balls and their hitting each other then you also have all the preventions and omissions that occur" (p83). Illari (2011) has recently looked at this problem and offered a production account of information-transmission. This chapter explains and investigates these particular issues.

Chapter 6: Information, Knowledge and Confirmation Holism.

Published in: *Current Issues in Computing and Philosophy*., IOS Press., 2008.

Problem 13 in Floridi's "Open Problems..." article is a fascinating one and one that opens several cans of worms. The problem with The Continuum Hypothesis: Should epistemology be based on a theory of information? Many information technology text books restate the problem as follows;

Data ➔ Information ➔ Knowledge

Floridi terms this the information continuum hypothesis and there are interesting things to note about the assumption. It suggests that data precedes information which in turn precedes knowledge. It suggests a linear and asymmetric relationship between the involved elements. And from a philosophical point of view it seems to suggest that information supervenes upon data, and likewise knowledge supervenes upon information. In this chapter I focus on these relationships but am primarily interested in the relationship between knowledge and information. In the chapter I challenge the asymmetry assumption calling into question the idea that there can be informational states without epistemic states, or indeed vice versa?

A significant problem for epistemology relates to belief states. Gettier famously argued in his 1963 paper *Is knowledge justified true belief?* that is, even though an individual could be justified in believing P and if P turned out to be merely accidently true then P could never constitute knowledge. This finding effectively turned epistemology on its head and thus the search for a non-doxastic or externalist account of knowledge began in earnest.

If the asymmetrical assumption is true, that is if knowledge does indeed presuppose information, then an information-first epistemology could be plausible. However what if the asymmetric assumption was incorrect? The chapter develops a counter argument (following Timothy Williamson, 2000) that the relationship between information and knowledge is in fact co-extensive. Further I draw upon Gareth Evan's work arguing that in order for an informee to understand any information-yielding event, be it a proposition or the perception of an object, photograph or whatever, some prior information must already be held by the informee.

Later in the chapter and as a result of the prior discussion, I highlight the implications confirmation holism has for the philosophy of information. Confirmation holism is the premise that all theories and consequently the statements that make up those theories are underdetermined by the data or evidence supporting them. The idea here is that statements or theories can never be tested in isolation. That is, we are always reliant upon other information, or knowledge in order to confirm (or disconfirm) any scientific or factual claims.

Chapter 7: Information Ethics and Entropy

Paper presented at: Australasian Post Graduate Philosophy Conference, Auckland University, 2012.

This final chapter brings prior work here to bear on Floridi's Information Ethics. The focus is primarily on what I see as two significant issues for IE, both related to our construal of the information object.

The first issue relates to Floridi's concept of entropy. We develop the argument that essentially all human activity in some form or another contributes to entropy (interpreted in the Floridian sense). This being the case we are faced with making decisions based upon which entropy contributing behaviours are morally blameworthy. This dilemma is inevitable since Floridi instructs us entropy ought not to be caused, or be minimised, or be removed and all informational objects (human and non human) deserve some minimal respect (Floridi, 2006, p24). If this turns out to be the case then it may be that IE ultimately collapses into

consequentialist-like claims as to which contributions to entropy we can accept against those we cannot.

Secondly, and related to the first issue – it seems on the face of it that some informational objects will simply not qualify as being intrinsically valuable, or at least such entities or objects will conflict with more traditional interpretations of intrinsic value. This objection can be cut off with an appeal to a metaphysical level of abstraction – an information object *qua* information object deserves minimal respect however (at a different level of abstraction) a Porsche Boxster certainly does not. This interpretation is the price we pay according to Floridi if we are to develop a foundationalist theory. However for levels of abstraction to really be a force majeure we need to be able to explain how an object at one level can have intrinsic value but not have it on another level. This chapter explores these issues and offers some solutions.

2 What is Philosophy of Information?

2.1 Introduction

The primary aim of this chapter is to introduce the reader to the terminology and current controversies regarding the Philosophy of Information. While some historical background is provided, typically the chapter will focus on controversies that are dealt with in more detail in later chapters of this thesis. Where this occurs it will be noted.

Luciano Floridi sums up the endeavour that is Philosophy of Information as follows:

> [Philosophy of Information] is a new philosophical discipline concerned with (a) the critical investigation of the conceptual nature and basic principles of information, including its dynamics (especially computation and information flow), utilisation, and sciences, and with (b) the elaboration of information-theoretic and computational methodologies and their application to philosophical problems. (2004, p555)

Chapters 3-6 of this work are primarily concerned with (a) the conceptual nature of information. The kinds of questions considered include:

- From an ontological perspective, just what comprises an information object? Can they be anything like the kinds of structures we find in object oriented programming as Floridi suggests?

- How far can Floridi's OO analogy go in explaining the nature of what he wishes us to accept as a fundamental metaphysical element?
- Can we accept informational objects as concrete entities, or is the concept a pure abstraction, that is, merely a useful kind of tool?
- What relation do such entities have with reality and knowledge?
- Is knowledge comprised of a multiple of informational nuggets?
- Can such an ontology successfully supersede previous metaphysical theories?
- What relationship does information have with causality and scientific explanation?

Chapter 7 moves towards (b) in its approach. I draw upon previous work and consider the implications for information ethics. The position established in the preceding chapters is applied to the question of whether an informational ontology can form the basis of a meta-ethical theory, namely information ethics (IE). Moral respect generally relies upon an argument regarding the intrinsic worthiness of the object in question. Thus, any onto-centric informational ethical theory must establish the intrinsic value of the entities it purports. Floridi argues that all informational objects have a minimal moral intrinsic worth *qua* information objects. However, this position has recently been challenged by Brey (2008). Brey considers how much sense it makes to confer moral value upon the kinds of things we generally consider to have no value whatsoever. His examples include such things as lint, email spam and toxic waste.

Considering Brey's examples, *prima facie* the claim seems absurd. If the elemental ontological unit is the information object and if every

informational object has some minimal moral intrinsic worth then everything has intrinsic worth. In chapter 7 we investigate this question.

Further, Floridi's use of the term *entropy* within IE is questioned. Physicalist explanations of the term entropy abound yet Floridi's use of the term bears little resemblance to such definitions. The chapter seeks to naturalise the use of entropy within IE.

Trying to find a unifying definition of information is problematic. To give a plausible answer is difficult since it has come to mean so many things across diverse disciplines. Nevertheless, it plays a key role in how we interpret, explain, model and represent our realities. There is an emerging tendency towards the intuition that it [information] must somehow be related to knowledge. Computer programmers and database analysts link it to the concept of data – information is said to be meaningful data. Indeed a semantic account of information as opposed to the more austere probabilistic mathematical accounts given by Shannon and Weaver (1949) and later Bar-Hillel and Carnap (1953) seems to have more relevance in our information age and it is this semantic construal of information that is at the heart of the present work. Nevertheless we should expect nothing less of a philosophical analysis of any notion to strip it back to bare basics. What are its fundamental components, relations and properties? And so when we ask "what is information?" we are in essence asking a metaphysical question not too different to the perennial philosophical anxieties about truth, causation, knowledge or meaning.

Floridi's elementary constituent, the informational object, plays a pivotal role in his metaphysics as well as in IE. It draws considerably upon ideas recently developed within the philosophy of information and is essentially the embodiment of his semantic notion of information (2004) but its inspiration also comes from concepts and narratives that will be familiar to computer scientists most notably those associated with object oriented (OO) programming and database theory.

There are of course various other theoretical and methodological positions on the notion of information. While Floridi frames information as

being semantic, qualitative and factual in its nature, other approaches include quantative and entropy based forumulations and deal with related concepts such as knowledge and belief in a formal logical way.

Shannon and Weaver's (1949) approach is mentioned above and considers information from the perspective of its encoding, storage and transmission. Earlier Fisher (1925) developed a probabilistic account of information whereby the amount of information associated with any random variable could be measured against the probability upon which that random variable depends. Fisher information, as it is sometimes called, has contributed to various areas in mathematics such as Bayesian statistics.

Computational conceptions of information have also been forwarded, most notably by Kolmogorov (1965) and Chaitin (1969). These approaches have informed algorithmic complexity theory, providing a foundation for artificial intelligence and human cognitive theories.

Floridi's construal differs from these in that it defines information semantically as comprising of well-formed, meaningful and truthful data. A significant distinction is that none of the formal quantitative approaches mentioned above imply truthfulness or wellformedness (Adriaans, 2012).

This thesis will primarily focus on Floridi's semantic notion of information and as such we shall leave further detailed discussion or comparison of other informational theories and concepts to other researchers.

In the computer science and information technology disciplines talk of entities, properties, relations, classes and objects is commonplace. And so we could well contemplate just what could the concept of the informational object entail. What kinds of properties are common to the concept? Is the notion in any sense explanatory? What kind of correspondence with real world entities do informational objects have and, what is the structure of such relationships? If we consider for a moment the view that the informational object is purely a theoretically constructed concept, how rational might it be to offer some kind of ethical theory based upon a seemingly abstract object and if so where does this leave real world entities in the debate? Is there indeed a rational distinction to be made? Are the

two somehow one and the same but just named differently? Could knowledge be somehow composed of informational objects? The following investigations we hope shed some light on these questions and the importance or otherwise of the informational object to the emerging field of information and computing ethics.

While the philosophy of information certainly concerns itself with the nature of informational (and sometimes digital) phenomena, in general information is most typically interpreted as semantic information and it is primarily this context upon which the current thesis will focus. For the time being let us think of semantic information as simply being information consisting of well structured data, that is, organised or expressed linguistically or by some other symbolic means, therefore conveying some meaning about itself. Thus semantic information could simply refer to a theoretical point on an abstract two dimensional plane or, alternatively, the entire set of plans for an A380 Airbus. That is, semantic information is information *about* some thing or state of affairs.

Of course the concept of semantic information is closely related to more formal sciences of information and computing and it is increasingly implicated in various epistemological issues involving the nature and range of postulated relationships between data and theory.

Information itself thus clearly represents an obvious link between IE and philosophy of information and is the core theme at the heart of this work. Where philosophy of information seeks, amongst other things, to find robust definitions of information, computer and information ethics employs these definitions in the service of illuminating ethical and moral problems related to computing and information technologies. As such we should hope that our construals of information *per se* tend toward correctness. Otherwise errors made about the elemental nature of information are likely to be transitive, that is, they will be carried through and perhaps amplified when developing metaphysical or ethical theories.

This chapter proceeds as follows. This following section describes the recent emergence of the information ethics, its current state, defences and

critiques. Later we take particular notice of the role the informational object plays in the debate. It should be noted that I intend to focus more sharply upon issues of importance to IE as developed by Luciano Floridi rather than computer ethics in general. Whilst not intended to be a comprehensive historical summary it will be useful to provide a brief outline of the development of information and computing ethics in general.

The second section introduces in more detail the philosophy of information. In a similar fashion I consider the main arguments and assumptions of this emerging field as it relates to this particular body of work. The discussion will primarily focus on the emergence of the general definition of information which lays the foundation for forthcoming chapters and focuses on the information object.

2.2 Information and Computing Ethics

Recently Floridi (1999, 2006), and Floridi & Sanders (2001) have offered a more tightly defined theory of information ethics. This came about in part as an attempt to provide some foundational basis for the disparate set of controversies that seemingly exist within computer ethics but also quite specifically as "a macroethical theory for analyzing particular microethical issues in computer ethics" (Tavani, 2008, p148). The aim of this section is to explore some of the motivating factors behind Floridi's development of information ethics. Additionally, I introduce the terminology and some of the controversies related to the topic. In section 2.4 I explore Floridi's ethics in more detail. We come back to Floridian information ethics in chapter 7. That chapter will address what I consider significant shortcomings of the theory and offer a more robust account.

Information Ethics as Luciano Floridi (2008) points out has come to mean different things to different researchers working across divergent as well as collaborative domains. This includes not only computer and information technology sciences but medical ethics, local and global

governmental, business and commerce studies as well as more formal studies of ethics within philosophy.

The roots of computer ethics are popularly traced back to Norbert Wiener's ground breaking work in the 1940s and 1950s. Witnessing unprecedented progress in machine automation Wiener considered what appropriate reactions we might have to the emerging new world (1954, p12). Of significance was his suggestion that we develop an ethical stance in parallel with the development of technologies. He was particularly concerned with what he described as communication and control issues between humans and machines. In many senses a modified version of this anxiety still exists today, no doubt shared by a great many more people. Wiener was primarily a scientist and a mathematician but also held keen philosophical views on the social and ethical implications of information and communication technology. His fundamental proposal was that "machines can and do transcend some of the limitations of their designers, and in doing so they may be both effective and dangerous" (1960).

Although it was at least two more decades before the use of the term "computer ethics" was in common usage it is quite clear that Wiener's thoughts were well ahead of their time and laid the foundations for thinking about computer ethics. In his text *The Human Use of Human Beings* (1950), he considered not only issues familiar to ethicists such as justice, 'the good life' and the like but he also developed a step-wise methodology for applied ethics with direct consideration to the implications of the forthcoming 'information age'.

> Where a man's word goes, and where his power of perception goes, to that point his control and in a sense his physical existence is extended. To see and give commands to the whole world is almost the same as being everywhere.
> (1954, p97)

Wiener's trepidation is directly linked to the use of information. He argued "society could only be understood through the study of messages and

the communication facilities which belong to it" (1954, p16). Wiener correctly predicted a future whereby messages between humans, between humans and machine and even independently between machines themselves, were destined to play an increasing role in society (*ibid*).

It remains debateable today as to whether machines (in particular computers, either today or in the future) do in any sense exhibit intelligence comparable to human beings. Wiener was not so concerned with this popular preoccupation of the time. He preferred to dispense with such arguments as being essentially "semantic" in nature. The definition of terms such as 'intelligence', 'life' and 'purpose' did little, he argued, to avert the drawing of behavioural analogies between machines and living organisms. Wiener preferred to explain the behaviour and performance of such entities (human or machine) by way of their interaction with the external world. Humans were just another (albeit special) kind of system (1954, p57). Of course it is acknowledged that such complex issues regarding machine intelligence and issues surrounding the nature of consciousness etc. are at the heart of cognitive science. Without wanting to diminish their significance, these issues are by and large beyond the scope of this work.

No doubt Wiener considered such points somewhat extraneous when compared to the seemingly far more immediate, practical and social implications brought about by the increasingly widespread automation of the workforce at that time. It is from these perspectives that information and computing ethics proceeds today. That is, computers can perform colossal numbers of operations at levels of detail vastly more intricate than the general limited abilities of most humans. Clearly a double edged sword exists. Quite how our world today, in all its complexity, might proceed without such functionality is difficult to fathom. Yet any intelligent analysis of such performance is often retrospective in nature since the development and use of such technologies progresses at a speed difficult for any single person to monitor. How are we to predict what future devices may be created and to what uses such devices will be employed?

Weckert (2007) characterises this as a tension between an "ethics-last" versus a more proactive model with respect to computer ethics. According to Wiener (cited in Weckert, 2007, p13), adopting a wait-and-see stance may well result in "disastrous consequences". As such the agenda of computer ethicists has typically been to at least promote the development of ethical and moral understanding alongside the development of the technology. This approach will be in part predictive as well as reactive as Weckert (*ibid*) points out. We cannot be expected to accurately predict all of the possible implications and consequences, at the same time we would not wish to unnecessarily stifle technological development with the assumption that such development would generally produce undesirable results. Instead most computer ethicists would adopt a "proceed with caution" approach.

So although computer ethics in some form or other has been around for several decades, mused over primarily by a handful of computer scientists and technicians, library and information professionals and philosophers, it has only relatively recently emerged as a topic earning serious consideration by applied ethicists and philosophers. The numerous reasons for more recent and much wider interest in the topic ought to be somewhat obvious to most whom have used a computer, and perhaps more so, the internet. Only a few imagined the technological transformation that has taken place over the last half century or so culminating, certainly within the developed world, in an almost unprecedented commercial, societal and cultural mind shift. Our now adolescent 'information age' is increasingly blurring the division between our so called digital or online lives and natural lives raising bona-fide ethical concerns on an almost daily basis.

Information and communication technologies have become an unpretentious extension of our day to day lives to the point where many of us simply cannot do without them. Commercial as well as social activities are relentlessly conducted online 24 hours a day 7 days a week, year in year out. Yet it seems only recently have we had grounds to pause to consider the nature of the impact of this revolution upon our lives. Increasing public

as well as academic awareness of the various moral issues and the consequent desire and urgency to develop appropriate ethical responses has essentially been driven by the mere fact that the technology is both increasingly commonplace and globalised. Thus, from a traditional perspective at least, we might begin with the following question: what are the key ethical issues related to computers and information technology today?

Richard Mason (1986) suggested that although the ethical issues involved were many and varied they could be broadly outlined within four categories (summarised by the well known acronym PAPA) Privacy, Accuracy, Property and Access. Each of these categories are worthy of consideration Mason argued, with the ultimate moral imperative that human dignity be enhanced or at least preserved. Mason's categories have been expanded in recent years to include security, computer crime, issues related to freedom of expression, equity and issues of responsibility and professional ethics. What follows is a brief summary of the most popular topical issues for computer and technology ethics. We include and expand upon Mason's original list of categories. The focus will be mostly based upon the role information plays within each.

2.2.1 Privacy and Anonymity

We expect widespread recognition of our right to privacy. To say our privacy has been invaded infers some kind of loss of a good that many individuals feel entitled to exercise sole control over. On the other hand many often consent to a relaxation of their privacy, usually with specific conditions. It is common to provide various details about ourselves in order to fully participate within society. However, technological advancements, most obviously the Internet, combined with the seemingly relentless gathering of information appear *prima facie* to constitute serious threats to privacy. In particular, privacy related to the information that specifically pertains to individuals. This special category of privacy has been described as informational privacy (Floridi, 1999., Tavani 2004., Tavani 2008).

There are potentially vast bodies of data directly related to ourselves that might be collected, stored, analysed, altered and consequently shared, matched or sold. Such data might include our online internet browsing habits, records of retail purchasing history through loyalty schemes, credit card and eftpos records, medical history, beauty salon history, vehicle registrations and parking tickets, airline ticket and hotel bookings, academic achievements, email, telephony not to mention our Twittering, Facebook and Tumblr posts, Skype, online chat records and the like.

Floridi (1999, p52) describes informational privacy as the "freedom from epistemic interference or intrusion". When data about particular individuals is appropriately restricted or not known, or indeed unknowable then informational privacy is preserved. However, within today's information rich world the issue is more complex. Don't look for TV cameras in your bedroom Larry Hunter argues (1985) but for those analyzing the information that is already willingly shared.

Unknown information regarding individuals often exists only due to the disparate and separated nature of the possibly related data. Much data that exists about individuals is often sparsely represented in varied purpose built, geographically separate, individual databases which manage information for specific functions, such as banking, medical records etc. Although this in itself may present privacy issues due to inappropriate use or access, it is the combinatorial effect inherent in the correlation, matching, processing, manipulation and reconstitution of such disjointed sets of data about individuals that is often more concerning.

Furthermore, Nissenbaum (1998, in Weckert) points out that there is virtually no limit to the amount of information that could be gathered. Similarly, the range of data mining or analysis that could conceivably be performed upon such data along with the fact that the resulting information could be stored virtually forever only serves to increase the possibility for abuse.

Floridi considers the point that in such cases it is not that the derived information from multiple sources might have legal consequences for

individuals, or might be embarrassing, shameful or harmful, it is the fact that care and respect for an individual, (whom Floridi somewhat radically claims "is their information" (1999, p53)) has been violated. Informational privacy according to Floridi is the maintenance of the personal integrity of information which constitutes "me-hood", or in other words the dynamic evolving informational object that is representative of me.

Brennan and Pettit (2004) note that many of us develop internet-specific identities, that is, many people create single or multiple virtual identities completely distinct and separate from their real world identities. I may choose to represent my gender, age, my species or various other aspects about *me* in a completely different way on the internet. Such information often appears to only ever represent something imaginary or simulated with very minimal links to ourselves. The question as to whether the same privacy boundaries apply is a topic currently attracting considerable debate[1].

Virtual environments offer unprecedented affordances unconstrained by traditional real world limitations. Furthermore, opportunities for interaction by such agents are not merely limited to other humans and their immediate worldly environment but artificial agents and the virtual environment itself. Questions as to whether traditional theories of privacy apply in such environments remain largely unanswered. The question of identity itself looms large for agents embedded in virtual environments.

Privacy concerns also extend to other forms of technology. Brey and Soraker (2009, p55) suggest that we consider the implications of video surveillance and raise the question as to how much privacy individuals might expect in public areas. They go on to claim that "The notion of 'public' itself has changed in light of information technology" (*ibid*). The debate within current mainstream and social media regarding an individual's

[1] See Brey, Chp 15 in Tavani (2008) for example. Brey considers the relationships between virtuality and reality and whether virtuality can or should function as a substitute for ordinary reality.

expectation to privacy in public spaces, whether those spaces are virtual or "real world", is particularly salient.

Questions of anonymity, particularly with regard to internet use, are often included within the privacy debate. The benefits of anonymity are similar to those of privacy (Bynum, 2008). For example, anonymity may offer individuals some protection while using the internet to obtain information or participate in disucssions on sensitive topics such as gay rights, abortion, sexually transmitted infection, as well as political and religious issues. (*Ibid*). On the other hand these issues must be balanced with the understanding that anonymity can also be used to harm. Examples include individuals using anonymity for illegal personal gain or amoral purposes. Questions also arise regarding the use of anonymity within social media contexts where the anonymous user subjects others to cyberbullying.

2.2.2 Cybercrime and Security

Issues to do with security and cybercrime have a particulary high profile. These involve offences committed against individuals, groups, governmental and private sector organisations with the intent to cause harm. While there has been vigorous debate regarding the ethics of cyberwarfare and cyberterrorism these particular issues only represent a small fraction of the real concern. Issues related to spam, hacking, malware, viruses, copyright infringement, identity theft, child pornography and the use of social media and web based forums to groom children are far more prevalent.

Many of the issues here once again relate to the protection of personal information. For example, Van den Hoven argues that many of the informational vulnerabilities in society relate to the harm that can be done on the basis of personal data – "theft, identity fraud, or straightforward harm on the basis of identifying information." (2008, p311).

Perhaps the most recent addition to this category is the security related issue of state monitoring and control of informational infrastructures with

the aim of protecting against terrorist attacks. The debate involves getting the balance right with regard to invidual's civil liberty and privacy.

2.2.3 Accuracy

Inaccurate information can result in negative impacts for individuals and organisations alike. Who is accountable and what responsibilities do individuals or corporations have for informational accuracy? In particular accuracy is raised as a key issue in surveillance environments (Sewell and Barker, 2001). For example Introna (2005) points out that even under ideal situations the facial recognition systems, widely in use in the US to identify 'terrorists', have a capability only to ever achieve a 70-85% accuracy rate with performance degrading significantly in crowded environments such as airports, malls etc. Introna following Brey (2000) argues that the internal operation of information technology is generally not "evident, obvious, transparent or open to inspection by the ordinary everyday person affected by it" (2005, p75). Of course the reasons for this often relate to the corporate entity protecting their intellectual property interests and these protections are not just limited to the realm of IT applications[2]. A special problem arises however when proxy decisions are made via an algorithm or "system" which adversely affects the rights of an individual and where the rationale behind such decision making is hidden or obscured from external parties including the individual. Furthermore, even if an individual could inspect the inner workings of a particular algorithm or computer program it is still not clear that anyone but an experienced software coder in the particular language used could readily understand how such processes and subsequent decisions are made.

On another front Vedder and Wachbroit address the increasing problem of determining the reliability of information on the internet (2004), from simple content to search engine results rankings. This relates to Introna's

[2] I thank Morgan Luck for pointing out this equally applies to for example KFC's 11 different herbs and spices and various other corporate product secrets.

previous comment regarding the black box nature of such technology – search engine algorithms are typically a closely guarded secret. Thus, what constitutes informational reliability and where does the responsibility rest?

2.2.4 Property

The informational property debate has been typically characterised by a tension between those that advocate for an intellectual commons versus those that argue for legislative protection of intellectual property (Tavani, 2004). According to Tavani (ibid) there are two key concepts related to informational property rights: the notion of an intellectual object and the intellectual commons. Typically intellectual objects (a certain kind of informational object) are non-tangible and, they might be digital in nature or expressed in print such as books, musical recordings or software programs. The concept of scarcity that applies to tangible objects and resources do not exist in the same sense in the case of intellectual objects (*ibid*, p6). Such objects can be copied many times over with little or no consumption of raw resources.

The "intellectual commons" can be explained as the realm of intellectual material, knowledge and information open and accessible to everyone and not owned in any real way by any individual. Ethical approaches to property attempt to understand and resolve issues related to property rights and the intellectual commons.

2.2.5 Access and Equity

There are two senses in which 'access' can be interpreted in relation to computing ethics. In one sense the debate is centred on access to information and communication technologies (ICT). Typically labelled the digital divide, Compaine (2001) characterised the issue as essentially the gap between those that have access to ICT, that is, "information tools" and the ability to utilise those tools and those that don't, or where such access is severely limited for various reasons.

Himma (2007) outlines in some detail the empirical evidence for the "information gap" suggesting that such a division perpetuates poverty in afflicted nations. He goes on to show how our ordinary intuitions as well as normative ethical theory entail a moral obligation on the part of affluent nations in closing such information and digital divides.

Kvasny (2007) on the other hand approaches the issue from a more socio-historical perspective depicting the digital divide as having underlying existential causes. Simply providing more physical access to ICTs according to Kvasny does little to close the gap or understand the unique socio-cultural needs of people who have traditionally been excluded from information technology and the digital world. "We produce discourses that discount their values and cultures and show them why they need to catch up" (Kvasny, 2007, p205).

The other sense in which the access issue can be construed relates to access to freedom of expression. The debate centres on the competing issues of censorship versus possible harms caused by access to certain kinds of information. Whilst most of us believe that our access to freedom of expression ought to be largely unconstrained we also realise that this must be balanced against the consequential harms related to such access. As an example we see restrictions emerging upon antilocution or what is commonly today known as "hate speech". This is typically defined as threatening, abusive matter or discourse likely to incite hostility against groups of people on the grounds of their ethnic origin, skin colour, religious faith, sexual orientation or the like. Censorship represents a limit to our access to such kinds of human expression.

Cohen (1993) (amongst many others) argues that human interests are best served if there is free and open access to the expression of others. Cohen suggests this concern is inherently linked to three fundamental human interests: our interest in expression; in deliberation; and in our inherent interest in information. This "informational interest" described by Cohen presumably allows us to deliberate and reflect against a background of alternate viewpoints expressed by others. Such viewpoints may not always

be palatable but it is the full range of alternative views, Cohen points out, that allow us to deliberate upon and pursue perhaps more admirable beliefs, values and actions.

Mathiesen argues "that someone censors when they think that it will be a bad thing if other persons access certain expressions" (2008, p579). She goes on to argue that typical normative descriptions of "badness", that is when someone accesses some informational content that one might think is inherently bad, or that it might be considered instrumentally bad if someone were to access such expressions, are in fact usually not reasonable grounds to justify censorship. Mathiesen's argument turns on the premise that we may not be so good at distinguishing between what we simply don't like and what might cause actual harm.

2.2.6 Freedom of Expression and Content Control

The issue is often framed as the balance between censorship and free expression. Questions arise as to what level of control should the state have over content control and censorship balanced with the individual's rights to freedom of expression and freedom of access to such content. There is no doubt the internet has become a significant player in shaping social consciousness.

For example, the role the internet played and the subsequent governmental censorship and control in the Arab Spring is unprecedented. Stepanova (2011) argues that while information and communication technologies (ICT) and social media had little to do with the underlying socio-political and socio-economic factors there is no doubt that ICT and social media networks had an initial mobilizing effect. Stepanova goes on to argue that new ICT networks are likely to have a critical effect in countries where the governing regime has little or no social base (2011, p4).

There are several types of speech which are identified as being candidates for censorship. For example, hate speech such as those associated with facist and racist groups (Brey & Soraker, 2009, p56). Other types of speech might include forms that attempt to harm or undermine the

state or individuals such as information about how to build bombs. The role of ethics relates to the levels of permissibility and methods by which such expressions might be controlled.

2.2.7 In Summary

Upon reflection we note that the issues discussed above, with the exception of cybercrime and security, are not specifically computing or information technology related issues, notwithstanding the examples provided. The dissenter might be tempted to raise an objection, questioning the need for an ethics of computers in order to deal with issues that have been around a great deal longer than computers themselves. The fact that we have witnessed an "information revolution" does not in itself sustain the argument that we need a separate field of ethics.

Moor (1985) recognised this objection and argued that unlike previous technologies, computers were "logically malleable", that is, computers could be "shaped and moulded to do any activity that can be characterised in terms of inputs, outputs and logical operations." Additionally Moor argued that computers seem to be the closest thing we have to a "universal tool" thus the limits to the uses that computers might be put to seem only to be limited by our imagination.

This logical malleability has allowed us to use computers in ways we never imagined, performing functions that were not possible only a decade or two ago. The question therefore arises, ought we do such things? Often no laws or policies exist to regulate such use resulting in what Moor termed "policy vacuums". Such policy vacuums Moor went on to argue create "conceptual muddles" (1985).

The seemingly limitless capability afforded by computers, Moor argues, offers an equal array of possible choices with regard to our use. However, the kinds of discourse about how we ought to conduct ourselves, given such technology environments and opportunities, often seem inadequate in the absence of considered policy. Moor pointed out that even when prior policy existed it often seemed inadequate for conduct within such newly created

situations. Following Moor's analysis, the need for a computer ethics is not only driven in response to various ad-hoc ethical dilemmas related to computing, but additionally due to the increasing globalisation, ubiquity and speed of computing, information, the internet and communication technologies, as well as the virtually limitless ways in which the technology can be applied.

2.3 The Uniqueness Debate

On the face of it Moor seemed to have established a good reason to study computer ethics, however quite early on in the debate the most significant challenge to the idea was expressed by the question we have already hinted at, that is, are computer ethics issues any different in their fundamental nature to other moral questions? For example, we could conceivably argue that privacy violations are privacy violations regardless of whether a computer was involved or not (Tavani, 2001). Generating a great deal of robust discussion, this particular methodological issue came to be known as the uniqueness debate.

Prominent in the uniqueness debate have been Don Gotterbarn (1991, 1992) and Gotterbarn and Rogerson (1998). In response to Moor's malleability argument Gotterbarn argued that "the underlying flexibility of math and logic is greater than that of the computer, but we did not develop 'logic ethics' and 'mathematics ethics' " (1991). The solution Gotterbarn suggested was that any ethical or moral considerations relating to the use of computers ought to be dealt with by 'professional ethics'. Although not ruling out "the examination of critical concerns like the impact of technology on the nature of work or computer fraud" ultimately computer ethics according to Gotterbarn could be adequately dealt with by existing professional ethics and codes of practice (*ibid*).

Gotterbarn is perhaps most known for his claim that we have developed many technological devices that have had significant impacts on society in the past but have found no such requirement to develop a parallel ethics.

Arguably there have been comparable revolutionary societal shifts to that of the computer revolution. Nicholas Carr (2004, p16) for example points out "various contemporary commentators … have identified particularly strong parallels between the rollout of the rail network in the mid-1800s and the expansion of information technology, particularly the Internet, in the late 1900s". Indeed Moor (1985) and Gorniak-Kocikowska (1996) compare the industrial revolution in England with the computer revolution. Gorniak-Kocikowska argues that like the printing press the computer revolution is global in its impact thus we need a new global or universal ethical theory to respond to the problems created by computing, information and communication technologies:

> The more computers change the world as we know it, the more irrelevant the existing ethical rules will be and the more evident will be the need for a new ethic. This new ethic will be the computer ethic. (Gorniak-Kocikowska, 1996)

There are no doubt other technologies (other than computing and information technology) that have profoundly affected our lives. For example the tractor revolutionised agriculture, the telephone, television, and as mentioned, Carr suggests the locomotive. Computers and the internet are the revolution of our age, however is this in itself a good enough reason to establish a new field of study? Tavani suggests the question might be a red herring. The ethical issues he argues are "both philosophically interesting and deserving of our attention, regardless of whether those issues might happen to be unique ethical issues" (2002, p37).

Walter Maner (1996) on the other hand makes the distinction between ethical instruction and moral indoctrination. The latter, that of "socialising" students into "professional norms" he argues is the corollary to Gotterbarn's approach. The unfortunate result being students developing a political rather than academic bias and therefore more likely to create an environment of

impartiality which, he suggests, could overwhelm honest enquiry. Of course we note that this is more of a methodological issue than one related to the uniqueness of the moral dilemmas related to computing and information technology.

In any case, Maner argues that a number of new and unique ethical issues do in fact emerge that simply could not have arisen if computer technology did not exist. A concerning example is the ability of either government or private enterprise to gather, aggregate and data-mine data from virtually any online or offline source. An example is found in the Spokeo web site. A social network aggregator that gathers and makes publically accessible data from many sources such as telephone directories, social networking sites, real estate listings, government census data and a wide variety of other sources. It is the speed that such information can be accessed and the accuracy of the data that transforms this particular issue into a unique problem.

In his 1996 paper Maner presents a variety of reasons culminating in the idea that we ought to study computer ethics "because the set of novel and transformed issues is large enough and coherent enough to define a new field". The failure to find satisfactory non-computer analogies to such problems Maner goes on to say, is testament to the uniqueness of computer ethics.

Other prominent positions on just what computer ethics comprises include Deborah Johnson's genus-species argument (1994). Johnson, neither surrendering entirely to Gotterbarn nor Maner's ideas, considers the ethical problems related to information and computing technology are best understood as a new species of existing problems and, thus worthy of consideration. For example, Johnson argues that information technology has given rise to new kinds of entities, such as digital databases and software and in recent debate (Shoemaker, 2009, Dillon, 2009) suggests information technology is changing the very way we construe personal identity. The kinds of ethical questions raised are related specifically to these entities, for example, do very large databases containing personal

information constitute threats to personal privacy? Johnson, rather than concede that such issues are new ethical problems, considers them a "new species of old moral issues". As such there are specific ethical questions about these new kinds of entities that arise, but they are not entirely in themselves unique. For example, should we protect software and data ownership by law and, do very large databases that contain personal information pose a threat to personal privacy?

It should be noted that the recurring theme of computing and information technology ethics from Wiener forward is primarily concerned with promoting and protecting standard core values such as life, happiness, freedom, resources, knowledge and so on. This focus certainly has its roots in the classical and typically western anthropocentric ethical theories such as Kantianism or duty based ethics, as well as varieties of virtue and contract based and consequentialist ethics and utilitarianism. The motivation according to Tavani (2008, p36) has been towards developing ethical approaches that anticipate harm to human values and aims to develop new technologies whereby that harm is somewhat mitigated. Nevertheless, it is clear that information and computing ethics is *prima facie* an applied or practical ethics. Most of the issues and debates are related to the ethical questions surrounding data privacy, software (errors, ownership, etc.) access to the technology and autonomous agents or artificial intelligence. The stake holders are of course computer professionals, policy makers, software and internet users as well as professional and academic ethicists and philosophers. Clearly the approach from this perspective is a human-centric approach.

It is this largely pragmatic approach that concerns Floridi. Pragmatic approaches are often considered short sighted, relativistic and ad hoc. The criticism goes something like this: while the application of a particular ethical solution to a specific ethical problem can often be found, there exists no foundational basis from which to analyse the problem domain as a whole. Thus the problems associated with the development of information and computer ethics have been largely methodological in nature. Floridi (2002)

categorises the problem as follows; computer ethics tends to move from a problem analysis focus toward tangible tactical solutions, and why not? This is the generally accepted methodology in computer science and information technology. In response to unforeseen problems that have caused "policy vacuums" computer ethic's rational, yet largely inductive response, has been to develop practical solutions embodied in professional codes of practice, technical standards and new legislation. Floridi points out this constitutes a bottom up approach, destined to fighting spot fires emerging out of the seemingly "random collection of ICT-related ethical problems" (2002, p2).

We have considered some of the issues surrounding computer ethics in general, let's now consider in more detail Floridi's suggested solution, Information Ethics.

2.4 Floridian Information Ethics

What follows is an outline of Floridi's Information Ethics; we consider problems and objections and begin our investigation into the Floridian conception of the Information Object. Often overlooked by those considering Floridi's information ethics is the role his theory of informational realism plays. Thus the aim of this section is to introduce the general terminology and to clearly establish the link between information ethics and Floridi's informational realism.

In 2008 Springer's Ethics and Information Technology journal devoted an issue to Luciano Floridi's philosophy of information and information ethics. In the issue Charles Ess notes "it is difficult to overstate the contribution and significance of Luciano Floridi's work – first in conjunction with Jeff Sanders and then in his further developments of the philosophy of information and correlative information ethics". Floridi's overriding goal over the last decade or so has been to develop a basis for Information and Computer ethics which adequately addresses what he calls

the foundational problem. That is, the concern that computer ethics proceeds in an ad hoc manner addressing issues and problems as they arise.

Underpinning Floridian information ethics is his conceptualisation of a philosophy of information within which information is reified in everything that exists. Central to the theory is the notion of the informational object. In developing this notion Floridi draws heavily upon object oriented programming concepts, thus following object oriented theory, the answer to the question 'just what is an informational object?' Floridi must answer 'everything'.

In a recent paper titled, *A Defence of Informational Structural Realism*, Floridi defends an ontological commitment to "a view of the world as the totality of informational objects dynamically interacting with each other" (2008, p219). Earlier however (cf. 2004, p572) he flags this as an open problem for the philosophy of information. The specific questions relating to this and questions considered in detail in this thesis are "what is the ontological status of information?"[3] and, "can nature be informationalized?" (p574) "We are asking whether the universe itself could essentially be made of information, with natural processes, including causation, as special cases of information dynamics..." (*ibid*).[4]

From a traditional perspective, ethical theories approach moral problems from a largely anthropocentric basis. They are often said to be agent-focussed and so the *objects* at the centre of moral consideration and respect are human beings; thus the principles are derived from core human values, such as life, happiness, rationality and the like.

More recently however the standard ethical theories have been extended to include bio-centric and even eco-centric standpoints. Other living organisms and certain ecosystems as such have been included as the kinds of things that are said to have intrinsic moral value. That is, the idea that

[3] I deal with this particular issue in some detail in Chapters 3 and 4.

[4] The issue of causality and its relationship with information and informational processes is dealt with in chapter 5, *Explanation, Causality and Information*.

something possess some value 'in itself' or 'for its own sake' and therefore commands respect regardless of and more important than its material (or instrumental) or extrinsic value. A recent eco-centric approach has been termed Deep Green Theory (Curry, 2005). Deep green theory holds the entire ecological community as being equivalent to the ethical community, thus the value-orientation of deep green theory is earth-centred as opposed to strictly human-centred. Curry does not intend that deep green theory replace human-centred ethics; rather it is meant to compliment and extend, thus enabling behaviour not accounted for under the traditional anthropocentric approaches.

Floridi's IE steps even further away from the traditional anthropocentric and newer eco-centric approaches. IE is an "object oriented and ontocentric" theory (Floridi, 1999, p43). The controversy surrounding IE is that it further enlarges the idea about what can count as objects worthy of moral consideration to include not only people and animals and ecosystems but anything that exists present, past or future "from paintings and books to stars and stones" (*ibid*). The ethical concern for IE rests not on the question 'what ought I to be' or 'what ought I to do' but instead 'what ought to be respected and/or improved?' It is the 'whats' welfare according to Floridi that matters the most (*ibid*, p49).

We have noted already that the methodological criticisms of computer ethics arise largely from an argument (the uniqueness debate) that suggests there are no special ethical problems exclusively related to computer ethics, rather there are ethical concerns in which information and digital technology are involved. Thus computer ethics at best, from this point of view, would be considered a Microethics, that is, a domain specific, applied or professional ethics and as such problems could be adequately dealt with from traditional ethical perspectives. Floridi's approach has been to forward a Macroethical framework of Information Ethics which he argues is superior to other Macroethical approaches such as philosophically significant, field independent ethics such as Consequentialism and Deontologism.

He begins this counter position in his landmark 1999 paper with the following claim:

> The essential difficulty about Computer Ethic's philosophical status is a methodological problem: standard ethical theories cannot easily be adapted to deal with CE-problems, which appear to strain their conceptual resources, and CE requires a conceptual foundation as an ethical theory. (1999, p37)

The tendency has been to categorise CE as an application of action-centered ethics. The concept most central to the action-centered view is that of autonomy. Thus the event critical to the action-centered account is a free and intentional mental action, most typically a decision made by a person. Autonomy is necessarily linked to agency and rationality. And so the autonomous will of rational agents or persons is a basic construct of an action-oriented ethical approach, hence the anthropocentric focus.

The 'strain' seems to come in the form of the "projection of human agency, intelligence, freedom, and intentionality (desires, fears, expectations, hopes etc.) onto the computational system" (Floridi, 1999, p39). Without such an 'object oriented' approach, Floridi argues, computer ethics seems destined to anthropomorphise computational systems.

Bynum (2008, p37) agrees that the "human-values approach" needs to be widened to include much more than human beings and their actions and intentions. There are clearly difficulties associated with the application of Kantian imperatives where free will, rationality and autonomous agency are absent; at least any attempt would require a considerable abstraction of such principles. Floridi (1999, p40) presents two types of problems as counter-evidence to a successful application of Kant's moral imperatives. "Neither the law of impartiality (the Golden rule) nor the law of universality (behave as a universal legislator)" (*ibid*) Floridi argues, are sufficient to address the following issues:

1. Computer ethics problems not involving human beings.
2. Computer ethics problems that are ludic in nature.

Let us first consider the first statement, problems not involving human beings. Any motivation to develop an ontocentric macro-ethical theory must come from an opinion that there exists entities, other than and independent to humans, that are worthy of moral consideration. In consideration of the first statement Floridi offers an allegory (1999, p54) involving a young boy playing destructively in a dumping ground. In the allegory the boy amuses himself by smashing dumped car windows. Floridi argues this behaviour pointlessly increases the level of entropy within the dump yard and according to a maxim prescribed by information ethics, 'do not increase the level of entropy'; we can make an objective ethical judgement against the boy's behaviour.[5]

It is important to note that Floridi's definition of entropy differs somewhat from the physicist's concept of thermodynamic entropy. It refers to any form of destruction or corruption of informational objects, or "any kind of impoverishment of being" (Floridi, 2004, p30).

A significant problem for information ethics is Floridi's construal of entropy and the application of when (and when not) it applies. In chapter 7 I argue that virtually all human activity increases the level of entropy; thus the problem for IE is to explain which entropy raising activities are more morally blameworthy.

Floridi's allegory is essentially an attempt to illustrate the central claim of IE, that is, everything that exists (due to its status as an information object) has at least some minimal level of intrinsic value and is therefore deserving of moral respect. Most traditional ethical theories are human or

5 Floridi's allegory, its implications and his use of entropy within information ethics is explained in full in chapter 7.

agent-oriented whereas IE, as Floridi explains, must be patient-oriented.[6] Rather than being concerned with how agents ought to behave, IE considers the question; which kinds of objects deserve moral respect? And in answer to such a question Floridi argues by default, all objects.

With respect to the second issue, computer ethics problems that are ludic in nature, Floridi considers the difficulties in applying traditional ethical approaches to CE problems having a ludic or game-like nature. Within virtual environments this is relatively easy to conceptualise. No one is ever charged with murder after killing virtual entities inside a computer game. Even the virtual killing of avatars representing real world people in online gaming environments can be completely acceptable within the bounds of the game, in fact some games entirely dedicate themselves to this mode (Modern Warfare II, and the massive online multiplayer environment, World of Warcraft come to mind). The consequence, Floridi argues (1999, p40) is that not only do most people feel no responsibility for their virtual actions, but they are usually willing to extend the Kantian universal maxim to all participating agents. Everyone is free to partake in the carnage. Having said this, games such as Modern Warfare and World of Warcraft have their own internal ethics or rules of engagement. Disregard for the established etiquette and standards of behaviour is likely to get you excluded from the group or even the gaming environment.

Abdication of responsibility when the agent believes their actions are insignificant or not fully real, just in case those actions are being carried out in a virtual or informational computing based environment, is not grounds in itself to morally excuse certain actions. Floridi argues, for example, that the increasing incidence of computer crime committed by seemingly respectable moral individuals illustrates the limits of CE. Such individuals,

6 The 'patient' is typically in an interactive relationship with an agent, the agent being the initiator of some morally burdened action and, the patient the receiver, reacting (more or less) to the situation.

Floridi argues, rarely understand the implications of their actions independently of their technical competence (*ibid*).

2.5 Regarding the Intrinsic Value of Information Objects

Critical to IE is the issue of whether or not things *qua* information objects, no matter how trivial, have (even minimal) intrinsic moral value. Thus the main point raised in this section is to introduce the controversy regarding the attribution of intrinsic value to information objects.

As mentioned earlier, Floridi employs Object Oriented (OO) terminology and methodology in developing IE and it will prove helpful to briefly summarise this approach. Object Oriented programming is a data-centric view of computer programming where data and behaviour are strongly linked. Objects are class variables responsible for their own behaviour. We can talk about a class as denoting a certain predicate, and its objects as the extension of that predicate, alternatively the true propositions obtained from the predicate. Hence objects are often described as instantiations of classes. A simple way to understand this follows. Consider the class of birds. The class of birds is a somewhat abstract notion; however an instantiation of a bird, for example the pigeon wandering near our table at a café, is a concrete instantiation of that class.

A central and simple conceptual idea in OO programming is that of inheritance. It is certainly not uncommon in nature to encounter situations where a particular class of objects has certain properties in common whilst some objects in that particular class have local properties of their own. For example all ellipses have an area, and some ellipses, namely circles, have a radius (Date et al, 2000). This example not only addresses the concept of inheritance but also introduces sub-typing in OO. Thus, all the properties that apply to ellipses (the supertype) are inherited by circles and so a circle is a subtype of ellipse. So it would seem to object-oriented analysts and programmers and Floridi alike that such concepts are suitable for modelling reality.

Floridi applies these ideas in general to his bottom-up, minimalist approach to moral worthiness – the intrinsic value of a particular object is considered with respect to its local and inherited values.

> An object can be treated as a sufficiently permanent (a continuant) information monad, a description of the ultimate primal component of all beings. The moral action itself can be constructed as an information process, i.e. a series of messages invoked by an agent, that brings about a transformation of states directly affecting the patient, which may variously respond to those messages with changes or other messages, depending upon how the message is interpreted by the patients methods (Floridi, *paraphrased*, 1999).

Utilising the OO concepts of inheritance and sub-typing, Floridi goes on to argue any object that can act as a moral agent can also qualify as a moral patient, but not vice versa, for example animals can be moral patients but not moral agents. Therefore it is better (from an ethical perspective), Floridi suggests, that the focus be upon the informational nature of an object as a possible patient rather than a possible agent. The moral worth of any object as a moral patient, "in the most universal and abstract terms" (*ibid*) allows us to extend to any object our general conclusions, thereby including any possible informational element that may be affected by the behaviour of a moral agent in any given environment including the infosphere itself (*ibid*).

Floridi's minimalist, bottom-up approach looks for the minimal conditions of moral worthiness represented by an aggregate of local and inherited attributes relating to the particular information object in question. When focusing on the possible intrinsic value of the object by virtue of its local attributes only, particularly where those local attributes are essential properties of the class of human beings, Floridi argues, is similar to a Kantian maximalist axiology, where only a limited set of attributes (eg. intentionality, rationality etc) qualify the object as having intrinsic value.

Instead Floridi poses the question – what is the intrinsic value of O *qua* an object constituted by its inherited attributes? A pluralist approach to the question refers a variety of inherited attributes perhaps including intelligence, sentience and feelings right down to mere existence itself. Floridi imagines a hierarchy of inherited attributes and their associated priorities in moral standing – hence the minimalist theory. The resultant claim is that *minimal* intrinsic worth is universal and shared by all entities that are more or less morally respectable. Although this respect, Floridi points out, may be overridden with regard to other levels or degrees of dignity afforded to other objects.

Floridi argues IE is the last stage in the development of the trend away from anthropocentrically focused ethical theories, maintaining that even biocentric approaches are morally biased and too restrictive in scope – inanimate objects can now share intrinsic value and are thus entitled to some minimal degree of respect. The set of inherited attributes that would constitute the "least biased and most fundamental solution" to the problem of an object's least intrinsic worth, Floridi suggests, is its abstract nature as an informational object. That is, an object *qua* an informational object is the minimal condition of possibility of moral dignity and hence normative respect (Floridi, 1999). This raises an interesting question: could there be legitimate degrees of intrinsic value, that is, do some things have more intrinsic value than other things? If so, this seems to violate the principle of intrinsic-ness; surely something either has or does not have intrinsic value. This issue is considered in more detail in Chapter 7, *Information Ethics and Entropy*.

Floridi's theory is certainly controversial and consequently is not without its detractors. Brey (2008) presents two objections to Floridi's IE. On the one hand Brey considers the validity of attributing intrinsic value upon information objects. Later in the paper Brey questions Floridi's materialist ontology of information objects. Let us consider Brey's arguments.

Following Baird Callicott (1995)[7] Brey questions Floridi's commitment to confer intrinsic value on non-human entities:

> IE tells us that we should be equally protective of human beings and vats of toxic waste, or any other information object, and that we have an (albeit overridable) duty to contribute to the improvement and flourishing of pieces of lint and human excrement. At best, this suggests very little guidance in making moral choices. At worse it suggests that IE gives us the wrong kind of guidance. (Brey, p112)

Although Brey agrees an extension of traditional anthropocentric ethical approach has some merit, he argues that Floridi makes a mistake in inferring from the fact that biological organisms and some inanimate objects (following bio and eco-centric ethical approaches) deserve respect that everything is now deserving of moral respect and is intrinsically valuable. Floridi's inference is invalid according to Brey because it conflates the idea of deserving respect and intrinsic value.

The account of intrinsic value Floridi develops is of information objects set against the background of what he terms the infosphere. The term captures "the environment constituted by the totality of information entities – including all agents – processes, their properties and mutual relations" (Floridi, 1999, p44). This construal follows computer and information science's object-oriented approach and is favoured by Floridi because it begins with the assumption that everything is an 'object'. Furthermore only objects can communicate with other objects by virtue of their relationships or where explicit messages are passed from one object to another. Floridi

[7] Baird Callicott an environmental ethicist raises fundamental theoretical questions regarding the intrinsic value in nature and the need for a nonanthropocentric environmental ethics. If we are to have such a theory we need to have good reasons to do so, particularly since many may not consider certain objects as having any intrinsic value in the first place. Although Callicott concedes intrinsic value is not a wholly objective property, he suggests value like other natural properties is actualised by a situated observer/valuer. (1995, p73).

wishes to extend the class of some specific entities (those perhaps deserving of respect, and that we might adjudicate to having intrinsic value) to all of reality – that is the infosphere within which all objects exist.

The infosphere is analogous with an ecosystem or the biosphere and so requires at least our minimal respect. Furthermore, we have a duty of stewardship towards the infosphere – to allow it to flourish and to ensure we do not contribute unnecessarily to entropy, in fact according to Floridi we are obliged to prevent and remove entropy from it where we can. But does this make sense? Could even trivial informational objects have intrinsic value? Furthermore as we mention earlier (and discussed in detail in Chapter 7) if virtually all human activity ultimately contributes to entropy, then some kinds of entropy will be more ethically acceptable than others, but which kinds?

The concept of intrinsic value has garnered considerable philosophical interest in particular over the last century or so. Typically the notion centres upon the idea that whatever is intrinsically good must be something worthy of being valued purely for its own sake. Anthropocentric interpretations of intrinsic value usually include appeals to the unique and often difficult to define properties associated with humans, properties such as free will, rationality and the like. Floridi points out however, that even when humans don't exhibit such properties (such as free will and rationality), we still seem to attribute intrinsic value upon those individuals. He gives the example of individuals such as with severe brain damage or with other severe injuries or illness. Therefore, he goes on to argue it must be some other property that engenders the attribution of intrinsic value. These other such properties, outside of free will and rationality, however, do not seem to be unique to humans alone and as such Floridi argues that as these generic properties exist in other things we ought to extend the class of moral patients to all entities exhibiting such properties.

Brey argues this approach is a mistake. "First, it wrongly equivocates between deserving respect and possessing intrinsic value, and second, it overgeneralizes from specific classes of entities to all of reality." (2008,

p111). Part of the problem may be in Floridi's method. In developing a list of properties that might constitute intrinsic value Floridi argues that they must be the "most general possible common set of attributes which characterises something as intrinsically valuable" (2002, p297). The answer Floridi suggests is the informational nature of any entity that may act as a patient of a moral action – this is the lowest threshold of inherited attributes that constitute its minimal worth. Thus being "an information object *qua* information object is the minimal condition of possibility of moral worth and hence normative respect" (*ibid*). This being "the central axiological thesis of any future Information Ethics that will emerge as a Macroethics" (p298).

It could be that the validity of attributing intrinsic value upon any kind of object at all, least not conceptual objects, is not merely litigious but erroneous. Some think we can decide upon which things have intrinsic value on the basis of their relationship with other things or events. G.E. Moore (1903) for example argues that the intrinsic value of a thing is independent of its relationship with anything else. So, the intrinsic value of a thing would not be reduced by the removal of everything else that it is related to, "if they existed by themselves in absolute isolation, we should yet judge their existence to be good" (1903, p187).

Clearly Moore's argument presents problems with regard to the relationship between conceptual information objects and their respective environments. For example one kind of information object, that kind that object oriented theory proposes, object classes and their instantiations are purely digital phenomena and are certainly not actualised in nature. That they reference something real, that is, something that exists in a spatio-temporal sense seems problematic on Moore's account. The "real world object" exists separately from the object oriented model which references it but on Floridi's account the instantiated object model is as worthy of respect as much as the object it represents. Indeed Floridi actually considers these objects one and the same but just interpreted at different levels of abstraction. Nevertheless, on the OO account, such informational objects

rely physically upon the computational environment within which they exist and logically upon the object which they reference. Removal of such relationships would render the objects meaningless. That such referentially based models could be worthy of respect in themselves independently of the things they reference seems on the face of it absurd.

On the other hand, Beardsley considers the possibility that "if anything is intrinsically valuable then it is not an external object, but an experience or psychological state" (1965, p3). He goes on to argue that something that has value is to say it deserves to be valued. This concept of intrinsic value seems closer to Kant's approach. For Kant it wasn't the object that possessed any intrinsic or absolute value but instead it is good will itself that is the only thing or attitude worthy of ascribing such esteem. Everything else, Kant argued, could only ever possess extrinsic or economic value to varying degrees but is essentially morally irrelevant.

On the basis of such scepticism Dewey (1939) and Beardsley (1965) argued that intrinsic value does not exist at all. Both forwarded pragmatic challenges to the concept of intrinsic value. All value according to Beardsley is extrinsic value. Beardsley's argument seems to hinge on an issue related to epistemic access. Even if some objects do have intrinsic value; we cannot know which do hence the role of intrinsic value is redundant in our reasoning about the certain types of value objects can have.

Prima facie the wholesale ascription of intrinsic value upon external real world objects is problematic. Transitively attributing intrinsic value upon concepts, models or notions representative of those real world entities then seems even more troublesome. An informational representation for example, it would seem, would not immediately appear to have the same level of value than the actual thing the representation represents.

In order to sort these issues out we first need to clarify the problem of reference and its relationship with informational objects. This particular issue plays a central role in this thesis and is dealt with at length in chapters 3, 4 and 5.

This section has summarised some of the issues related to ascribing intrinsic value to things *qua* informational objects. The issue is re-visited in more detail, along with Floridi's use of the term *entropy* as it relates to his ethical theory, in chapter 7.

2.6 Ontological Issues

This particular topic is related to the problem of reference which mentioned briefly above. Chapters 3 and 4 return to ontological issues; however, at this point I would like to sketch a broad outline of the problems.

We can and often do, for sake of argument, concede to certain entities. Sometimes those entities refer to actualised objects in space-time, sometimes not. For example some people like to talk about UFOs as little flying saucer like objects containing aliens. Those that would deny the existence of such things argue that this is just a fanciful idea in people's minds. We seem to be talking here about two different kinds of things, an actual real world object and a mental idea. The mental idea isn't what people are talking about when they deny the existence of UFOs.

It seems reasonable to conceive of two significant kinds of things, purely mental ideas and on the other hand objects, states of affairs or events that exist in a spatio-temporal kind of way. Into this ontology we are now asked to admit of a third kind of entity, the information object. Information objects can represent both actualised spatio-temporal types of objects as well as mental ideas. And there is nothing to stop us imagining information objects representing other information objects or indeed packages of information objects, even self-referencing informational objects.[8]

In fact there is infinite scope for a somewhat unlovely overly populated universe full of imaginable let alone possible (or indeed impossible) unactualised information objects, all of which, according to IE, we are

[8] In object oriented programming as well as within relational database theory the facility exists to conceive of entities with recursive relationships with themselves.

required to have some respect for, all of which come under our stewardship, all of which have at least some minimal level of intrinsic worth. Quine may well have objected to such a universe arguing that it "offends the aesthetic sense of us who have a taste for desert landscapes" (1953, p4).

The question as to whether such entities have intrinsic value or are deserving of respect is one issue, another question, one asked by Floridi himself in Open Problems... (2004), "What is the ontological status of information?" (p572).

I therefore flag three problem areas associated with ontology, to be dealt with in greater detail in later sections of this thesis:

2.6.1 Concerning identity.

Can two informational objects be identical? If so would they be one and the same? Could no two possible informational objects be strictly identical? What about identical digital copies of *information packets* or object oriented objects?[9] The issue of identity, I don't believe, is adequately addressed with regard to informational realism. Most likely the identity of an informational object will be explained with reference to its properties and relations. And this approach is more likely to be what Armstrong (2010, p55) calls a "loose and popular" construal of identity. This approach allows for the perpetual flux that is the time-space continuum. People grow and eventually die, mountains erode down into sand. In general macroscope items are always changing. In the philosophical jargon this is often called this *perdurance*, put simply, the idea that objects persist with 'loosely' the same identity over time. Nevertheless formal metaphysical theories must in someway address the issue of identity. In chapter 4 I offer some insights as to how informational realism might deal with this.

[9] Note Floridi uses the Object Oriented (OO) concept to try to explain the informational object. I think this is a mistake, a full explanation as to why is given in chapter 3. In chapter 4 I look at how OO deals with identity within OO programming theory. This provides further evidence as to why OO is not a reliable methodology for explaining informational realism.

2.6.2 Concerning reference.

The things and events that informational objects reference are purportedly informational objects themselves. Yet it seems to me at least, that there is a physical separation between the things themselves and the informational object that is representative of that thing. Which of these entities deserves more moral consideration? Is there a distinction? That claim might be made that that information is more akin to reference than to reification. Floridi may argue that the reification fallacy is diffused by his concept of levels of abstraction (explained in detail in the next section). This tack however, on the face of it, creates problems for any proposed ontology – in information science at least and certainly from an object oriented programming perspective, information has always been a somewhat abstract concept referencing (more or less accurately) the real thing or state of affairs it purports to model or represent. These issues are dealt with at various points throughout this work.

2.6.3 Concerning other modalities.

Informational realism typically deals with actual entities, that is, it considers real objects in the world such as tables, chairs, mountains and the like to be informational objects. However, is it possible that alternative modalities such as necessity, probability, possibility, and impossibility could apply to informational objects? Metaphysical theories often either include provision for or at least consider the implications of such modalities. What kind of implications could such modalities have for the moral respect IE requires us to extend to informational objects? While Floridi's philosophy of information instructs us to exclude 'false information' do we have a duty of stewardship toward 'not necessarily false' but possible or probable informational objects? Some might consider the question to be nonsensical; I disagree. For example, if we assume we truly have freedom of choice, then those choices result in very real future possibilities; this could well entail future possible informational objects. With regard to informational

realism, can *possible* or *probable* informational objects have any ontological basis? These issues are briefly discussed in chapter 4.

2.7 Levelism – The Method of Levels of Abstraction

> The moral value of an entity is based on its ontology. What the entity is determines the degree of moral value it enjoys (Floridi, p294, 2002).

Floridi's method of levels of abstraction plays an important role in his metaphysics as well as his ethical theory. This section outlines in some detail the method of levels of abstraction that is referred to at various points in this thesis.

The conceptual analysis of informational objects within Floridian IE (and, their consequent intrinsic worth) is considered at a specific level of abstraction (LoA). That level of abstraction is the one whereby all objects are considered informationally. The idea of levels of abstraction, sometimes termed levelism in philosophical literature, presumes other points of view are not only possible but equally valid. Thus on one view objects might be considered fundamental or prime, however Floridi's Information Ethics considers an entity's (minimal) intrinsic value as being fundamentally related to the entity's ontological status as an "information object" (2002, p287) but not as an LCD screen TV, wooden spoon, piece of lint or whatever.

We note that most recent claims by Floridi seem to move away from the universal position above and somewhat limit the scope of an informational objects intrinsic value. Most prominently "... ontological levelism is probably untenable. However, epistemological levelism should be retained as a fundamental and indispensible method of conceptual analysis" (Floridi, 2008, p304). I will briefly outline Floridi's initial argument regarding levels of abstraction before discussing the recent change of tack.

At a particular level of abstraction we are warranted in our analysis of an entity, agent, or a patient as an informational object, constituted by a collection of common inherited (informational) attributes. Floridi draws on analogous terminology in Object Oriented Programming (OOP) where the concept of levels of abstraction is familiar. OOP is a computer software development approach which considers all processes and data interactions within the program to be handled by "objects".

Thus we can make a clear distinction. Object oriented computer programs are (only ever) representative of their corresponding problem domains.[10] That is, a customer database program for example, might model PERSON object classes and PRODUCT object classes. The instantiated objects reference the real world objects more or less accurately – they are essentially just a computer model of the real world, built for a specific purpose – to record data about people and products and sales and the like. There is never an isomorphic relationship between the object model and the real world objects – the two are separated in the sense that a business does about its business of selling products to people and the computer system records these transactions and events. The relationship between the actual business happenings and the data inside the OO computer program is an asymmetric referential relationship.

In other words the computer program utilising the object oriented model seeks to solve some business problem, related to recording various information, relationships and the like occurring between the real world objects (people, products and so on). Although such object oriented models are useful utilities, particularly to organisations wishing to develop information technology systems that help them manage their data, there are theories of the physical objects they represent that are epistemologically superior.

[10] Note, I deal with this issue and Floridi's use of object oriented programming theory and terminology in considerable detail in chapter 3.

Getting back to Floridi, the attributes he wants to talk about are of course fundamentally informational in nature and are the common observables at the particular LoA being considered at the time. So, just as the bits and pieces of the human body are comparable when comparing humans with humans, an object *qua* informational object is comparable to other objects in that it shares the same kinds of properties, that is, informational properties, as opposed to blood, bone, various cellular materials, water, and so on.

Following the OOP paradigm Floridi places informational objects in their rightful environment, the infosphere, this being the "context constituted by the whole system of information objects, including all agents and patients, messages and their attributes and mutual relations" (2002, p289).

The idea of levels of abstraction is not new; indeed one popular formulation of the idea comes from computer science. While the idea was discussed earlier it was perhaps first formalised in papers by Dijkstra (beginning with 1968). The concept behind levels of abstraction is somewhat reductivist in nature. We can analyse realities at differing levels; thus at one particular level at least, ensuring we are comparing apples with apples. At higher more abstract model-theoretic levels however, analysis produces a broader scoped understanding with lesser correspondence between the model and reality. Lower levels of abstraction yield a higher degree of correspondence between the object and reality. Thus, Johnny could be analysed as a complicated multicellular organism, or as a mammal, perhaps more commonly as a male person. Alternatively Johnny could be analysed (at a clearly more abstract level) in a system-theoretic way, as an informational system with inputs and outputs, internal processes, or indeed "as an information object that interacts and shares a number of properties with other information objects, like a customer digital profile." (Floridi, 2002).

Quite differently to Floridi's LoA concept, Dijkstra's 1968 paper describes five levels of abstraction, 4 of these levels corresponding to actual programmatic implementations, level 5 being the computer operator "not

implemented by us" Dijkstra jokingly comments. At level 0 Dijkstra discusses the goings on inside a computing processor.

> At higher levels we find the activity of the different sequential processes, the actual processor that had lost its identity having disappeared from the picture (1968, p343).

Moving up through the hierarchy of levels the supposed identity of the previous level disappears as it is supplanted by the current level of analysis.

Dijkstra was by no means the first to offer such a pragmatic analysis. Quine (1953, p10) argues "One's ontology is basic to the conceptual scheme by which he interprets all experiences". How else are we to judge anything outside of some particular agreed upon conceptual scheme. Judged in another conceptual scheme Quine argues ontological claims may be adjudged false (*ibid*).

2.8 Food for Thought: Absences and Information.

The final section of this chapter very briefly introduces a metaphysical issue that is dealt with in detail in chapter 5.

The problem of *absences* according to Armstrong "are really rather horrible, ontologically speaking" (2004, p83). Yet absences can be significantly informative and explanatory. I use the term 'absence' mostly in the common sense understanding of the word. So an absence is an omission, or a prevention of a particular action or event taking place – such things do not exist, they are absent. Armstrong gives an example of a father pulling a child out of the way of an oncoming car. If the father had not pulled the child out of harms way, the child would have been hit by the car. The car hitting the child, that is, "the non-existent thing that was prevented has no causal efficacy." (Armstrong 2010, p83). Phil Dowe (2000) makes a distinction between omissions and preventions arguing that both involve counterfactuals. If the child had not been pulled out of harm's way the

counterfactual is reversed and the omission becomes 'the grabbing of the child' and it also has no causal efficacy.

On the other hand Jonathan Schaffer, in an influential paper, argues that absences or what he terms *negative causation* can serve as a cause, effect or causal intermediary. Negative causation according to Schaffer is genuine causation and his argument attempts to show how causes need not be physically connected to their effects (2004, p197). Schaffer goes on to argue that scientists across the disciplines routinely talk of negative causation. Illari (2012) discussing Schaffer's work offers examples from bio-chemistry. I will discuss the examples in detail within chapter 5. Suffice to say, there is *prima facie* evidence that in the case of absences an information exchange is occurring. This in turn seems to suggest that absences can indeed be causally efficacious. Thus my approach in chapter 5 will be to forward a thesis that where an absence is part of an explanation we can give an information-theoretic account.

3 The Floridian Notion of the Information Object

This chapter appears as a chapter in Demir, H. (ed.) *Luciano Floridi's Philosophy of Technology: Critical Reflections*, Springer, 2012.

The world is the totality of facts, not of things.
(Wittgenstein, 1922, [1.1])

3.1 Introduction

Floridi's *information object* concept is at the heart of his informational realism and information ethics. This chapter critically examines the concept. In particular I argue that the use of object oriented terminology to explain information objects within a metaphysical frame work is ultimately flawed. While the chapter does cast some considerable doubt upon the ontological status of the information object, chapter 4 will go on to offer a solution that might strengthen informational realism.

Ontological questions are questions about the nature, existence or reality of objects. And whilst there is a deceptive air of simplicity about the most basic ontological question[11], 'What is there?' the equally simple and somewhat obvious answer, 'Everything' leaves us somewhat unsatisfied. Obvious controversies arise when a scientist or philosopher argues that there is something or other which she purports exists, to which I or another scientist or philosopher would not agree. Thus with regard to questions of ontology Quine reminds us "there remains room for disagreement over cases" (1953, p1).

It's perhaps no coincidence that the Object Oriented (OO) programming community has adapted a similar maxim to their own end. To the question

[11] This question was famously coined by Quine in his 1953 article "On What There Is"

"What is an object?" the OO analyst would also answer "Everything" yet just exactly what an 'Object' is, is still by and large up for grabs, not only to ontologists, and information theorists but perhaps surprisingly to OO programmers themselves who one would have thought had a mortgage on such terminology. Thus like information, meaning, truth or belief our concept of 'object' is multifarious, capturing both abstract ideas as well as concrete ones. Correspondence between what computer programmers often call "real world objects" and the more or less analogous abstract predicates they like to call "object classes" has always been a tenuous one. The ultimate aim of the OO designer is to capture as much semantic information as is possible about the *real world* objects she hopes to model yet there is always the recognition that any OO model, whether at conceptual design levels or logically implemented as a working OO application, is merely a representation of some (usually tightly defined) subset of external reality.

Further, whilst it is probably true that most OO designers are out and out *realists* about the *real world* objects they seek to model there remains considerable debate both within the OO fraternity and in philosophy in general about just what an object is. The fact is all of our computational modelling efforts, whether they are Object Oriented by nature or that they might employ some other methodology *du jour* really only represent a flimsy link with reality tainted thoroughly with all the human-made ideas, theory and language associated with that particular conceptual scheme. This is similar to the problem which, at least from the philosophical perspective, led Quine to maintain that "the very notion of *object*, or of one and many, is indeed as parochially human as the parts of speech; to ask what reality is *really* like, however, apart from human categories, is self-stultifying." (1992, p9).

Drawing heavily upon OO theory and terminology Luciano Floridi utilises these human made constructs and language in order to clarify his *information object* notion. The notion is not only a critical element of his Information Ethics (IE) but plays a central role in the service of his wider theory of Informational Realism. The Floridian information object however

seems be an unusual ontological case[12]. This is not just because an information object is not a spatio-temporal kind of thing; we might find and (perhaps) agree upon the existence of all manner of abstract objects, indeed we agree that an information object is an abstract, non-spatiotemporal kind of thing. And whilst there is plenty of debate surrounding the ontological status of abstract objects in general that is not the primary focus here instead our attention will converge upon the nature of Floridi's information objects and his use of OO theory and concepts to clarify this concept.

The information object concept seems unusual because the information contained by an information object seems to be *about* something else – that is, it seems to be *about* the actual object or state of affairs that is described or referenced by the information object. At least this is certainly the case for OO objects in that any object class, following Quine (1992, p6), is just an abstract entity to which each real world object bears a tenuous relation, in other words, my object class for 'cat' for example may well be quite different from yours even though we generally agree upon every cat. This is because cats are almost without exception learnt about via ostension whereas structural abstract objects such as object classes are known only with regard to their role in cognitive discourse and never by ostension (*ibid*).

Thus there is a significant difference in the way Floridi talks about and wants to utilise information objects and the way in which an OO designer or application uses OO objects and their progenitors, object classes. Accordingly the *information objects* unusualness is amplified by the way Floridi seems to want to treat information objects, that is, as independent and external[13] objects of themselves, almost as if they were something more than abstract and worthy of genuine ontological status. It may be that this

[12] We are obliged to point out that Floridi does limit the scope of his adoption of OO concepts and theory by saying "OOP is not a viable way of doing philosophical ontology, but a valuable methodology to clarify the nature of our ontological components" (2004, p5).

[13] By "independent and external" I mean something whose existence is independent to human thinking or perceiving, and therefore would exist whether or not (for example) humans existed, in other words 'observer-independent'.

talk is merely a convenience or some kind of metaphor about information objects; if indeed this turns out to be the case then our job will be to clarify such talk.

Floridi acknowledges questions surrounding the nature of information as legitimate threads of enquiry (2004, 2008c) if information is not an independent ontological category then to which category could it be reducible. On the other hand if it (information) indeed does constitute a valid ontological category then another problem emerges, just how does it relate to the objects to which it usually refers? Such questions lead to enquiry *vis-à-vis* the nature of information *per se*, its relationship with meaning and its status as a natural human independent phenomenon or entity.

Although Floridi relies upon the terminology and the conceptual framework that is representative of the object oriented (OO) programming and design paradigm his literal application of the concept of the information object differs considerably from the service the OO object[14] is put to within OO computing. That the concept warrants any ontological status seems, on the face of it, at odds with the OO conception of the *object* which considers objects to be *referents*. To explain, just like Floridi's Information Objects, the object concept in OO programming is inextricably linked to the concept of abstraction yet there are some significant differences. Firstly the sense in which an OO object is more or less representative of some *real world object*. By the term 'real world object' I mean to refer to any material or non-material (abstract) object, concept, state of affairs or idea that may be represented (or modelled) using the usual constructs offered by OO-theoretic means. Indeed this tenuous relationship, between OO objects and

[14] When I use the phrase "OO Object" I am talking about the structure and function of Objects in the service of some OO application, design or model. I want to distinguish this from the phrase "Information Object" which embodies the meaning explicit in Floridi's IE and Informational Realism and while Floridi uses OO terminology and method to explain his conception of the information object I want to show how, even if we do accept the 'information object' concept, they cannot really be like OO objects.

their real world counterparts, constitutes a necessary condition of all OO objects – such objects are always modeller (observer) dependant.

Secondly that the OO object is a logical and abstract structure that acts as a *referent* to data elements physically stored on disk. The OO object does not and cannot refer directly to any real world object insofar as it is a computationally implemented model-theoretic interpretation represented by a particular class of structures – those structures being an Object Oriented domain class.[15] Just as a landscape painting is a more or less accurate representation of the landscape it seeks to represent so too is an OO domain class a representation of some business, engineering or scientific problem.

Finally the work that such an object performs being delimited by its duty and scope within the OO application within which it resides. Such objects cannot meaningfully exist in isolation, either outside their particular domain class or without some supporting technology framework. Nor can such models come into existence without careful prior analysis of the laws and relationships governing the behaviour and the nature of the entities and attributes themselves of the extant system that the object model seeks to represent.

Thus there is no doubt whatsoever that all object structures within the OO paradigm, modelled or instantiated within an OO application, are through and through human-made entities or artifacts. And so to suggest that reality could be something like OO class structures or objects seems to me to be like saying that the extant reality represented by a landscape painting could be something like the painting. Of course whilst appearances are identifiable with the model (or the painting) and this suggests some level

[15] Development of an *object oriented domain class* begins with the systematic identification and modelling (and diagramming) of all the entities or objects, attributes, operations and relationships that an OO designer perceive to be important about a particular problem domain, be it a business oriented problems such as invoicing or accounts receivable or a scientific problems such as the modelling of biological or genetic systems and so on. Individuals within the domain class are then generalised and represented as object classes which characterize the structure and behaviour common to all objects in that class.

of empirical adequacy of the model our structural representations applied to either concrete or abstract objects within our vicinity are "rooted in innate predisposition and cultural tradition. The very notion of object at all, concrete or abstract, is a human contribution, a feature of our inherited apparatus for organizing the amorphous welter of neural input" (Quine, 1992, p6).

At this point we flag an important distinction. My position so far is not directly in conflict with Floridi's claim that information *qua* information is prime, prior or observer independent, however we ought to be wary of accusations of theory-ladeness [16] associated with directly comparing seemingly naturally occurring phenomena, in this case information *per se*, presuming such a thing does exist independently to human perception, with information *qua* OO objects. We shall have more to say on this specific distinction later.

To summarise, where the OO paradigm considers the object concept to be structural, abstract, referential and artifactual by its very nature, Floridi's informational realism attempts to apply this idea to the age old philosophical dilemma, "what is the nature of reality?", to which he answers the "totality of informational objects dynamically interacting with each other" (2004b, p1). This informational ontology is then explicated with extensive and explicit reference to Object Oriented Programming (OOP) methodology, "OOP provides an excellent example of a flexible and powerful methodology with which to clarify and make precise the concept of the informational object" (*ibid*, p5).

The point is this, OO objects are explicit about their referential relationships, OO objects are indeed *referents*, but informational objects, if

[16] The idea of theory-ladeness comes from philosophy of science whereby scientific observations are said to be theory-laden when the language and terminology used to describe such observations in question is largely derived from the theory itself. Thus, discussions about the nature of *information* using object oriented terminology could be accused of being non theory neutral. Having said that it is difficult to see how any discussion regarding information could not be influenced by various aspects of culture and language.

they are to be considered as basic ontological units surely could not refer to any other object. There seems to be a fundamental contradiction going on here, and it is to this that this paper directly speaks.

3.2 The Properties of Object Oriented Objects

This section offers a brief examination of the OO programming and design conception of the object. Developing an understanding of the terminology and background of OO programming and design is required in order to critically review how Floridi's notion squares with the theory.

Any review of the OO literature reveals a key notion; *abstraction*. The *abstraction* concept is the underlying idea from which all other OO terminology emerges. Common concepts (terminology) used in OO circles include the concepts of *class, inheritance, encapsulation, polymorphism* as well as the *object* concept itself and these are all implemented via an abstraction layer realized by the *method*.

Although the concept of class and object are interwoven there is an important conceptual distinction. Class structures are considered abstract in that they are defined and realised as a generalisation of the concrete, instantiated individual objects of that type. Of course instantiated objects are not *really* concrete; instead they represent individual members of the particular abstract class which defines them. Thus the class *REPTILE* represents the properties common to all reptiles, cold blooded, scaly skin and the like. While we typically identify (real world) individuals in any class via ostension, this particular lizard or that particular crocodile, OO objects are always an abstract expression representing an instance of the defining class. The distinction between abstract classes and concrete instantiated objects is relative. By analogy an object class is a predicate and an object a proposition.

Inheritance is a function built into an object's structure whereby objects in a hierarchy inherit the data elements and behaviours of their parent object class. Thus subtypes in the *REPTILE* class, such as *CROCODILE*, inherit

all the reptilian attributes and behaviours and then add a few properties specific to crocodiles.

Encapsulation is an OO specific mechanism whereby an object's components are restricted from being directly accessed. That is, the internal representation of the object is hidden from the outsiders view – just how the attributes and behaviours of the class *REPTILE* are implemented within the OO application is not available for scrutiny by users of the system.

Polymorphism refers to the ability of an object's methods (which might also be implemented as operators), as designed by an OO designer or programmer, to be utilised in more than one way depending upon context within which it is used. Thus, consider the operator '+', appropriately defined we might use to add up number data types or alternatively if presented with text strings the operator may concatenate or append them to a list depending upon the usage desired by the designer.

The *method* is the mechanism by which all object interactions and manipulations are performed. Note here that under all self-respecting OO development environments objects are instantiated *via* what is often termed a class constructor method, they certainly do not pop into existence spontaneously just because some corresponding real-world object needs to be represented. It is directly due to the abstraction approach that one cannot have direct access to data that might exist inside an object; instead the *method,* sometimes utilising a *message* (often called *parameters* in programming), must be invoked. The message may also contain some identity condition – that is a way of identifying which object you wish to refer to. Furthermore object creation within OO is explicit. Objects are created as per the needs of the application or the application user.

Persistent objects usually refer to real world objects, or states of affairs – something that the OO designer wishes the OO application to represent. Objects may persist in which case they would be correspondingly represented in a database somewhere after the application closes or they may be temporary, existing only while the OO application is running.

Temporary objects are often associated with the general operation or running of the application. For example in a Windows application scroll bars, dialog boxes, menus and the like are all instantiated objects which exist and are represented in the memory of the machine (or server) running the application only whilst the application is running.

There is no doubt much more to say about the OO object concept however this brief overview should provide a starting point for our discussion and comparison. It should be clear by now that OO object classes and their instantiated objects are nothing much like their real world counterparts. The OO class *REPTILE* is an abstract representation, something more akin to "reptileness" than any individual reptile and any instantiated OO *REPTILE* object is a highly stylised, conceptual and extremely simplified model of a reptile, nothing like an actual reptile. Furthermore each instantiation of an OO *REPTILE* object, provided it carries the same attribute values, is logically identical to every other similarly defined *REPTILE* object. This of course can never be the case for actual reptiles of the same species.

These concepts and rules are general to all OO systems, nevertheless (perhaps surprisingly) consensus regarding the concept of the *object* within OO design is far from agreed upon. For example, after introducing some key controversies observed in the OO literature, Date and Darwen ask "So what exactly *is* an object? Is it a value? Is it a variable? Is it both? Is it something else entirely?" Due to the alleged ambiguity they go on to assert, "As a matter of fact, it is largely because of this confusion over what objects really are that we prefer ... not to use object terminology at all, except in a few very informal contexts." (2000, p10). Instead Date *et al*, prefer to rely upon a vocabulary that draws upon predicate logic and set theory to explicate their model of data representation.

Whilst the sheer variety and volume of OO programming and design literature available no doubt contribute to the confusion, even across a small sample we see inconsistent points of view emerge. Booch for example (1994, p35) coins a simple truism, "What we can agree upon is that the

concept of an object is central to anything object-oriented." Martin (1992, p241) perhaps in the tradition of Berkeley,[17] prefers, "An 'object' is anything to which a concept applies", and "A concept is an idea or notion we share that applies to certain objects in our awareness". On a first reading James Rumbaugh, one of the founding object-oriented methodologists, appears to get closer to the mark with, "We define an object as a concept, abstraction or thing with crisp boundaries and meaning for the problem at hand" (1991). Whilst this approach is certainly useful when it comes to defining an object class model for a well defined problem domain to be implemented as an OO application or database, it seems problematic as the basis for a universally favoured ontology. Indeed quite the reverse seems to be the case. Concepts and ideas seem to have vague rather than crisp boundaries. This is particularly the case when the applicability of a predicate to its subject is tolerant, for example when does a child cease to be a child and begin being an adult? The fact is most concepts do not have easily defined boundaries; reality is not crisp. On Rumbaugh's definition most of reality would be thrown in the too hard basket with regard to object modelling. One might be tempted to argue; all we need is a set of clear *semantical rules* that apply to our artificial OO style language and we could by and large eliminate such vagueness and ambiguity. However this approach seems to point to a required preciseness in meaning which naturally gives way to appeal to definitions and hence does not appear to solve our problem.

The confusion only deepens if we consider the historical discussion in philosophy regarding objects, classes and attributes. Whilst the Floridian notion of the information object posits attributes, collections of which constitute the nature or essence of an object and classes (following the OO paradigm) as "a named representation for an abstraction, where an abstraction is a named collection of attributes" (Floridi, 2002, 2004b, 2008),

[17] George Berkeley famously argued in his *Treatise Concerning the Principles of Human Knowledge* that material objects were merely ideas or concepts.

Quine on the other hand argues that the job of attributes can be adequately handled by classes. "Classes are on a par with attributes on the score of abstractness or universality, they serve the purposes of attributes so far as mathematics and certainly most of science are concerned; and they enjoy, unlike attributes, a crystal-clear identity concept" (Quine, 1969, p19). I accept (and imagine many others to agree) that any identity problem for concrete objects can be generally cleared up by appeal to concepts of reference but there seems to be a lack of an identity concept for attributes. In the case of OO programming however there is a clear identity link between attributes belonging to certain object classes. Attributes in this case can only be exposed via the methods associated with each class using pointer mechanisms and the unique object ID construct; however this tells us nothing about how attributes can be identified by non OO information objects. Thus talk of attributes in the OO sense seems to be quite different to the philosophical and ontological concept of attribute.

The point is this; in the OO world an object is a *referent*. Method invocation represents a layer of abstraction between some actual representation[18] that is implemented in some arbitrarily complex way on a computer disk (which is of no concern to an OO modeller) and the encapsulated object provides some kind of pointer reference to that data via the declared methods. Moreover object classes and their instantiated objects stand independent to reality and attempt simply to model reality in so far as *something* exists in time and space.

3.3 Floridi's Object Conception

The previous section introduced some of the controversies regarding Floridi's use of OO methodology. Furthermore we have outlined, albeit briefly the OO side of the story from a broadly definitional perspective, we

[18] By *actual representation* I mean the physical or internal codification or implementation of the data as it exists on disk.

now need to tease out how Floridi develops his concept of the informational object.

Floridi presents, via analogy, a picture of the information object using a well known icon, the chess piece, the pawn (2002, 2008). The identity of the pawn he argues is known not by its somewhat arbitrary properties as a physical object but rather by its function in the game of chess. In fact we could simply replace the pawn with any placeholder (Floridi suggests a cork) without any semantic loss. Alternatively we needn't have a physical placeholder at all – a pair of chess players with good enough memories, imagination and with a minimum of communication between the pair, could visualise an entire game thus suggesting that the real pawn is not really a material thing but a mental entity or an entity constituted by a bundle of properties (2008, p30). Unsurprisingly the pawn makes for a very good analogy in the application of OO theory, its role is simple and its functional boundaries are conceptually crisp. Indeed there are a great many OO programs that have been written to represent electronic versions of the game of chess, defining an object class with the appropriate attributes and methods that represent the physical pawn is relatively trivial.

Informational analysis according to Floridi relies upon another computational concept, that of levels of abstraction (LoA). Put simply we can discuss computational systems at differing levels of abstraction, high conceptual levels often involve abstract diagrammatic models, at a lower level we might imagine some written computer code or logical statements in SQL or the like, at lower levels still strings of scalar variables and combinations of bits and bytes. Thus according to Floridi, "The choice of LoA pre-determines the type and quantity of data that can be considered and hence the information that can be contained in the model." (2008, p16). The entire notion of the information object thus is couched within the levels of abstraction concept.

Yet this reductionist construal still seems odd to me. I wonder what utility the use of terminology at a particular LoA might have, such as the level whereby we construe objects as OO-like informational objects, other

than in merely naming the various concepts, objects or structures and their relationships or linkages across varied levels of abstraction in OO terminology. This OO-like terminology is usually reserved exclusively for the development of OO applications. The various other valid levels of abstraction within the OO paradigm include unified modelling language (UML) constructs such as use-case diagrams, system sequence diagrams and state machine diagrams and each of these levels develop structures and models that have very clear abstraction relationship rules linking them with domain class diagrams and their consequent computational implementations.[19] Nothing however in Floridi's literature suggests that informational objects exhibit similar relationships or rules across analogous levels of abstraction; this seems to be particularly the case between information objects and the real world objects that define them.

Thus, whilst I agree that a pawn (with all its requisite behaviours and attributes) may be imagined and that a pair of chess players with sufficient memories could somehow visualise an entire chess game, this is not the same thing as a pawn being represented in terms of OO theory, nor is the same as a physical pawn. Of course there are certainly *some* properties that each pawn representation shares but there are a great many differences also. Floridi however clearly takes a certain selected set of properties of the pawn quite seriously and these seem to be more definitive or significant to him. By way of example when I (mentally) imagine moving a pawn on a chess board I visually imagine a three dimensional space as well as a prototypical three dimensional image of a pawn. I fully expect this mental image to be quite different to a MRI scan of what is going on in my head at the time and we can be sure the MRI scan will be slightly different on each consequent imagining. It is not clear that such different construals could be explained

[19] I do not intend to rehearse the literature on the development of OO models and their abstraction relationships and rules across differing levels of abstraction however any review of UML OO modelling literature will suffice should the reader wish to read further. The UML Wikipedia page perhaps might be a good starting point.

away as being merely different levels of abstraction of the same thing, because they are not the same thing.

To further complicate the issue, consider the following; a Google image search of "pawn" returns a wide array of images of pawns not only within chess contexts but as an iconic image in its own right and these are all valid representations of pawns which bear little resemblance to the OO-like informationally austere pawn. There is no doubt at all that the pawn *can* be reduced to a set of behaviours and basic properties that can be modelled within an OO programming language but this (to me at least) only represents a minimally sufficient condition for an object to qualify as a pawn, in fact a *virtual pawn* as the current case may be.

The worry deepens when we consider imagined pawns. It is certainly not outside the realms of impossibility that mental representations of pawns between individuals vary greatly. Such representations of a "pawn" are not identical at all with the set of attributes, methods and other OO constructs that defines a pawn in an OO program. The same goes for physical pawns in physical chess sets. I cannot for example throw the information object that represents (or is) a pawn at my opponent when I'm losing thus suggesting that the set of pawns described purely in informational terms is not identical with the set of either physical pawns, nor with any mental representations of pawns whatever our normative notions about how pawns are meant to behave. Himma (2004) takes a similar tack as follows:

> Indeed, it is hard to see how a pawn could be identical with the information object that describes its properties and operations. If we conceive of the pawns as nothing more than information objects then all of the propositions in the set constituting the relevant information object are propositions that describe that set. Such an assumption would, of course render some of the sentences obviously false (information objects lack spatio-temporal location and hence can't be moved around) (p148)

This analysis I believe raises several questions. Firstly how do the information objects defined by things such as mental images relate if at all to all the other information objects that represent real world objects which by ostension we'd agree are part of the same class or set other than by loose consensus? They surely relate at some level since they are all supposed to represent the same thing – but this, I contend is largely folk talk. What Floridi seems to be talking about with regard to pawns is a relativity of identity of type. The OO method of defining a pawn *necessarily* relies upon pawns being of the same type and thus sharing some well defined properties. Floridi relies upon this methodology to clarify his information object and as such is quite serious about these particular properties, he takes it that these properties do exist and that they are constituent properties of pawns. Hence two different tokens, be it a cork or a carved piece of wood, have the same properties – such properties constitute the tokens identity transcending the material properties, that identity presumably in this case being pawnhood. It seems that what Floridi is talking about with regard to the information object (taking 'pawn' as the example), is something like the universal concept of pawnhood for which we already seem to have a theory albeit a controversial one.[20]

However it is clear that the cork pawn and the wooden pawn (as well as the imagined pawn and the OO pawn) whilst they share *some* properties, these properties are not identical across all pawns. Each set of properties relating to each pawn are particular to that pawn.

Another question raised is this; although the LoA approach is well proven within object oriented and computer systems design this is because there exists very explicit rules about how differing levels of abstraction are

[20] The problem of Universals was originally discussed by Plato and Aristotle and as a topic has captivated philosophy ever since. Universals are generally considered repeatable or recurrent abstract entities that can be instantiated in individual objects, classic examples are considered to be qualities shared by entities, such as two green chairs sharing the quality of "greenness" and "chairness".

linked to one another. Such rules about how real world objects and their informational counterparts are linked via LoA does not seem to be addressed by Floridi. Instead he offers us a conceptual discussion regarding ontological commitment and levels of abstraction (2008, p17).[21] The structural approach taken by Floridi works well for classes, and perhaps the abstract entities that are *information objects* but we initially learn about pawns not through abstract structural discourse but via ostension. While we might agree upon what qualifies as a pawn, my class of pawns, whilst probably exhibiting many identical features could also be quite different to yours.

There are two categories of problems with regard to the Floridian Information Object. I will outline these problems now. The first I call the *Methodological Problem.* Whilst Floridi draws heavily upon OO programming and design terminology in order to explicate his informational realism (as well as supporting the role of the information object within his IE), the object concept itself within OO programming or design, (issues of clarity aside) is heavily contextualised and specific. The rules linking different levels of abstraction from high level conceptual models to much lower level compiled object classes and programs and their consequent representation at the disk level are very explicit. To extract the OO object concept from its own theoretical environment leaves us wondering as to the explanatory value such discourse could have outside metaphor and analogy. The application of OO concepts is specific to their domain and use of them outside this domain requires considerable ad-hoc addition and modification which mostly ends up in confusion and misunderstanding. This is even the case within computing circles, a clear example of which can be seen in

[21] Floridi does attempt to address ontological commitment to different LoA by attempting to reconcile epistemic and ontological structural realism. However in this paper I am concerned with the relationships between information objects, OO concepts and real world objects, as such this is outside the scope of this particular paper.

recent attempts to apply object oriented concepts to the relational model of data.[22]

The object concept seems to have such a wide range of applicability that it ends up somewhat ambiguous. We have noted that Date and Darwen (2000) dispense with OO terminology in favour of vocabulary based on set theory and predicate logic in their discussions on data representation. Certainly on Rumbaugh's description it seems difficult to understand how OO type objects could represent anything within the ranges of our normal understanding of language.

The second difficulty I am calling the *Identity Problem*. This issue relates to where and how a Floridian information object's data members are represented or manipulated. While Floridi bases his notion of the information object on OO concepts and terminology his goal for the information object is clearly quite different to that of an OO application or data model designer. Floridi certainly does not want his objects to be referents in the same way an OO object is. The issue seems to be a problem related to identity or correspondence relations between the abstract information object and its real world counterparts.

The methodological problem is only a problem when OO concepts are used outside an OO context. OO programming for all intent and purposes, 'works' and all the philosophical anxiety in the world over just what an *object* might be or whether it accurately addresses ontological problems doesn't really matter, certainly not to the OO designer or programmer who is simply solving what is usually an information management problem using a particular development/design environment. In other words OO theory or at least parts of it, is instrumentally reliable with regard to the creation of 'working' *Object Oriented* programs and their corresponding object structures. The question of whether they are (or not) ever directly representative or answerable to any external truth about the real world is not

[22] Date and Darwen (2000, p371) call this a "great blunder" arguing that it both dilutes OO concepts and undermines the conceptual integrity of the relational model.

at issue. Object oriented programs (or OO databases) are structured collections of relatively simple facts represented by sets of values and governed behaviourally by simple computational procedures and functions. However the truth of such facts, or more precisely the correctness of the data representations depends entirely upon whether or not the values are *(a)* consistent with the rules (usually the 'business rules') upon which the OO application has been designed, and *(b)* correspondent with the external environment, that is, not erroneous. Of course *a* is exclusively the responsibility of the OO designer, whereby *b* is almost certainly contingent.

Whilst there are philosophical disagreements regarding a precise definition of just what constitutes an object, the points above and the distinction between *a* and *b* is generally accepted within the OO community. An OO model is designed at a conceptual design level by an OO designer conceptualising object classes which are somewhat representative of the external (data) environment or as it is sometimes termed, 'the problem domain'. Moving through the requisite LoA appropriate to OO modelling, the OO model is implemented internally at a logical level as a set of definable structures specific to a particular vendor's database or programmatic development environment. The resulting OO application or database by and large serves some clear delimited business, engineering or scientific function.

Of course Floridi is *not* saying OO objects are the *same* as informational objects, the kind that he supposes reality is comprised of. He is not trying to directly co-opt the OO object concept into the service of ontology. Yet he does borrow and rely upon OO terminology and unless clarified it seems reasonable to argue that the information object could only suffer the same ambiguous issues that the OO object endures. Thus it is not at all clear that OO terminology can play any pragmatic role in explicating a universal ontology other than providing some kind of loose model of the ontological components of his informational realism.

The Identity Problem is a more complex philosophical problem and concerns the issue of how we can have knowledge of abstract entities, the

role their composite attributes play as well as how and what such components reference. We shall consider this problem as part of the next section which considers the philosophical position with regard to abstract objects and some objections.

3.4 Abstract Objects and their Problems

Already stated is the notion that philosophy views 'abstract objects' as those which are non-spatiotemporal in nature. It also often considers abstract objects to be causally inert, that is they generally have no direct ability to affect the 'real world'[23]. Whilst there are varied theories about the nature of abstract objects, this thesis proceeds with the general assumption that all objects fall into two categories, either those which are concrete or those which are abstract. Of course there is much debate surrounding the concrete/abstraction distinction however following a general approach let us consider for the most part concrete objects to be spatio-temporally extended, (like tables, chairs and mountains), and abstract objects not (like sets, prime numbers, predicates, fictional characters and informational objects). Further under the general assumption, no object can straddle the distinction. Yet whilst an information object itself is abstract it can represent both concrete and abstract objects. Thus an information object is equally capable of representing Wonderland's Alice as it is Aoraki.[24]

The information object however seems to exhibit a special kind of abstractness. As discussed in the previous section, Floridi's use of OO theory to "clarify" the notion of the information object leads us to conclude that he must think there can be groups of things of the same *type* (pawns, for example), and these things are all members of the same class (which is

[23] In particular the causal inefficacy argument, that is, an object is abstract (if and) only if it is causally inefficacious. This argument forms part of what Lewis (1986) calls "the way of negation" in that abstract objects lack certain features that are evident in concrete objects (Rosen, 2012)

[24] Aoraki is the indigenous Maori name for Mt Cook, New Zealand's highest mountain.

something like an OO class), and that this sameness is taken in a strict sense, that is these things share some important selected identical properties. This must be the case since OO programming requires strict relations of identity with regard to object construction via class methods. Whilst this notion of abstraction in the OO world is taken for granted in philosophy it does not seems to be quite so simple. I tend to think (as per the illustrated case of pawns above) that notions of identity across groups of resembling objects are not strict at all, and whilst it is difficult to avoid talk of like-properties across individuals or particular objects such talk of similarity is largely an artefact of our language, culture and some innate human predisposition for organising and describing our world.

In philosophy, one form of *nominalism* entails the rejection of the notion of strict identity across objects. This somewhat extreme form of nominalism amounts to a rejection of universals. For Floridi's notion of the information object to work, particularly with regard to the OO concept of object it seems he needs to be a realist about universals, otherwise we cannot form a concept of pawnhood representing the identical properties shared across the class of pawns. Another form of nominalism rejects abstract entities just because the nominalist generally disagrees that non-spatiotemporal, causally inert objects exist at all. Part of the reason for this rejection is as follows. In science causality and explanation are closely linked and whilst there is an entire domain in philosophy devoted to this topic we will simply say that in many cases to explain a fact is to identify its cause.[25] In other words, the intellectual understanding we have about our world and all its systems and organisms and so on is often summarised in our scientific explanations, these explanations are often causal by nature (*cf.* Salmon, 1998). The corollary to this is that it is difficult to explain how we

[25] Generally the claim in this context refers to non-trivial or non-tautological facts. For example, 1+1=2 and, 'unmarried men are bachelors' qualify as facts but they have no cause. (I thank Morgan Luck for pointing this out to me). We might note here as a by-line that Floridi following including many others doubts the informative nature of tautologies or necessary truths.

could have knowledge or understanding about an object if there are no causal relations between it and us. Such causal relations could simply be the reflection of light from the surface of some such object impacting upon sensory apparatus (in that case our eyes). Quine argues "Science itself teaches us that there is no clairvoyance; that the only information that can reach our sensory surfaces from external objects must be limited to two dimensional optical projections and various impacts of air waves on the eardrums and some gaseous reactions in the nasal passages and a few kindred odds and ends" (1974, p2). All explanations no matter how complex or convoluted ultimately sheet back to these basic causal interactions.

Goodman and Quine do espouse a form of nominalism based upon a simple philosophical intuition, "What seems to be the most natural principle for abstracting classes or properties leads to paradoxes. Escape from these paradoxes can apparently be effected only by recourse to alternative rules whose artificiality and arbitrariness arouse suspicion that we are lost in a world of make believe" (1947, p105). Compared to the OO object conception whereby objects are derived from declared object classes the Floridian conception falls prey to at least one causally related paradox we alluded to above. That is, from an OO point of view an object *qua* referent points (literally and programmatically) to a physically implemented value (or variable). There is a logical difference between the appearance of that value exposed by the objects methods and the encoded version that is used to represent its appearance (via various layers of abstraction) internally (on a computer disk). But, and this is the rub, while the referent doesn't point to the external thing it attempts to model, in the case of our example above, the OO pawn, it does point literally to some arbitrarily complex piece of data, and that is what the referent is physically referencing. Furthermore there are direct, explicit, and traceable causal links between the various levels of abstraction.

But what does an information object point to? Clearly the intention for it is that it either points to or emerges from the thing that it references.

However in the case of an information object there is no implementable (as there is with an OO object) physical layer of abstraction which provides mappings between logical representations at higher levels of abstraction and physical, or the actual material data representation. Thus in the case of an information object the abstraction mapping process is either a behavioural or a cognitive process, relating some vague sensory input between the extant object and the agent or organisms experience of it. This experience will be different for each organism. Humans experience exposure to the sun for example in an entirely different way to the way algae *experience* it, if we can even call it that.

3.5 Natural Entities, Artifacts and Technology

There is a final related distinction I will make between OO objects and Floridian information objects. In this section I draw upon an argument made by Deborah G. Johnson (2006) regarding the distinction between natural phenomena or natural entities (which according to Floridi can be explained as 'dynamically interacting informational entities') and human-made entities or what are often termed artifacts. The section continues on from previous discussion in developing the argument that informational objects, the kind employed in Floridi's informational realism and ethics cannot be anything like object oriented objects. I first outline Johnson's argument and then apply it to my general discussion.

Johnson's paper develops the thesis that computer systems can be moral entities but not moral agents. Although I don't intend to discuss the main thrust and conclusions of Johnson's argument regarding the moral status of computer systems in this paper, whilst developing her argument she highlights two important distinctions which *are* salient to the present discussion. The first distinction is between artifacts or what would normally be considered human-made entities and those entities which occur naturally. A second distinction is made between artifacts and technology. Although she concedes that these distinctions are inherently problematic in the sense

that there is no sharp boundary they are nevertheless significant. Any "rejection or re-definition of these distinctions obfuscates and undermines the meaning and significance of claims about morality, technology and computing." (2006, p196) Johnson asserts.

The challenges are illustrated as follows; a stick used by a tribesman as a spear to hunt an animal for instance is a naturally occurring object which has been used as a tool. Thus the stick whilst seen as a natural entity is also utilised as a form of technology. Newer technologies such as genetic modification and nano-molecular technology may also appear to blur the line between nature and technology. Just which parts are naturally occurring parts and which are human-made artifactual parts may be difficult to assess. The only difference Johnson notes, between biotechnology and other types of technology such as computer systems and the like is the extent to which they manipulate, or the level at which the manipulation of nature occurs (2006, p196-197). Thus while challenges can be made to the distinctions Johnson argues it doesn't mean the distinction is incoherent or not tenable. On another level, according to Johnson, these distinctions allow us to make sense of the kinds of questions related to what kind of effect human behaviour has on the planet versus something independent of human behaviour – nature.

Although absolute definitions of *technology* are problematic, referencing Heidegger, Johnson attempts to avoid some of the debate by simply arguing that technology is a contrivance and inherently refers to human-made things. This of course includes computer systems at both logical and physical levels. While Johnson confines the term artifact to physical objects,[26] I think however we can extend Johnson's definition of artifact. Computer systems are made up of both physical and logical components and the logical components are just as artifactual in nature as the physical or material

[26] Johnson, (2006, p197) "A common way of thinking about technology – perhaps the layperson's way – is to think that it is physical or material objects. I will use the *artifact* to refer to the physical object"

components. This is clear since there are explicit design and development processes associated with the creation of all logical components of computer systems. The resulting objects have unambiguous designed functions, can be identified and map to specific arbitrarily complex data representations on disk or in the memory of the computer, they are indeed virtual artifacts.

Further Johnson points out that "technology is a combination of artifacts, social practices, social relationships and systems of knowledge... sometimes called socio-technical systems" (2006, p197) Thus artifacts (as components of socio-technical systems) cannot exist according to Johnson without systems of knowledge, social practices and human interactions and relationships. Artifacts are created, distributed, utilised and have meaning only within the context of human social activity (*ibid*). Whilst there are differences between logical and physical artifacts they are both clearly the result of human behaviour and decision processes.

Object oriented objects are artifacts. They are nothing but the result of an explicit design processes by humans. On the other hand informational objects seem to exhibit some of the characteristics evident in both natural entities and human made artifacts. Information objects however may be just a way of talking about and interpreting the things and events around us, mental entities, indeed this is a view I develop further in the next chapter.

This points toward the contingency of information objects since the existence of such objects relies upon some vague correspondence between the object's internal structures, whether we describe such structures as "attributes" and "methods" and so on, and their relationship with the real world physical entity (table, chair, mountain etc.) they seek to represent. We should point out that Floridi's construal of the information object seems to differ from the role classes play in a similar discourse, where the information object seems to pick out individuals the role of classes, to use Plato's metaphor are our attempt to carve the beast of reality along it's joints.

The problem becomes more complex when we include informational objects *qua* abstract entities into the mix, the kind that nominalism would typically reject. For example, there could be any number of imaginable

information objects representing anything for any possible interpretation. To paraphrase Hayaki (2006, p81) who considers similar problems associated with contingent objects, we are not counting actual possible physical objects; we are counting the ways in which an object might be represented (by an information object) by any possible agent.

The information object suffers from an identity crisis. Following Johnson, in order to identify an information object as an independent entity requires us to separate the object from its context however in this process we are extracting it from the context that gives it its meaning and function. This appears to be a problem the analysis of information *qua* information suffers from in general.

Thus whilst I agree it does not make sense to ask the question "Where are these information objects you talk of Luciano?" abstract objects such as information objects do not exist in space, I do think it legitimate to ask "how do every day (concrete) objects map to their information object counterparts?" This question is answered in the OO case since there are clear and explicit abstraction relations between different levels of the model. This level of detail seems obscure with regard to informational objects. I hope I have made it clear that OO system entities are never isomorphic with any kind of external reality (but informational objects are supposed to be). OO objects are merely a more or less accurate *model* of reality. Indeed OO models are pragmatic by nature – the goal is to solve a business, engineering or scientific problem, that is, a problem that can be adequately solved with an OO application. The Floridian account however seems to suggest an object *qua* information object does indeed reference the real world object it purports to represent but just how this works is not explained. Floridi argues, "the ultimate nature of reality is informational, that is, it makes sense to adopt a level of abstraction at which our mind-independent reality is constituted by relata that are neither substantial nor material (they might well be but we have no reasons to suppose them to be so) but informational." (2004b, p5).

The problem is not that informational objects are not materially evident, neither are classes but they are an essential part of the natural sciences. Quine's nominalism for instance admits of such abstract objects such as classes, numbers and sets and the like into his physicalist ontology because science simply could not proceed without them – this in essence reveals Quine's pragmatism. However is such a comparison any argument for the admission of Floridian information objects as a legitimate ontological category? While most of us would agree that there are concrete objects in the world that are both substantial and material, the question remains as to how seemingly mind-independent, non-material, abstract relata causally interact with the material world. For this gentle nominalist at least Quine's (as well as Occam's) suspicions are aroused.

3.6 Conclusion

Whilst OO programming and design was in its infancy Gaifman, (1975, p329) a philosopher, summed up our predicament, "objects are notoriously theory-laden; an informative discussion of objects of this or that kind presupposes already a whole conceptual scheme". The conceptual scheme which supports a notion of the information object isn't surprising, we live in the information age – our economy, social structure and culture are virtually defined by information and information technology. This technology is implemented in computer systems that are developed using object oriented development environments such as Java and C++. These environments are supported by a popular design methodology Unified Modelling Language (UML) which explicitly lays out rules for mapping high level conceptual models to implementable object oriented class models and eventually computer code. The notion of the most basic primitive ontological category being informational by nature is somewhat attractive; it fits with our current and popular world view, it appeals to our natural desire to impose order upon our world. Having said that implementation of levels of abstraction within computing environments are explicit and functional. A fundamental

and necessary property of OO objects is that they are referents and this referential property is implemented directly via the implementation of LoA between various conceptual through physical layers ending with bits mapped to hard disk or memory addresses. Identity relations with regard to OO objects are explicit and clear. OO objects of the same type, sharing the same attributes, are logically (and digitally) identical. It isn't the case that nature operates in a way anything like an OO application. Floridi doesn't say as much but this does raise the question as to what extent a seemingly primitive natural kind, the information object, can be made clear with appeals to human made artifacts. Careful derivation of semantic information from objects and states of affairs led to the development of the OO model. It doesn't then make sense that that OO model can clarify nature in any way outside loose metaphor.

What is not clear is why Floridi would explicitly choose the currently popular programming paradigm. It could merely be that object oriented programming and modelling literature, boasting high levels of semantic representation, seemed to offer something contemporary and novel. However, there are a variety of other data modelling approaches within computing that while being conceptually more austere may have presented fewer difficulties. The abstract concept of the bare data type itself may have proved a smaller target. Indeed this seems to be the approach Armstrong (1989) takes and is the topic for discussion in the chapter that immediately follows.

I have argued in this chapter that the inherent ambiguities of OO theory end up being transitive, that is, the problems we see in defining just what an object is and what its behaviour ought to be within OO theory are equally problematic for any ontological theory based upon the paradigm. The point is we can still hold on to classes, properties and relations, essential components for any metaphysical framework without having to talk about object oriented objects. Bas Van Fraassen advises us that "Theories with some degree of sophistication always carry some metaphysical baggage" (1980, p68) just like hidden variables theories in quantum physics the hope

is that carrying the baggage will eventually pay off. What I have done is presented an argument as to how informational objects, if we at least entertain the possibility of such things, do not seem to be much like OO objects.

The ontological argument for such objects and hence informational realism is ultimately metaphysical by nature but Floridi adopts a structuralist methodology thus what we can salvage with some conviction is the properties and relations that are part of these postulated entities. Whether these are captured by universals (as in the property of *pawnhood*) or are ultimately unique to instantiated objects (what Armstrong (1989) called *particulars*) is an issue examined in the very next chapter. It should be evident (with respect to my own declared nominalistic tendencies) that my analysis of these problems has been cautiously parsimonious with regard to postulated entities. Nevertheless the topic of *informational realism* is certainly an intellectually fascinating one for anyone interested in what D.C. Williams termed "grubbing around in the roots of being." (1966).

4 Informational Realism and Metaphysics

A shorter version of this chapter appears in the proceedings 2011, *The Computational Turn: Past, Presents, Futures?,* International Association of Computing and Philosophy Conference, Aarhus University, Denmark.

4.1 Introduction

In the previous chapter Floridi's information object came under criticial scrutiny. It was concluded that if we are to admit informational objects into a metaphysical ontology they cannot be much like an object oriented object. While Floridi argues the comparison is only meant to serve as a useful way of clarifying the information object he does extensively apply OO theory and terminology to the concept. At best this strategy obfuscates the concept. This chapter seeks to provide an alternative, more robust method of explaining objects and states of affairs from an informational level of abstraction.

Sometimes, seemingly simple questions end up being the most puzzling. For example, the question, *what is information?* elicits a few billion more results in a Google search than there are individuals on the planet. This may well tell us nothing much more than whatever it is there is a lot of it about. It certainly indicates there is plenty of room for debate.

The question 'what is information?' is essentially a metaphysical one, that is, one that neither science nor dictionary definitions (let alone Google) offer an adequate answer for. An educated thinker, perhaps one educated in the information sciences, may appeal to a set of terms and concepts in support of their answer, such as; data, properties, classes, relations, attributes, entities and so on. A more philosophically inclined reply may point to distinctions between the physical and material worlds, or the abstract and concrete, to questions of veridicality, whether such a thing is

mind dependent or occurring naturally in the world, and perhaps enquiring to its causal properties. [27] Both approaches and their associated terminologies form the starting point for what has come to be called *informational realism*.

To begin the metaphysical story we might frame our problem by asking, can the world be fundamentally informational by nature? On first suspicions a sceptic might persuade us that in consideration of such a question we must be aware that any answer may exhibit *confirmation bias*. That is, swayed by the fact that we live in the so called *Information Age*, it is hardly surprising that we might attempt to interpret the world in such a fashion. However our protagonist in return might wish to offer some empirical examples in support of her argument. Patricia Churchland begins her 1986 text Neurophilosophy with one such simple and compelling example:

> Squirming out from the primordial ooze, our evolutionary ancestors harboured within themselves a perfectly astounding invention – the excitable cell. ... From the very beginning, mobile creatures whose excitable cells were capable of conveying information about conditions outside the body had a survival advantage over those whose movements were independent of whatever was going on outside... the organism that flees in absence of predators and feeds willy-nilly is doomed to be prey for those more lucky organisms fitted with cells coordinating representations of the world with movement in the world. (*abridged*, 1986, p1)

Central to Churchland's thesis is not just the profound ability of minds (our own, as well as those of other species) to so richly represent the

[27] The relationship between causation and information is examined in more detail in Chapter 6.

external world but the seemingly simple ability of the non-sentient "excitable cell" to fundamentally interpret and respond to basic requirements of life such as feeding, fleeing and reproducing. At much higher cognitive levels we tend to call this perception but low level examples such as Churchland's seem to illustrate the fact that information exchange or transmission is critical to the mechanisms by which excitable cells operate. This example certainly seems to provide some empirical support for a kind of informational realism from an organic or perhaps a system-theoretic point of view. By system here I mean a complex of various elements and relationships and of having some kind of behaviour which involves inputs and outputs and some form of processing – that is, information processing. This kind of analysis is a popular methodology in many of our attempts to understand how things work. But what about the amorphous welter that constitutes the non-organic universe? Can that also be informational by nature? Or are we to explain it by some other method?

Perhaps we should start our analysis with a more general question; if we are to propose a metaphysical system that has information or informational objects as a central focus how should we formally specify such a system? Such a specification is almost certainly going to involve a discussion regarding the postulated objects, properties, related entities as well as any relationships that hold between such things. We would hope that our specification transcends the current zeitgeist, that is, it should not be swayed by our current "information age" biases. Just because we live in the so called "information age" it does not necessarily follow that the universe is informational by nature.

Metaphysics however comes in different forms, informational realism is one recent attempt but there are other equally compelling theories. David Armstrong (2010) for example forwards a view that differs from the informational realist approach. His view proposes that the most fundamental structure can be explained by what he calls "states of affairs". We shall examine Armstrong's view in more detail soon. However, briefly, a state of affairs is simply the instantiation of a particular property in an

object. In its simplest form this is a two-term relation between an object and the property (or properties) that particular object has. While there are other metaphysical theories, I have chosen to focus upon Armstrong's state of affairs approach. Armstrong's approach is both compelling and contemporary, and hence offers an interesting comparison to Floridi's informational realism.

An obvious question to begin with might be this; which version is correct? Should we pick Armstrong's states of affairs view, or the idea that space-time is fundamentally informational by nature comprising dynamically interacting informational objects. This chapter examines both of these competing, contemporary views with the aim of striking some rapprochement. Firstly (section 4.2) we introduce informational realism as a form of structural realism. We quickly find common ground with Armstrong's view whereby any entities we do postulate must be causal in nature. This view squares with my analysis of information and its relation with cause outlined in chapter 6 of this thesis. Section 4.3 looks at the similarities (and differences) between these two competing metaphysical theories, in particular the reification of Floridi's concept of the informational object and by contrast Armstrong's purely semantic conception of the "state of affairs". Section 4.4 looks at the problem of identity for informational realism and finally 4.5 considers other modalities such as imaginary and probable informational objects.

4.2 Structuralism as a starting point

The most contemporary and comprehensive account of *Informational Realism* is that forwarded by Luciano Floridi (2004, 2008). Following a variety of thinkers that are sometimes called "structural realists" (Maxwell, 1970., Worrall, 1989., Ladyman, 1998) Floridi argues that the universe can indeed be interpreted from a structural perspective and the most accurate way of interpreting these structures is as *dynamically interacting*

informational objects. The general idea of structuralism is that the only knowledge we can have of the world is structural knowledge.

There are two primary forms of structuralism and while a full discussion is beyond the scope of this paper we briefly describe the two forms as follows: *Epistemic Structuralism* is the idea that the objects themselves, while acknowledging that such entities exist, are inherently unknowable. So while we can know all about the properties and relations of particular objects, that is in fact all we can know. On the other hand *Ontic Structuralism* is a somewhat more extreme view holding that there are literally no objects, only properties and relations, and therefore *structure* is all that there is. For example, if we were able to remove all the individual properties of a particular object, say a solid cube, or a chess piece, there would literally be nothing left according to the ontic structuralist.

Van Fraassen (2006) unsurprisingly calls this radical structuralism. Some of the arguments forwarded for ontic structuralism draw upon model-theoretic methods in science, in particular, idealisations in physics. As such the view has garnered some support with philosophers and physicists. Floridi argues the two forms are just two interpretations analysed at different levels of abstraction (LoA).

There is another contender in the ontological arena. David Armstrong (1998, 2010) begins his argument with an assumption, in fact this is the same realist assumption I would expect Floridi to begin with – "all that exists is the space-time world, the physical world as we say." (2010, p1). Whilst this is a hypothesis that we are all intuitively comfortable with, once we start peeling away the layers of what this actually means we are usually forced to include the existence of more difficult and abstract hypothetical notions; notions such as classes, properties, relations and the like. Trying to explain the nature of things in the world without appeal to such concepts appears almost impossible. It is often these abstract notions that notoriously muddy the metaphysical waters.

Early on in his discussion Armstrong makes an important distinction. A restriction of what it means to exist in the space time world entails the

rejection of what philosophers term "abstract objects" (*ibid*). Armstrong postulates that it is difficult to determine what kind of *causal* role an abstract entity has. Drawing upon Graham Oddie's 1982 work Armstrong argues that any entity we postulate must have some kind of causal role in the world. Oddie termed this the Eleatic Principle. The argument roughly states that if an entity plays no causal role at all it is difficult to forward a case for its existence. This idea originally comes from Plato's Sophist where the "Eleatic Stranger" argues that "the mark of being" is inherent in power – the power to influence or change something. Entities that lack causal power arguably ought not to be included in any ontology.

It is not entirely clear whether Floridi's ontology extends to such entities. However from the structuralist's point of view Armstrong's approach, and earlier Oddie, appear reconcilable with Floridi's ontology. Esfeld (2009) has forwarded an argument in support of Armstrong arguing that fundamental physical structures possess causal essence, or powers. Further Esfeld argues that structures themselves are indeed the basis for scientific realism, supported by physics and "provide a complete and coherent view of the world" (2009, p179).

To exclude abstract entities from our ontology however does require some explanation as to how we deal with such things as classes. Quine (1953) argues, such things are merely a "convenient myth". Classes it should be remembered are not actually instantiated instead it is the *particular* objects of a certain class that are. The decision upon which *class* a particular object belongs is the job of experts in the particular area. For example we might rely upon an ornithologist to determine which species of Albatross a particular instance belongs to. Thus the role classes play is in making some sense out of what Quine calls the conceptual scheme of physical objects. They help simplify our categorisation efforts.

While I do not intend to begin an extended discussion of classes here, we can say with some confidence that classes are abstract entities, as such they in of themselves have no causal efficacy. Rather it is the instantiated objects of certain classes that carry causal import. And so states of affairs or

informational objects are certainly not classes. Armstrong argues that while states of affairs have attributes (or properties) as their constituents they are particulars. "The same goes for classes of *first-order* particulars. Such a class, whether or not it is a singleton, is not repeatable, it is not predicable of many" (Armstrong, 1998, p188).

Floridi's account of classes, following a popular concept in computer programming is an *object oriented* version, "A class is a named representation for an abstraction, where an abstraction is a named collection of attributes and behaviour relevant to modelling a given entity for some particular purpose at a certain LoA." (Floridi, 2008, p30). While Floridi does not provide any discussion regarding the instantiation of particulars of given classes, my previous discussion is compatible with object oriented accounts and by and large follows the accounts given by Armstrong and Quine amongst others.

The version of structural realism to which Floridi appeals holds that the structures are physical, hence concrete and can be described in terms of networks of qualitative physical relations among objects – the objects themselves are simply what stands in those relations. Furthermore the objects do not possess any intrinsic identity distinct from the relations within which they stand (Floridi, 2008., Esfeld, 2009). Floridi goes on to argue, a way of making sense of these structures "is as *informational objects*, that is, cohering clusters of data, not in the alphanumeric sense of the word, but in an equally common sense of differences *de re*, i.e. mind-independent concrete points of lack of uniformity." (2008, p26). Thus with the backdrop of an already respectable metaphysical theory, structuralism, Floridi begins to build the picture of the information object.

An interesting parallel with the duality issue in the philosophy of computation can be drawn here. The dualism in computer science relates to the apparent distinction between abstract and physical computational entities. The typical way to explain the distinction is in identifying hardware and physical devices at the physical end of the spectrum, and software, that is, algorithms, compilers, code and the like as being abstract.

There are however objections to the distinction with some arguing there is no real distinction to be made (Suber, 1988., Moor 1978). Colburn (2000) however, argues that software has a dual nature and Turner (2013) argues that this duality seems to apply to all computational artifacts.

While Floridi postulates the information object as a physical entity, for any instantiated object its properties and relations appear to be abstract concepts. Certainly the concept of class is an abstraction and this is universally agreed upon within computer science. Following Turner (2013) it seems that informational objects, as Floridi wants us conceive of them have both abstract and physical guises.

There is perhaps another way the duality can be addressed. In section 2.7 of this thesis Floridi's method of levels of abstraction is explained. The method allows for finer distinctions to be made than the rather blunt instrument offered by a simple abstract/physical distinction. Appropriate analysis at both physical and logical or abstract levels depending upon analytical requirements can be made. The issue for this particular chapter is to decide whether or not the informational object is ontologically significant. While I discuss in detail the nature of Floridi's information object in chapter 3 the following section considers the information object within the context of structuralism as being a *bearer* of properties.

4.3 Informational Objects, Universals and Bearers.

The idea of the information object as being ontologically significant has gained traction recently not only in computer programming circles but also philosophically. We could attribute this newfound popularity, particularly with regard to philosophical interpretations, with the fact that we live in the so called information age. We, at least in the developed world, view the world through information-coloured spectacles these days. Adding some substance to this claim is the fact that our information systems are designed and developed using fashionable object oriented (OO) methodologies. Information modelling is now the accepted process by which facts or

propositions, the sentences that demarcate the various states of affairs and "things" of which the modeller is interested, are defined via "object class" structures. Such structures in turn represent various properties, behaviour and relata.

The information object in this sense is an intuitively fitting and elegant way of representing the problems we attempt to solve via computational means. OO design and development is "instrumentally reliable" – it works albeit within a limited scope. The majority of modern implemented information technologies across the entire gamut of industries and applications typically employ object oriented approaches. The focus has shifted from procedural algorithmic processing to an object driven methodology and as such states of affairs and "things" are abstractly modelled as self-contained (encapsulated) object structures, responsible for their own identity, relations, properties, states and behavioural rules. It's perhaps not surprising then that we might ponder; could the universe be interpreted and/or represented in such a way?

From a wider perspective what is often termed the computational turn has given rise to the informational object concept central to and emerging as fundamental in an informational ontology developed primarily by Luciano Floridi (2004, 2008). The concept is important for Floridi since the information object plays a role central to his Information Ethics (IE) and Informational Realism (IR). But more than this, the idea of the information object seems to offer new ways of understanding epistemology, semantics, scientific explanation, and ethics. Floridi has developed a detailed picture of the information object (or entity as he sometimes calls it)[28] employing Object Oriented programming and design methods and theories to clarify the concept.

Floridi wants us to consider the notion of the information object as analogous to the OO conception of an object however, I argue in Chapter 3

[28] In more recent work Floridi (2012) has called the information object an *i*Object.

for a variety of reasons that information objects, certainly with respect to developing a metaphysical theory about reality, do not seem to be much like OO objects and hence OO as a methodology is the wrong way to describe informational objects.

Arguably the most significant difference is that OO objects act unequivocally as referents. They are referents in two ways, firstly they are logical representations of arbitrarily complex data stored physically on computer disks and secondly they refer to, at a logical level, the real world objects they purport to model. This clearly is not the kind of object Floridi wants informational objects to be. Thus while there may be some similarity between OO objects and Floridi's conception of the information object I argue (Chp 3) that the similarity is more harmful to the idea of the information object holding any independent ontological status or existing independently as a particular category. The similarity is that both object concepts are conceptual and referential and therefore abstract by nature. This is a problem for Floridi since he clearly wants to confer a stronger ontological status upon the information object.

Problems arise if the information object is indeed conceptual. Following Lauden (1977, p48) such conceptual postulates can have no existence independent of the theories which postulate them. Nevertheless the concept of an information object is certainly a convenient and relatively intuitive way of bundling up constituent properties and relations belonging to the particular in question. Those properties and relations are in fact what philosophy sometimes calls *universals* and it is each particular (a distinct informational object) that instantiate those universals. Universals themselves are the constituents of information objects shared across many objects.[29] It is the universals that help us categorise and classify objects, such as all things that are red, or all objects weighing 25kg.

[29] The denial of the existence of universals (nominalism) holds, following John Locke, that "all things that exist are only particulars" (individual objects) and as such there are no strict identities across different objects, that is there are no universals.

To be a realist about universals is simply to take talk of such properties (sometimes called attributes) seriously. We might say that such properties of objects actually do exist in the world, and on this point it seems Armstrong and Floridi agree. As stated earlier Floridi avoids direct discussion of universals. This could be because he wishes to avoid any controversy associated with the identity of properties. Recall, on Floridi's account, objects do not possess any intrinsic identity distinct from the relations within which they stand. Where Armstrong and Floridi differ is that Floridi's concept of the information object is based upon his notion of information representing a *difference*. He asks as to imagine a universe comprising simply of a two dimensional white surface, and nothing else. "So far nobody has been able to solve the problem of finding a principle of identification for something like that, by definition is undifferentiated" (2008, p40). We can't treat universals as individuals because there are no criteria to identify them as such. Nevertheless it is the constituent properties, that is the relations and structural properties of the object that are knowable but only as they stand in relation to the object in question.

Thus we come to an important distinction between Armstrong and Floridi's metaphysics. The information object for Floridi is essentially the *bearer* of its constituent structural properties and relations. Armstrong however questions the need to recognise a bearer as an independent category of particulars. He argues that whilst properties and relations can be known "the bearer of properties and relations, it is alleged, cannot be known. Why then postulate a bearer?" (Armstrong, 1998, p95). This line of reasoning follows the structuralist's approach somewhat in that all structuralists want to admit into their ontology is structure. The postulation of the bearer for Armstrong lacks ontological and epistemic economy (*ibid*). This raises the question, is the Floridian information object the same kind of thing Armstrong terms a bearer? Armstrong's properties and relations compare favourably with the structuralist's elements mostly of the same

name, that is classes, objects, properties, relations and the like. Thus while it is convenient to talk of objects, structuralism holds that such entities are ontologically secondary to relations and properties which are primary. "The relations bear all the ontological weight: Objects are literally constituted by the relations in which they stand" (Esfeld & Lam, 2010, p4). It's also important to note at this point, I suspect with desire toward ontological economy, that some structuralists wish only to talk of relations. For many structuralists, attributes and properties simply reduce to relations. Stachel for example begins his explanation of structural realism with "Structure refers to some set of relations between the things or entities that they relate – the relata", in its strictest sense "the totality of structural relations reduces to relations between relations" (Statchel, 2006, p2). However Floridi (2004, 2008), Ladyman (1998) Ladyman and French (2003) all develop arguments whereby properties play roles, it seems, independently to relations. Certainly from a computational perspective properties and relations play quite different roles and it's unsurprising that Floridi, given his fondness for OO programming terminology, sticks to this ontology.

Getting back to the issue of the bearer, let us consider Armstrong's epistemological difficulty with the bearer, or its counterpart in informational realism, the information object. While we acknowledge that it is convenient to be able to talk of objects *qua* information objects and we also note that our perceptual experiences intuitively point us in the direction of an acknowledgement of objects (Armstrong uses the term particulars). Perception, Armstrong argues, is just the fact that we recognise that 'something is the case'. The question of veridicality is independent since we know that our perceptions can be overruled by higher cognitive faculties, thus accounting for all manner of phenomena from optical illusions to mirages. Armstrong adds there is no simple relation between the content of our perception and the acquiring of beliefs. The content is always that something is the case. That is, what is perceived is always a particular *state of affairs* external to the mind (Armstrong, 1998, p96) and this is typically propositional in nature. Basic perceptions Armstrong argues are simply

propositions taking the form 'this has certain properties and/or this has certain relations to that'. It may be an object having a certain colour and shape and being in a certain spatial position with regard to other objects (*ibid*). Armstrong concludes the section by arguing it is not obvious that all that is given to us in perception is properties and relations, certainly our intuitions point towards particulars. The structuralist's argument seems to beg the question by assuming straight off that properties and relations are the only things perceived.

Floridi bridges this gap by postulating the information object; however on the face of it the information object idea only seems to get us so far. The fact that something happens to be the case, or that a particular state of affairs is true as a result of our perception, lends itself to a propositional nature. We are able to state such things, such as it is raining today, or the box is red. Information objects are essentially bundles of properties and relations. Whilst no information object can be strictly identical with another, the properties and relations can and are identical across multiple instantiations of like objects (many objects can have a mass of 25kg and there are many red objects for example) and it is the properties and relations directly to which we talk. The information object seems to be an abstract notion that somehow demarcates a particular set (bundle) of properties and relations and appeals to our intuitions regarding extant objects in the world. Furthermore whilst they do not exist outside their instantiations it would seem properties and relations hold a more fundamental ontological position than that of the information object. Armstrong's state of affairs notion avoids the ontological trap Floridi's information object appears to fall into since it avoids the postulation of bearer.

It could be that Floridi's information object has no ontological implications at all but this doesn't seem to square with what he wants the notion to be. His claim is that all there is in the world is dynamically interacting informational objects, he means this literally. He is very clear that his is a structuralists view – information objects supervene on their constituent properties and relations, and if this is the case, an argument

could be made that information objects are not ontologically significant or additional to the things (properties and relations) they supervene upon. Armstrong on the other hand makes his position quite clear, "there certainly are no such entities as facts, that is states of affairs" (Armstrong, 1998, p114). The need for states of affairs in his ontology, Armstrong argues, is due to the contingency of a particular instantiating a universal. The idea of the state of affairs is simply propositional to Armstrong as we have already mentioned and so states of affairs are what hold their constituents together in a non-mereological form of composition (*ibid*, p118). To put it another way, Armstrong's ontology admits both universals (properties and relations) and particulars, and what connects the two is *instantiation*. So for example *a*, a particular, instantiates F a universal, *a* is F, for example the *box* is RED. A state of affairs is nothing additional to a particular and its properties; it is merely the linguistic term of convenience. States of affairs according to Armstrong supervene upon particulars. This being the case, states of affairs are not ontologically additional to the things they supervene upon (1998, p117). Given Floridi's structuralist position he may be going too far in postulating the information object as a distinct entity with any ontological significance.

4.4 Identity

Armstrong admits the *state of affairs* concept has no ontological import, rather it is the *particular* instantiated via a universal which exists as the "persistent object". The particular and its associated universals are constituents of states of affairs. This points us towards the idea that the counterpart to Floridi's informational object is not the state of affairs but rather the *particular*. "Simply put, micro- and macro-objects are analysable as informational objects that satisfy OSR (ontological structural realism) specifications" (Floridi, 2008, p29). Unlike Armstrong's metaphysics, informational realism makes no formal claims regarding the identity of informational objects except via a somewhat austere example. As discussed

in chapter 3, Floridi asks us to consider the pawn in a game of chess. Its *identity* is not determined by its contingent properties as a physical body, including its shape and colour, rather the pawn is a well defined cluster of specific *states* (properties like white or black, and its strategic position on the board) and is determined by a set of *behavioural rules*, "it can move forward only one square at a time, but with the option of two squares on the first move; it can capture other pieces only by a diagonal forward move..." *etc* (Floridi, 2008, p30). A question looms, could this be a suitable method whereby an informational object's identity could be established? Recall Floridi does state that informational objects do not possess any identity distinct from the relations within which they stand. The rest, regarding identity is left open. It seems we need a theory of universals similar to Armstrong, alternatively Floridi has some kind of attribute/relational theory of identity – but nothing further is said.

Armstrong's ontology on the other hand does have a particularly robust method of identity determination. If one is a realist about universals any basis for arguing two or more particulars are the same in some respect is generally resolved by arguing they instantiate the same universals (Armstrong, 1998, p22). It is not clear this solution is available to Floridi since identity for Floridi seems to be determined by the states and behavioural rules of individual informational objects – there could be more to informational objects and we'll investigate options soon. It is unclear however how effective (or complex) the determination of states and behavioural rules might be in determining sameness. While Floridi's identity criterium seems to work for functional objects such as chess pawns it is unclear how this functional approach might work for physical objects that lack specific functionality. Consider for example buckets of water, mountains or handfuls of gravel. It's not even clear that informational realism employs this functional approach across the board. Identity is only discussed with regard to the identity criteria of the chess pawn in Floridi (2008, p30).

Given Floridi's use of OO terminology and theory in developing his informational realism we might look to how identity is established in OO objects as a hint on how to deal with the issue. OO programming allows for instantiated objects to be strictly identical except in the sense that they are not one and the same object (indiscernability of identity) in a physical sense. That is two instantiated objects are considered exactly the same in OO if they both have exactly the same properties, even if they are not the same physical instance. In OO theory this is sometimes called *structural equivalence*. Furthermore the concept of identity in OO programming is intrinsically tied to the concept of reference. That is, two identical objects are identical if they exhibit all the same properties exceptingly only in that they reference a different location on disk. Different objects are distinct in that they may relate to other distinct objects or that the attribute set (the variables) are different upon comparison. Unique identity regardless of the equivalence of attributes or the object's relationship with other objects is guaranteed by the object ID (OID) concept. Each individual object has a unique OID in OO. This approach is clearly far too simplistic to be of much use to a general metaphysical theory. This is further evidence that the use of object oriented terminology and theory falls well short of a method suitable for developing a metaphysical theory (see Chp 3).

On the other hand if we accept a theory of universals, a clear criterion of identity can be established. We should point out this sense of identity is what Armstrong calls identity in a "loose and popular" sense (2010, p55). For example consider a 100 year old tree. The tree has remained in the same place for those hundred years. We typically refer to it as being the same tree that stood there a year ago, or indeed 75 years ago, and we'd also agree that it is the same tree that was a sapling 99 years ago. However in the strict sense the make up of the tree, for example its biochemical make up is significantly different as each year passes by. Under any strict construal of identity we are forced to concede the following. For any objects *x* and *y,* if $x = y$, that is if *y* has exactly the same properties as *x* then we are talking about the same object (Priest, 2000, p65). However this would seem to

exclude any interpretation of identity relating to most objects that persist over time because put simply, time changes things.

Virtually all macroscopic particulars change over time yet generally (in the "loose and popular" sense) we refer to them as being one and the same. Thus, while a strict theory of identity and this would include the one applied to OO programming, would be difficult to adapt to informational realism I don't see Armstrong's "loose and popular" theory of identity as being incompatible at all with Floridi's metaphysics. Furthermore Armstrong's method is compatible with an empirical approach and thus sympathetic to my own naturalistic bias. Orderings of properties into class systems is typically how science is able to categorise phenomena as well as the entities it investigates. The job of deciding what universals exist is not an *a priori* task, it is the job of empirical science.

Let us say then that informational objects exhibit continuity over time and provided any changes to those objects are not too abrupt there is a general resemblance from moment to moment. Chapter 6 considers issues of causality with regard to informational objects however it is worth stating at this point that it also seems that at any particular snapshot in time an informational object will be *causally* linked to itself at the previous moment in time. Armstrong presents this case with regard to his "particulars". If we imagine a causal break at a certain time in the existence of a supposed particular, at such a point there is absolutely no causal (or nomic) connection between what goes before and what comes after (2010, p57). No one surely would be prepared to say such a thing was one and the same thing before and after. Thus all the while a particular (an informational object) exists it must grow out of its causal past (*ibid*).

4.5 Imaginary, Possible and Probable Informational Objects

Most of the discussion regarding informational realism pertains to *actual* informational objects. As such the set of such entities discussed seems to be limited to objects that exist in a spatio-temporal sense, things such as car

windscreens, pawns, every day kitchen objects and so on. Such entities are often described as concrete rather than abstract. While there seems to be some sympathy for abstract entities, things such as real numbers, classes and universals, possible or probable entities don't get any mention in Floridi's informational realism. Informational realism deals primarily with structural objects, and "structural objects need to be concrete things in the world" (Floridi, 2008, p41). Nevertheless while possibility, along with other modalities such as impossibility and contingency appear to have a place in semantic analysis and of course modal logic, on the face of it these modalities look to be problematic with regard to informational realism. If we contemplate the possibility of such entities it's not to hard to imagine a very crowded universe consisting of an infinite number of possible informational entities, from possible elephants in the room to Quine's imaginary fat man in the doorway.

It could be that such modalities apply only to propositions and not to informational objects at all. However consider this from David Lewis on possible worlds, "When I profess realism about possible worlds, I mean to be taken literally. Possible worlds are what they are, and not some other thing" (Lewis, 1973, p85). Floridi intends us to take informational objects as literally as Lewis wants us to take possible worlds.

There are some problems with regard to other modalities and abstract notions that informational realism may have difficulty in dealing with but that I believe it must ultimately address. While I flag these issues here I don't intend to resolve them except perhaps suggest some avenues worth exploring. It could turn out that such issues are just not particularly relevant to informational realism. However it seems that other metaphysical theories attempt to address them, thus informational realism ought not to shy away from them.

We equate actual entities to existence or being. There are no non-existent entities. However, it is true in a sense that possibilities exist. It is true that it is possible that it might rain this afternoon. Of course the existence of the rain is not actualised or instantiated until it rains but it

makes perfect sense to talk of the probability of rain. The question is how such abstract notions fit into informational realism. It could merely be a semantic problem. The probable rain this afternoon might be an interesting weather related discussion point but as a metaphysical category it seems far more problematic. We claimed earlier following Armstrong and Oddie, that the things in the world that we attribute existence to ought to exhibit some causal properties. Does the probable rain this afternoon have any causal properties? The answer is not so clear; it may cause me to pack a rain coat. In this sense however it's not the fact that at 3pm it rained, it is my belief that it may rain that has the causal effect.

Perhaps it will turn out that informational realism has to merely satisfy itself as being an ontology of *actual things*. On the other hand informational realism freely talks of the properties and relations (or universals) of structural informational objects and such things are clearly abstract in nature, the red ball that weighs 3 kilograms instantiates red and 3 kilograms yet there is no such thing as red or 3 kilograms independant to the particular it is instantiated within. Thus our ontology is seemingly committed to certain abstract concepts but not others. To add further fuel to the fire, science, let alone mathematics, as Quine puts it, "is up to its neck in commitments to an ontology of abstract entities" (1953, p13).

It seems that such entities, red, 3 kilograms, the imaginary fat man, and the probable rain this afternoon are psychological entities, what they might be beyond their mental nature is difficult to explain. However there is no doubt they (some of them at least) convey information. To omit them from an informational ontology seems arbitrary. For example there are plenty of cases of impossible-to-verify scientitic theories that we generally accept, take the big bang theory for example. Many would argue that the big bang is probably based on the balance of evidence and competing alternatives. But are probable entities as ontologically significant as probable explanations? Surely, at least, probable entities supervene on the corresponding probable explanations that they are involved in. If we accept

that the big bang actually occurred then surely we are committed to the postulated entities involved in the event.

Further, as Armstrong points out "there is a great, ever-growing, edifice of mathematics with its apparatus of axioms and proofs that yield a certainty that, if not absolutely certain, is the most certain knowledge that we possess" (2010, p71). The question is can all this be merely conceptual or analytic in nature, and if informational realism can only admit concrete entities into its ontology then where are we to place everything else? I don't currently have an answer to these questions except to say that a robust metaphysical theory ultimately must deal with such issues. Informational Realism, as a metaphysical theory, is in its infancy; there is much work to do and dealing with such tricky topics is surely part of that work.

4.6 Conclusions

The fact that I have cast suspicion on the information object as having any real ontological significance doesn't in any way diminish the importance of informational realism. The mistake I suspect that many might jump to is that Armstrong's state of affairs concept finds its equivalence in the informational object concept in Floridian informational realism. This is a mistake. The closest equivalence is found between Armstrong's particular and Floridi's informational object. Floridi does not consider his information object as a semantic concept; he considers it to be a real entity. On the other hand Armstrong makes it clear that there are no such entities as "states of affairs". Issues discussed by Armstrong, rehearsed in this chapter, regarding the postulation of a bearer arise. These issues I think are difficult for informational realism to sidestep and Armstrong's theory appears to have the advantage in this particular area.

Issues regarding identity are robustly dealt with by the state of affairs account. There is scant detail with regard to informational realism as to how to deal with these issues. I do not see the adoption of a theory of universals as being problematic for informational realism. In fact such a

theory fits in well with discussions regarding properties and relations. Within OO programming issues of identity are dealt with in this way and it is evident that the value of a variable, be it a property (attribute) or a relation can be identical across multiple instantions of different objects.

In the final section of this paper I dealt with the issue of possible, probable and imaginary entities. The issues as to whether such entities exist in any meaningful way is controversial. Certainly there seems to be some causal influence initiated by probabilities but these seem to be more closely related to psychological states rather than any future instantiation of a state of affairs or object. Nevertheless many metaphysical theories deal with alternative modalities. As I mentioned earlier it could turn out that informational realism is restricted to an ontology of physical/concrete objects. This would be a shame and it would be difficult to argue that Floridi's theory offers any real advantage over Armstrong's. The result would entail having to exclude psychology as a science based upon informational realism, furthermore there are a great many concepts in science and mathematics that are clearly conceptual and abstract in nature – just how informational realism might deal with such concepts is unclear. While the issues surrounding identity are relatively easily resolved there is certainly fruitful ground in working out just how abstract entities as well as alternative modalities can be incorporated into informational realism and Armstrong's 'states of affairs' approach offers plenty of clues.

5 Explanation, Causality and Information

A shorter version of this chapter appears in the proceedings 2011, *The Computational Turn: Past, Presents, Futures?*, International Association of Computing and Philosophy Conference, Aarhus University, Denmark.

Everything that exists makes a difference to the causal powers of something.
David Armstrong, (1997, p. 41)

5.1 Introduction

It seems patently obvious that *information* plays a critical role in explanation. Statements and models that aim to explain certain events, phenomena or facts typically convey explanatory information. We could say, to have something explained, in a sense, is to be informed. Such claims seem so self-evident that one might wonder whether they even warrant being stated at all. However to those of us interested in the philosophy of information it is somewhat puzzling that the literature on scientific explanation does little to critically examine or formalise the role information plays. Some might say well, there just isn't anything worth saying. However, earlier forms of explanation have relied upon deductive logic (most prominently Hempel and Oppenheim, 1948) and recent construals of semantic information hold the view that tautology and classic logic systems are informationally trivial (Floridi, 2004., D'Agostino 2013). More recent attempts at characterising scientific explanation have included appeals to statistically relevant factors (Salmon, 1971., Railton, 1978, 1981) as well as causality (Salmon, 1998).

In this article I hope to go some way in offering an information-theoretic account of explanation. This approach is not dismissive of previous accounts, rather it attempts to explain explanation from an informational

perspective whilst drawing on various aspects of the historical and received views. Attempts have been made at explanation utilising informational concepts, most notably Greeno (1970). However Greeno draws upon the Shannon-Weaver model of information which is by and large syntactic in its approach. The method discussed in this paper will draw upon the contemporary semantic construal of information and information transmission. It should be noted that my general approach tends towards an epistemic and pragmatic methodology rather than a metaphysical one. Nevertheless I believe examination of the various empirical examples I present lends support to an emerging metaphysical theory - *informational realism*. While this is not the primary goal of the chapter, it does represent an underlying theme within this body of work and as such I hope establishes the relevance of this work within the larger context.

The first task will be to consider the role information plays within scientific explanation. Thus we begin with a more detailed examination of the received views mentioned above. Secondly I will utilise the metaphysical topic of *absences* as a method to show how an information-theoretic account of explanation is the right one. I will argue that information plays a critical role in the explanation of events when an absence of something serves as a cause. My starting points with regard to these related topics is as follows:

1. Peter Railton, with regard to explanations, talks of the related concepts "the ideal explanatory text" and "explanatory information" (1981, p. 240). Railton openly admits in his 1981 paper that whilst it is typical to speak of sentences or texts conveying *information* he knows of "no satisfactory account of this familiar and highly general notion" (1981, p. 240). Further he admits that neither do the notions of information defined by Wiener and Shannon appear to fit his explanatory theory. Given that Railton's work continues to influence attempts at theories of explanation, in particular Kitcher's (1989) unificationist account, an enquiry into Railton's "explanatory information" seems overdue.

2. Phyllis Illari has recently offered a production account of causality based on information transmission (2011). Her theory is interesting in that it attempts to deal with a significant criticism of production accounts, what is often called the problem of absences. These are cases which have puzzled philosophers with regard to causality where there seems to be a gap between cause and effect. Such cases are not limited to events that occur by chance, that is, events that occur outside the bounds of what we'd consider to be the normal laws of nature, scientists routinely talk of processes in terms of negative causation. Such cases are typically processes that occur within systems *in-the-absence* of some contributing factor. Illari suggests these cases can be explained informationally. We will explain what Illari means by a "production account of causality" and then consider the role of information with regard to negative causation or absences. If indeed it turns out that absences can be explained informationally then I believe this will provide further support for informational realism, that is, the idea that at a most fundamental level the world can be analysed informationally.

In order to appreciate the utility an informational approach might offer it will be helpful to briefly review the development of *scientific explanation* over the past few decades at least briefly. In this section I flag some issues which will be dealt with in more detail in the following sections. Whilst the section is not intended to provide a comprehensive overview of scientific explanation, I will cover what I consider to be the most prominent theories and I will try to focus upon accounts of explanation where the role of information was considered, even if only implicitly.

5.2 Early informational accounts

The body of philosophical literature on scientific explanation is substantial beginning[30] with the Deductive Nomological[31] (DN) model (Hempel & Oppenheim, 1948., Hempel, 1965). On this approach scientific explanations were considered deductive arguments with the inclusion of at least one law of nature. The DN model sought to provide explanatory information by means of two constituents:

1. An *Explanandum*, that is a sentence describing event or phenomenon that requires explanation

2. An *Explanans*, that is a set of statements which form the premises of the argument. The *explanans* are the statements which form the conditions and include one law-like generalisation necessary entailing the conclusion (explanandum).

For example upon observation we may conclude (explanandum) "The infant has Down's Syndrome". The DN explanation is formed by the (explanans) initial condition, "The infant has three copies of chromosome 21" and the law-like generalisation "Any infant that has three copies of chromosome 21 has Down's Syndrome".

Hempel's approach is to treat explanations as a method of conveying information about the phenomenon or event in question from a nomological point of view. Note he has little to say regarding causal factors. As such Hempel's DN method generally offers a rather austere sketch of an explanation. There is no discussion of necessary and or sufficient

[30] Although the roots of scientific explanation and understanding can of course be traced back well beyond Aristotle, recent philosophical history regarding scientific explanation is generally considered to begin with Hempel and Oppenheim's ground breaking paper *Studies in the Logic of Explanation*.

[31] The Deductive-Nomological (DN) model is a formal method of scientific explanation which is represented as a deductive argument with appeal to at least one established scientific law.

conditions or of causally contributing factors. We might deduce from this that the DN explanation is only a partial explanation. This is particularly the case if we consider relevant causal information important to our understanding of the event in question. Furthermore we mentioned that deductive reasoning and logical truths have been considered informationally trivial, lacking empirical content [32] (Floridi, 2004., D'Agostino, 2013). Those that agree with such claims might want to raise the question with regard to how deductive reasoning might provide us with any new information at all. I do not fully agree, in line with my already declared pragmatic stance, I tend to think deductive schemas can carry some informationally relevant content, even if it turns out that all that information relates to is something about language, reference or definition. I shall come back to this controversy later in this chapter.

To most people, Salmon (1998, p3) argues, "The suggestion that there is a close connection between causality and explanation would come as no surprise". And while distinctions can certainly be made between the two concepts there are many convergences. Salmon goes on to claim, "In many cases to explain something is to state its cause." (*ibid*). Furthermore causality is a concept that permeates everyday thinking, planning and understanding about our lives, within technology and language and in our efforts to intellectually understand our world and universe. Scientific discourse and understanding is almost always causal. We need to know if smoking marijuana increases the risk, or is causally linked to mental health issues, if second hand smoke causes cancer, if eating fresh fruit and vegetables improves one's health. We seek to find causal explanations for accidents such as aircraft and vehicle crashes. Causal concepts are ubiquitous in the sciences and across many vocational disciplines from engineering to medicine and information and computing technologies.

[32] The degree of informativeness of a logically deductive schema is perhaps controversial, however given scientific explanation has moved on considerably from the Hempelian D-N approach we can safely leave this controversy to one side.

The role information plays in explanation and causation is critical. To have something explained is, at least from an ordinary language point of view, to be informed. To be informed about a particular event or phenomenon is to be in a position to be able to offer some kind of explanation as to the cause of that event or phenomenon. The circularity of these claims just serves to show that whilst explanation, causation and information are not the same things, the concepts exhibit major areas of convergence. Information and causality in particular appear to be intrinsically linked and we shall examine this in more detail soon.

The structure of scientific explanation involves various relationships between entities, processes, laws and theories. The concept of information might benefit from an investigation into the connections or relations that exist between it, causal concepts and scientific explanation and it is this particular can of worms that we shall begin opening now.

The informational quotient in explanation was increased with Salmon's Statistical Relevance (SR) model (1971). Salmon's approach was developed in order to deal with explanations of low probability that is, events not adequately dealt with by Hempel's deductive explanatory model. Hempel had previously attempted to deal with statistical explanations with his inductive statistical approach (1962). However, Salmon argued that the approach was flawed due to its requirement for high probability and its demand for expectability. Salmon's SR model attempted to also capture causal factors that are explanatorily and statistically relevant or that exhibited conditional dependencies with the phenomenon to be explained. The general approach is that all statistically relevant information is considered explanatory, and statistically irrelevant information is not. What is crucial for explanation, Salmon argued is not "how probable the explanans renders the explanandum, but rather whether the facts (relevant information) cite in the explanans make a difference to the probability of the explanandum" (Salmon 1989, p59). For example, in testing medical therapies, whether they are physical (pharmaceutical) or psychological, controlled experiments are used. That is, we then compare the outcomes of

an experimental group (those which receive the treatment) with a control group (those which do not receive the treatment). The resulting evidence determines whether we can justifiably attribute explanatory import with regard to the treatment and its effectiveness in remitting the disease. What is important in the analysis of the results is not whether or not the treatment resulted in a high probability, for example almost all colds clear up in patients within a week whether or not they take vitamin C. That is, if the remission rate is high but no higher than spontaneous remission then it isn't valid to cite the supposed treatment as the explanation. Additionally if the remission rate for a particular treatment is low, but significantly higher than spontaneous remission then we might cite the treatment as being at least part of the explanation (*ibid*). A high probability Salmon argued is neither necessary nor sufficient for a legitimate statistical explanation, instead statistical relevance is the key factor.

Salmon typically talks about statistically relevant information in terms of a "reference class", "When we choose a reference class to refer to a given single case, we must ask whether there is any relevant way to subdivide that class." (1971, p42) and, "the reference class rule remains, then, a methodological rule for the application of probability knowledge to single events". (*ibid*, p44). A reference class is the minimum set of statistically, causally relevant, non trivial information that explains a particular event. Some points to note with regard to the SR model, firstly it represents a departure from the deductive argument based structure of the DN explanation. Rather a SR explanation is a collection of information statistically relevant to the explanandum. This approach was considered an improvement on Hempel's model since as Salmon argued "irrelevancies [are] harmless in arguments but fatal in explanations" (1989, p102). Secondly the SR model does not limit itself or consider explanations for outcomes that only have a high probability of occurring. The SR model offers a method whereby events which have even extremely low probabilities of occurring can have some explanation. Although not acknowledged by Salmon, most likely because he simply did not think

about it, the SR model seems to represent the first information-based approach to scientific explanation.

It became clear to Salmon that statistical explanations needed to account for causal factors. That is, it seems that either statistical explanation must elaborate upon the causal relationships that exist between the effect, or the explanadum, to be explained and its causes, or that statistically relevant information supervenes upon causal factors. At this point we can raise a significant open question, one that perhaps did not occur to those working in scientific explanation but one that seems salient to the philosophy of information:

Do informational objects supervene upon causes?[33]

We shall come back to this question later within this article.

Salmon's concerns over causal factors eventually led him to propose a full blown casual theory of explanation (1998). His principal claim is that a scientific explanation is constituted by state of affairs predominantly recognised as patterns in the world where a pattern consists of at least one causal process. For example, we make a causal inference that the antibiotic penicillin eliminates infection-causing bacteria by observing the elimination of infection in a statistically significant number of cases – that is, an observable pattern emerges as a result of penicillin being administered to patients with certain infections.

Causal processes are described by Salmon as being continuous (in a spatio-temporal way). This view contrasts with the popular view of causality being a "relation" between particular events (the cause, and the effect). Salmon first explained his causal theory in a paper entitled The At-At Theory of Causal Influence (1977, reprinted in Salmon, 1998). The At-At theory claimed not only to resolve Zeno's arrow paradoxes but also

[33] See Chapter 2, The Floridian Notion of the Information Object in this thesis, for a full discussion regarding the concept of the 'information object'.

proposed a foundation for a concept of propagation of causal influence. While information played a significant yet largely unexplained role in virtually all of the models of explanation it seemed particularly apparent to Salmon as he began to explore in his At-At causal theory.

Causal processes, argued Salmon, necessarily transmit information (1998, p.16). Salmon initially explained this as the ability of a causal process to transmit a mark. An example is introduced in the form of a light beacon centred inside a circular building, something like the Roman Colosseum. The beacon rotates rapidly transmitting the spot of light on the inside wall of the building as it rotates. We can introduce a mark within this process by modifying it. For example if we are to place a red filter in the path of the light at a single point within its rotation, while the source light is still white, at the point where the red filter is located a red spot will appear upon the wall momentarily. Salmon used this example to argue that causal processes are unique in the fact that they are able to transmit information. However, the theory was eventually shown to have problems and Salmon was forced to modify this approach based on the fact that the theory had to account for counterfactuals. That is, if E had not occurred then C would not have occurred, or in the case of the rotating beacon, the luminous spot would not have changed if the mark had not been introduced into the process. The introduction of counterfactuals into theories of causation may have the potential for a claimant to draw absurd conclusions. For example it is true that had the filter not been introduced the mark would not have been transmitted but it is also true the mark would not have been transmitted under a vast variety of other contingencies, had the person introducing the mark never been born, or had the person been struck down with a heart attack just prior to introducing the filter, or had the light in the beacon failed and so on. The issue as Salmon describes it, is with the mark criterion itself as an explanation of causal processes. While the method might be useful in an experimental sense, that is for tracing or identifying causal processes, Salmon argues this is “only a symptom of the fact that it is actually

transmitting something else. That other something I described as information, structure and causal influence" (Salmon, 1998, p253).

A full discussion of the problems associated with Salmon's mark transmission account of causality and the issue of counterfactuals is beyond the scope of this paper however we shall say that Salmon, following Dowe (1992) moved towards an account of causality in terms of "conserved quantities". Under the new version, a process is now considered causal if it transmits a conserved quantity; examples of the properties conserved given are usually energy-mass, momentum or charge. The advantage of this account is that it avoids the notoriously subjective approach involved with counterfactuals which involve appeals to contextual and pragmatic concerns regarding the truth value of the causal explanations given (Salmon, 1998, p19). The most significant problem with Dowe's conserved quantities approach is that it relies upon these relatively few properties, charge, mass and so on in order to explain causal processes. However, for many causal processes, particularly in the higher sciences, these measures are not relevant properties at all. As Illari points out for example, "Charge, mass or momentum seem incidental to causal claims as 'smoking causes cancer' since the various sciences of cancer to not concern themselves with charge, mass or momentum" (2011, p4).

Salmon's mark transmission account and the Salmon-Dowe conserved quantities theory are considered to employ production accounts of causality. This is in contrast to what are called dependence accounts. It is worth briefly explaining the difference since dependence and production play different roles in explanations. Furthermore the two differing accounts tend to throw up counter examples against each other meaning the distinction is generally the point where disagreements begin. Illari provides useful illustrative examples in her 2011 paper. To illustrate the problem for dependence accounts Illari asks us to image Billy and Suzy throwing stones at a bottle. At one point Billy's stone hits and shatters the bottle, however an instance later Suzy's stone whizzes through the space and flying shards of glass. Had Billy's stone not been present Suzy's stone would have

shattered the bottle. However we can clearly trace the link back to Billy's stone as being the cause of the bottle shattering. Contrastingly to illustrate the problem for production accounts we imagine Billy and Suzy flying a bombing raid over an enemy city. Suzy flies the bomber and drops the bombs, Billy however flies a fighter shooting down enemy fighters trying to destroy Suzy's bomber. While we can argue that the success of the raid depends on Billy's actions we intuitively sense that there is no continuous link between Billy's actions and the bombing as there is with Suzy's actions. The point to be made is that production and dependence accounts of causality both appeal to differing intuitions about causality. As such a unified account has proven difficult.

We have outlined the general picture regarding scientific explanation beginning with Hempels DN method, the problems associated with this approach and consequently Salmon's through statistical relevance and causal approaches. We have also seen how Salmon proposed that information transmission played an important role in causal explanation. However, there has been another genuine attempt to bring an information-theoretic approach to bear on scientific explanation. Greeno (1970) drawing upon the Shannon-Weaver model of information transmission, develops a theory that proposes a "rough metric" for evaluating explanations. That metric is "an elementary quantity in statistical information theory" (*ibid*, p91). Greeno's approach was designed to overcome problems associated with explanations related to unpredictable or uncertain events. Two sets of variables are established; one set {M} represents the variable or set of variables whose values are to be explained, that is, the explananda of the theory. The other {S}, the explanans, is the set of variables used to explain occurrences of {M}. The idea being that the amount of information transmitted is the amount by which the uncertainty is reduced (*ibid*). Once again however, this theory, given its quantitative approach is rather austere from a semantically informational perspective. Salmon (1998, p98) provides a good example of Greeno's theory in practice in order to point out its shortcomings. We begin by imagining Smith and Jones have both

committed suicide. We can use our best psychological theories to gather all the relevant information about Smith and Jones such as sex, age, ethnicity, state of health, education, marital status and so on. In analysing such demographic data we might conclude there is a low probability that Smith would commit suicide but a high probability that Jones would do so. The general argument goes that risk factors associated with incidences of suicide are elevated for example in male indigenous individuals with poor or limited education, unemployed and so forth. Salmon points out however that while these may be causal factors that have some statistical relevance, they do not on their own mean that an explanation of Jones's suicide is any better than Smith's. There may be a very low probability of a particular event occurring but this does not mean that we cannot explain Smith's suicide unless we can cite the conditions under which it would be highly probable. The explanation of a low probability event is not necessarily any weaker than that of a high probability event.

Woodward (1987) also criticises Greeno's approach arguing he conflates "reason-seeking" with "explanation-seeking" questions in his attempt to measure explanatory power by transmitted information, although he offers no clarification of this distinction. Nevertheless I am happy to go along with Woodward's claim that not all reductions in uncertainty constitute explanations. I think this is illustrated in Salmon's criticism of Greeno's approach. That is, we are able to provide perfectly adequate and valid explanations for events for which the probability of occurrence is extremely low. For example the likelihood of an asteroid hitting earth causing widespread and mass extinction is very low, yet given all the relevant factors combined with laws relating to orbiting bodies and so forth such events are highly explicable. What is important, Woodward argues, is not a somewhat arbitrary set of factors relating to the frequency of occurrence of the explanandum phenomenon but rather information that is actually causally or nomologically relevant (*ibid*, p39). This will require a more semantic application of information than the Shannon-Weaver application that seems to be favoured in the early literature.

While there may be some value in developing these kinds of information theoretic approaches that might be used to predict (or reduce uncertainty) within certain contexts, and we see such approaches used for example by financial institutions in assessing the certainty (or uncertainty) of an individual's ability to service a loan, such approaches are not particularly explanatory. Conclusions drawn about a bank's decision not to offer a customer a loan are often informationally austere, "you did not meet our criteria", or "based upon the information you have provided we considered your application too risky". Similarly drawing conclusions about the cause of Smith or Jones's suicide based upon the factors listed above is tenuous at best. And if, as Salmon suggests, to have something explained is to understand its cause then an approach that merely reduces uncertainty seems flawed. Thus while such "information transmission" approaches may have practical value to specific events, Salmon (1998, p108) argues that such approaches "seldom, if ever, have genuine scientific import". The explanations that are scientifically interesting according to Salmon are those which relate to explanations of classes of events.

The above section should, I hope, provide enough background information on scientific explanation including attempts to incorporate informational concepts into the picture. What follows now is a more detailed consideration of Railton's concerns regarding information and explanation. This is important since it appears that Railton was the first to indicate a robust reconciliation between informational concepts and explanation was required. After taking into account Railton's concerns we shall consider and explain the problem of absences and their relationship with information, causation and explanation.

5.3 Railton's Informational Approach

The topic that seemed to fascinate many of those working on scientific explanation was probabilistic explanation. That is, how do we go about explaining, scientifically, things that seem to happen purely by chance? The

structure of explanations of such events seems fundamentally different to the explanation of a scientific theory that relies upon appeals to sets of laws or axioms such as DN explanations. While such kinds of explanation often take the form of deductive arguments, the axioms or laws forming the premises of the argument, many totally acceptable explanations of chancy events do not appeal to laws.

Railton (1981) provides the following example: imagine we spray a healthy milkweed plant with a dose of herbicide. The dose of herbicide alters the biochemical state of the plant from a normal, healthy state S, in which plants have a probability .9999 of surviving 24 hours, to a state S', in which the probability of survival drops to .05. However when we observe the plant 24 hours after it has been sprayed, we are surprised to find it still standing. How are we to explain its survival? Of course we don't ever attribute the spraying of herbicide to the survival of the plant, although we acknowledge the plant had a .05 probability of survival after spraying. While the spraying seems to be part of the explanation it in fact lowers the probability of survival. Thus, we typically we reserve 'probabilistic cause' for factors that raise the probability of a particular event happening. Had the plant failed to survive, we would have appealed to the spraying as the cause of its death.

The point to note about the above case is that it appeals to no theoretical laws yet it seems quite acceptable as a statistically based explanation. Nevertheless a more elaborate explanation, one that explains exactly why the plant survived after spraying rather than perished, might be required. While we have some relevant quantitative explanatory information, there are certainly some facts or qualitative information, perhaps as yet unknown, that might help explain the cause of the plant's survival. Put simply we want to know the why and how of the plants survival. Answering such questions will provide causal and mechanistic information directly relevant to the plant's survival.

Railton's somewhat pragmatic approach asks us that we consider a continuum of explanatoriness. At one end we find what Railton terms the ideal explanatory text. At the other end we have statements devoid of any explanatory information whatsoever (*ibid*). But what might constitute an ideal explanatory text? Railton suggests such a text ought to contain all the statements that enable us to give accurate answers to questions regarding the phenomenon p such that the statements provide explanatory information concerning why p. (1981, p240). However while we regularly talk about sentences conveying information in support of complete explanatory texts, Railton concedes he knows of no satisfactory account of the familiar and general notion that is information. It is clear, Railton concedes, that the notion is not the same one employed in Wiener-Shannon information theory, nor is it for that matter the method which Greeno's theory employs. What Railton seeks instead is a method by which the mechanisms and causes at work are exposed, those which elucidate the why and how of a particular event or phenomenon. The kind of information discussed in the Wiener-Shannon approach is syntactic and generally refers to some kind of metric with regard to the measurement of signals from an observed source.

The kind of information required to constitute an "ideal text" Railton argues is semantic information, information about something (*ibid*, p244). Thus Railton adopts the somewhat "received view" of semantic information. This account is often attributed to Bar-Hillel and Carnap (1953). We can summarise this position regarding the degree of informativeness of explanatory statements as follows. If a consistent statement S implies another statement S* then S contains at least as much semantic information as S*. If S is a tautology then S contains no semantic information since S eliminates no (logical) possibilities.[34] If S contains semantic information i(S), and S* contains semantic information i(S*) and if the conjunction of

[34] Railton claims to follow logic "if we take logic for granted" is his parenthetical claim. I take it Railton means the logical approach to semantic information as developed by Bar-Hillel-Carnap (1953).

i(S) and i(S*) is self-consistent, then the amount of semantic information contained in the conjunction will equal the sum of the two minus any overlapping information (*ibid*, p245). He goes on to state, within these properties of semantic information the quantity bears an inverse relation to probability. Propositions with a logical probability of one, tautologies, convey zero information. A conjunction raises semantic information but in independent events it lowers probability. For example, the statement "tomorrow it will be warm and cloudy" contains more semantic information than simply "tomorrow will be warm". Similarly illustrating the inverse relationship with probability Railton would argue, "tomorrow will be warm and the stock market will fall" has a lower probability of obtaining than the statements expressed as individual events. Disjunction on the other hand lowers semantic information but for independent events it raises probability, for example "I either saw the suspect at the 7/11 or at the service station". The individual constituents asserted independently may seem uncertain, however the probability of the disjunction in its entirety being true may be more likely.

It should be noted that any logical analysis of semantic information carries with it significant problems. The first point to note is on this account of semantic information logical inference is considered tautological and does not convey information or at least is considered informationally neutral. Thus we would immediately consider Hempelian DN explanations not to be particularly informative at all. Railton places such explanations towards the pessimistic end of the continuum of explanatoriness, in this case, the end that is largely devoid of explanatory information.

There are deeper issues with the Carnap-Bar-Hillel approach involving troublesome paradoxes. Railton tries to cut some of these issues off by including an informational consistency constraint into his definition of explanatory statements. That is, statements that constitute explanatory information must at least be self-consistent.

This is an attempt to cut off a significant paradox which applies to contradictory conjunctions of sentences. We can illustrate quite simply; for

example, consider the conjunctive statement involving a single event "tomorrow it will be warm and tomorrow it will be cool". Clearly the statement is a contradiction but under Carnap-Bar-Hillel notion of semantic information it is considered maximally informative since it reduces all uncertainty. Few of us would argue that such a sentence conveys any explanatory information, indeed quite the contrary; it conveys no real information about the weather tomorrow. Bar-Hillel and Carnap have this to say about the contradiction:

> It should, however be emphasized that semantic information is here not meant as implying truth. A false sentence which happens to say much is thereby highly informative in our sense. Whether the information it carries is true or false, scientifically valuable or not, and so forth, does not concern us. A self-contradictory sentence asserts too much; it is too informative to be true. (1953, p229).

I disagree, while some might want to argue that such a sentence eliminates all uncertainty, the form of the sentence could hardly be considered to convey anything useful or factual. Indeed the General Definition of Information (Floridi, 1999) would not hold such sentences to be informative.

Whether or not the Carnap-Bar-Hillel account of semantic information has any utility is debatable, however it does seem that such an approach is wholly unsuitable as a method that might form the basis of an information-theoretic account of scientific explanation. Put simply we expect the statements that comprise scientific explanations to be both consistent and true. There is simply no room in scientific theory or explanation for self-contradictory statements. Railton nevertheless seems to struggle with this asking "What sort of semantics should we use? Must semantic interpretation be a function of context? What constraints must be met in order for semantic

interpretation to exist?" (1981, p245). It seems to me as I go on to discuss presently, the answer to these problems is relatively straight forward.

While we expect all individual statements constituting explanatory information within an explanation to encapsulate truth, it would seem that the requirements of a semantic account of information for scientific explanation must also include some kind of consistency constraint. There is precedent for such a demand. Popper (1959) argues that the requirement of consistency "plays a special role among the various requirements which a theoretical system, or an axiomatic system, must satisfy" (*ibid*, p72). "From a self-contradictory statement, any statement whatsoever can be validly deduced" (*ibid*, p71). Thus we should rightly conclude that the informational content of inconsistent self-contradictory statements is null. However we should be wary of jumping from this to the conclusion that collectively all statements within an explanation must come with a strict requirement of complete consistency across the ideal text. While individual inconsistent or self-contradictory statements are not informative in any sense many scientific theories exhibit some internal inconsistencies. It could be that not all the information is in yet. While a particular theory may have considerable explanatory utility there could well be issues or problems that still require 'figuring out'. How are we to formulate just what constitutes an allowable level of internal inconsistency? I hope what follows goes someway in offering a solution.

The kind of semantic theory of information (that addresses the issues of consistency and truth) that Railton sought now exists. Floridi's (2004) theory of strongly semantic information largely addresses Railton's primary concerns. Floridi suggests such a theory of semantics must (a) not be as brittle as a classic logic system, that is, there must be some minimal ability for the system to tolerate some inconsistency, and (b) not have a progressive degree of fault tolerance. With regard to (b) Floridi suggests, for example, in the context of evaluating moral responsibility the errors are usually evaluated increasingly severely as each error occurs. For example, an individual making a single poor moral judgement might be forgiven, but

each consecutive moral mistake gathers an increased negative impact or is judged more harshly.

Classic logic systems, the kind I assume Floridi talks of above have recently been described as being informationally trivial (D'Agostino, 2013). There are a few reasons for this the most influential emerging from what Hintikka terms the scandal of deduction. Upon learning that deductive reasoning is tautological or analytical and that logical truths have no empirical content and thus cannot be used to make factual assertions, one might ask, in what sense can deductive reasoning provide us with any new information? Hintikka concludes, "it is not perfectly obvious there is some such sense, for what point would there otherwise be to logic and mathematics" (1973, p222).

If D'Agostino is right then it would seem that any informational semantics in the service of scientific explanation may only work as a useful kind of heuristic. Nevertheless a full explanatory history, something that might approach Railton's ideal explanatory text must give an account of all the causal, probabilistic and any other facts associated with the phenomenon to be explained. Such a text is likely to contain a mix of truthful and consistent quantative and qualitative statements as well as deductive nomological (DN) statements (tautologies) in some cases. In this sense I disagree with Railton (1981), Floridi (2004) and others. That is, I think we need to consider DN statements or tautologies to be to-some-minimal-extent informative although conceding perhaps not empirically so. While law-like, and in some cases tautological statements, are not particular informative, in conjunction with or as part of a complete explanatory text they seem to be minimally informative from a psychological perspective. That is, they are at least informative in the sense that they convey some kind of truth value with regard to explanatory statements (inferences) and in this sense I suggest are minimally informative but only in the context of their contribution to the greater explanatory text.

We should make a brief note with regard to our constraint of truthfulness of explanatory information. Floridi's theory of strongly semantic

information resolves the Carnap-Bar-Hillel paradox by requiring information to encapsulate truth. This requirement avoids the commonly cited Carnap-Bar-Hillel problem that their semantics has little to do with truth. Although Floridi has argued that false information is (approaching) oxymoronic (2004) and while this claim has proven controversial amongst some scholars (see for example Fetzer, 2004., Dodig-Crnkovic, 2005), such claims from a philosophy of information perspective may well be interesting. However, they bear little relevance to our investigations here and as such we will not consider them. The requirements for an information-theoretic account of explanation are that all relevant statements that constitute explanatory information are both true and self-consistent. Certainly with regard to scientific explanation any notion that we could draw informative explanatory conclusions from false information is incoherent.

A final point worth noting with regard to explanation, Railton points out there exists 'why' questions that simply do not have answers (1981, p248). Such questions relate to why one probabilistic outcome obtained as opposed to some other potential outcome and no part of an ideal explanatory text Railton argues provides sufficient reasons for why one particular probability obtained over another. Such questions are "requests for information that is simply not available." (*ibid*) That the information is not available does not mean that we don't know rather that there is simply nothing more to be known. For example Railton suggests, "Current physics offers us no reason why there is no negative gravity, despite the fact that there is both negative and positive charge." (*ibid*). It is at this point that Railton offers us a segue into the second part of this paper. "The point is that a theory may legitimately spurn certain requests for further explanation in both probabilistic and non-probabilistic cases, and we can now say what such spurning consists in: the absence of certain things from the purported ideal explanatory texts based upon that theory" (Railton, 1981, p248). While Railton's use of the term 'absence' here is not quite the same as the metaphysical problem of absences we are about to discuss, it does highlight

the point that an absence of certain things plays a role in scientific explanatory practice.

5.4 Absences

While the problem of absences has been discussed in metaphysics (D.M. Armstrong, 2010, Chapter 11., Dowe, 2000, Chapter 6) little has been said with regard to scientific explanation, absences and their associated degree of informativeness. Armstrong (2010, p83) following Dowe (2000) sidesteps the issue by arguing absences do not seem to have any causal power, instead a counterfactual account is given. Imagine a father grasping a child out of the path of an oncoming car so the child is not hit by the car. The nonexistent thing that was prevented, the child being hit by the car, Armstrong argues, has no causal efficacy. (2010, p83). If we reverse the case, that is the child is indeed hit by the car, the counterfactual is reversed, in this case the non-existent thing, the grabbing of the child has no causal efficacy since it never happened. On this account, certainly according to Dowe (2000) causation by absences is not really causation but a kind of quasi-causation. Schaffer (2003) on the other hand develops a case for what he calls "negative causation". "Negative causation occurs when an absence serves as a cause, effect of causal intermediary" (Schaffer, 2003, p197). The significant point of difference between causation and negative causation is that the latter involves no physical connection between cause and effect. Shaffer's thesis is essentially that causes need not be physically connected to their effects. (*ibid*).

There are many trivial examples we could draw on to illustrate negative causation. For example the failure of several debtors to pay their bills on time may be the cause of a manufacturer to fail to make a profit or indeed to remain viable as a business. Thus we might imagine the failure to initiate or interrupt a particular process as having a causal effect down the line. Hart and Honore (in Schaffer, 2003, p201) argue "The failure to provide others with opportunities for doing certain things or actively depriving them of

such opportunities are thought of in causal terms." (Hart & Honore, 1985, p2-3). Such examples may raise the suspicions of philosophers particularly those interested in causation and explanation, however we point out such folk theory cases illustrate the fact that instances of negative causation are ubiquitous within ordinary language and are supported by our everyday intuitions. Furthermore we can provide many examples of negative causation from the sciences and these examples seem to hinge upon what can be analysed at a fundamental level as an information exchange. That we can analyse such examples in informational terms, I believe, provides strong support for an information-theoretic account of scientific explanation as well as, from a wider perspective, the metaphysical theory of Informational Realism forwarded by Floridi (2004, 2008).

Prior to introducing scientific examples of negative causation I want to discuss an important relationship between information transmission, causality and reference. This will serve as a more formal analysis of what I called our folk theory of negative causation; that is, the popular intuition that suggests that negative causal influences or absences play an important role and thus feature prominently in causal discourse. However it also establishes a very important point regarding the causal nature of the informational link between an object or state of affairs and an individual's conception of that object or state of affairs and that this process appears to be neutral with respect to positive or negative causation. To explain further we briefly revisit the notion of reference discussed in Chapter 4 of this work.

In the chapter 4 I considered Gareth Evans's (1982) work on reference with respect to the importance information transmission plays in his theory of reference. Evan's work is notable in that it moves away from previously held descriptive notions of reference towards a more information-theoretic concept which involves causal (informational) links between word and

object. Earlier Kripke (1972) argued that reference is a causal notion[35]. We see similar reasoning in Evans in that he explicitly describes the process of demonstrative identification as being driven by an information-link between subject and object. For Kripke the reference of names (of objects) is transmitted by causal chains and thus the reference of any particular name in question isn't so much a description of the object but rather the outcome of the causal chain which produced the name. While it is clear my biases are informational in nature it does seems clear that the causal chains that Kripke talks about are by and large the same kind of thing that Evans describes as an information-link. Evans in a section of text on demonstrative identification begins "I think that an information-link between the subject and an object is a crucial necessary condition of the mode of identification we are trying to characterise" (Evans, 1982, p145). He continues, and this I take to be the crux of his argument:

> demonstrative thoughts take place in the context of a continuing informational link between subject and object: the subject has an evolving conception of the object and is so situated *vis-a-vis* the object that the conception which controls his thinking is disposed to evolve according to changes in the information he receives from the object (1982, p146).

Evans clearly views information as playing a causal role in the identification of objects in the world, but this is not a traditionally held view of causation, that is it is clearly not a Humean conception of causation as a necessary physical connection. Instead information transmission is the critical component in the process and it is also clear that this process is

[35] In 1973 Paul Benacerraf published and influential paper entitled *Mathematical Truth* in which he argued, following Kripke, for a causal account of knowledge. While I do not explore Benacerraf's paper within this work it is by and large compatible with the arguments presented here.

unconcerned with whether the naming was produced under conditions of positive or negative causation. Consider the final example given by Kripke in his 1972 text *Naming and Necessity*. In this example Kripke refers to an example by Evans (reference not given) regarding "reference shifts". I hope the example will make clear what Kripke and Evans mean by reference shifts. Madagascar was originally a native name for a part of Africa. Marco Polo however erroneously thinking he was following the native usage applied the name to an island to the extent that today the usage of the name Madagascar overrides any historical connection with its original reference (1972, p163). Kripke continues by saying that David Lewis argued the same thing could have happened had Madagascar designated a mythical locality. Thus reference can shift from a real reference to an alternative real reference or from a fictional one to a real one and *vice versa (ibid).* The point is that both positive and negative causes are pertinent to effects and in this case at least are inextricably related to information transmission.

The interesting thing to note with regard to *absences* is that they do not exchange *conserved quantities* as is the case with Dowe's account of causality. That is, with regard to absences no physical or material link (such as mass, charge or momentum) can be found. Absences in this sense are not the same kind of causal processes described by Dowe however they still appear to be involved in important mechanisms operating across many systems including meaning and reference within ordinary language as we have illustrated in the above example. And from an informational perspective there is nothing to suggest that our ideas ought to be confined to physical or real world objects. From the Floridian perspective, everything qualifies as an informational object and this of course includes objects and their referents.[36] Thus negative causation with respect to *reference* is just as legitimate as its empirical counterparts within the scientific examples that I will illustrate presently. Thus, I suggest a more inclusive account of

[36] See Chapter 3 for a full discussion of the Floridian Information Object.

causality, one that accounts for absences in terms of information transmission, can be given. Such an account will be required for an information-theoretic account of scientific explanation since absences play such an important role in many scientific explanations. Let us now consider examples from the sciences.

Illari (2011) draws upon an example from biochemistry "Cells routinely alter which enzymes they produce in response to which metabolites are available. A cell stops producing lactase, for example, in response to the absence of lactose diffusing into the cells cytoplasm." (2011, p15). Illari provides two biochemical examples including this by Craver, "the absence of the Mg2+ block does seem to cause Ca2+ to enter the cell. At least this is what controlled experiments suggest: when the Mg2+ block is in place, the Ca2+ does not enter the cell. When the Mg2+ block is removed, the Ca2+ current begins to flow" (Craver, 2007, p80). The absence of Mg2+ in this case is informative. It indicates with a high degree of probability that Ca2+ will begin entering the cell. Floridi has argued that at a fundamental level a datum is essentially a *difference* (2004, p44) and Wiener's claim is particularly important in this case, "information is information, not matter or energy, no materialism which does not admit this can survive at the present day" (1961, p132). There is no *conserved quantity* exchanged as per Dowe's account of causality yet the absence of Mg2+ is clearly causative in this process. We can say with confidence that the *difference* Floridi talks about is a discrete state, a datum. That discrete state is nothing less or more than simply the absence of Mg2+. An information-theoretic account is the most coherent in this case.[37]

Such examples introduce serious challenges to standard accounts of causality in particular Dowe and Armstrong's metaphysical accounts, but also Salmon's applied account of causal explanation which relies upon Dowe's conserved quantities approach. These approaches attempt to

37 For exhaustive examples see Schaffer (2003).

explain away absences with appeals to counterfactuals or in the case of Armstrong, *differences*. Of course philosophers familiar with the history behind causation will be aware that the strong focus upon physical or material connections between cause and effect began with David Hume who defined causation as a *necessary connection*. This clearly prompted Dowe to maintain that some kind of physically definable energy flow, such as a conserved quantity, energy, mass, momentum or such like is the thing that connects causes with their effects. Negative causation overturns this assumption. I go one step further and suggest that negative causation can be analysed in terms of an information-theoretic analysis of absence.

Illari argues a significant metaphysical problem for absences is that by introducing a gap between *C* and *E* seems to rule out the possibility of a link between *C* and *E* and that the solution to this problem is in proposing that information modifies what we think of as the gap (2011, p16). Illari goes on to argue that given we can show the causal role information plays with regard to absences this shows that our conception of causal connection is wrong. We need to update this conception to account for the kinds of causal connections that the sciences are telling us exist.[38] While Illari's 2011 paper is about providing an information-theoretic account of a production account of causality her work can be naturally extrapolated to strengthen an information-theoretic account of scientific explanation. Furthermore this is entirely consistent with Floridi's informational realism, the idea that the world can be described by the dynamic interaction of informational objects.

Salmon was partly right. His mark transmission approach was in part the right idea however he was also somewhat fixated upon finding the Humean necessary connection and this partly contributed to his decision to drop the

38 This begs the question, which comes first science or metaphysics? I should point out, my approach in this particular paper is pragmatic in nature and it seems to me that any system of metaphysics that did not allow or account for actual phenomena observed in the sciences would have little real value. Thus our acceptance of any metaphysical system surely must be determined in part at least in that it can accommodate science in a broad sense. I thank Morgan Luck for pointing this out to me.

mark transmission approach in favour of Dowe's conserved quantities method. Dowe's approach appeals more directly to a necessary and very much physical connection existing between cause and effect. However by adopting an information-theoretic view of causality rather than relying upon an *a priori* method of establishing causal connection we can leave it to science to establish empirical grounds for causal connections, the issue of whether they are negatively or positively causal is of little consequence. Furthermore the information-theoretic approach also accommodates Dowe's conserved quantities approach. Dowe's approach is simply analysis at another level of abstraction, one which within certain contexts makes perfect sense such as in the case of the usual Humean examples, billiard balls bouncing off each other where the exchange of a conserved quantity can be easily explained. On the other hand where absences are involved, as we see in Illari's biochemical example discussed earlier, the causal connection in this case can only be analysed in an informational sense. I believe this is compelling evidence to suggest the information-theoretic approach is the right one.

5.5 Food for Thought: Informational Objects Supervene on Causes.

Earlier in this paper we posed the question: do informational objects supervene on causes. For this to be the case is simply to say that the properties that we describe as informational properties differ with respect to or as a relation of the properties that we describe as causal properties. That is there cannot be a causal difference without there being a corresponding informational difference. Let's consider how this might be cashed out.

Central to this essay was Salmon's metaphysical claim that causal processes necessarily transmit information. We can take Salmon's claim one step further and argue that the transmission of information can be reduced to a difference, that is a difference in the background or environmental conditions. This according to Floridi, constitutes necessary and sufficient conditions for the existence of information. "A 'difference' is

just a discrete state (that is, a datum), and 'making a difference' simply means that the datum is 'meaningful', at least potentially" (Floridi, 2004, p44). We have provided some empirical evidence to suggest that such differences have causal influence. We could provide some explanatory information about what such differences consist in. However I suspect in some cases we may be faced with Railton's problem, that is, sometimes there may be nothing further to say. Let us try to frame the problem from a different perspective.

Floridi's informational realism proposes a world consisting of dynamically interacting informational objects, on this account everything can be analysed in informational terms. While I don't propose to forward a complete informational-theoretic metaphysical theory of absences here I think we can propose that the "absence" itself ought to be analysable from an informational perspective. If the absence itself can be claimed to be a self-contained informational object, in that an absence can clearly be thought of in terms of a datum or a difference this seems to suggest that informational objects supervene on causal relations. I expect this to be a particularly troublesome metaphysical issue since the problem with analysing an absence in informational terms would seem to be one of demarcation. Where do we draw the relevant informational boundaries? Conceivably one could argue that everything not relevant to the observed phenomenon requiring explanation is part of the absence. One suggestion is that we characterise the absence informationally in terms of the causal narrative that is what causal efficacy does the absence have? Thus we might argue that informational objects, in this case admittedly the quite abstract notion of an informational object that comprises the absence, supervene upon causes. There cannot be an informational difference without there being a causal difference.

One argument against such a claim is the argument to ontological innocence: whether A properties are indeed anything over and above B properties. We can explain this with reference to a metaphor given by Kripke (1972). Kripke imagines God creating the world "what does he need

to do to make the identity of heat and molecular motion obtain? Here it would seem that all he needs to do is create the heat, that is, the molecular motion itself" (1972, p153). Thus once the B properties are fixed (the causes in our case) then the A properties (information) come along automatically. A properties, the information in our case are necessarily instantiated as a direct result of the B properties (causes) and so, one might argue, information is nothing over and above causes. However, I think a clear distinction can be made. Our talk of causes are more closely related to what we might categorise as mechanisms, that is some particular event occurs as the result of or due to the cause. We have shown this works for absences or negative causes also. These phenomena are analysable from an informational perspective. While it is clear that we have an intuition about the distinct nature of the two, causes are talked about as being distinct from information, informational states or informational objects. Certainly it seems that while there is a clear relation between the two and that relation *prima facie* seems to be a supervenience relation, it is also clear that we are talking about two separate sets of properties, causal properties and informational properties and that these are numerically distinct.

5.6 Conclusions.

In this paper we explored the relationship between Information, Causation and Scientific Explanation. Historically scientific explanation, as illustrated by Railton, has struggled to find ways of incorporating information, from a more formal perspective, into its theories. However we have shown how the theory of strongly semantic information addresses the concerns of Railton overcoming the failings of the Bar-Hillel. We considered the issue of the explanation of causal phenomena and concluded that causal explanation relies heavily upon providing an informational account of the relevant factors. My analysis of absences, drawing upon Illari's 2012 work strengthens this argument by showing that information can bridge the gap between cause and effect where an absence is involved. These claims fit

nicely within the methodological framework or general definition of information developed largely by Floridi. Finally I briefly considered the issue of whether we can claim informational objects or informational states supervene on causes. While maintaining the two notions are distinct claimed there is a special relationship between the two. These claims I believe strengthen the case for an information-theoretic account of explanation as well as Floridi's informational realism.

6 Information, Knowledge and Confirmation Holism

This chapter appears as a book chapter in A. Briggle, K. Waelbers and P.A.E. Brey (Eds.) *Current Issues in Computing and Philosophy*, IOS Press, 2008.

6.1 Introduction

It may seem intuitive and self-evident to many that to have knowledge we first must have information. Although few in philosophy have attempted a detailed analysis of the connections between information and knowledge, the intuitive account at least seems to suggest a linear, asymmetrical, and somewhat hierarchical, yet thus far largely unexplained relationship structure existing between *data*, *information* and *knowledge*. Following Floridi (2004) this set of relationships is called *the information continuum.* While much of this thesis examines the nature and role information plays in Floridi's informational realism and information ethics, this chapter considers the relationship information has with knowledge. This chapter contributes to the underlying goal of this thesis in clarifying and contributing to an understanding of the nature of information.

We immediately note that although a 'continuum' *per se* need not explicitly exhibit a hierarchy or bias in any particular direction in this particular case a clear progression is imagined from *data,* perceived as having nominal epistemic value through to *knowledge* recognised as a much rarer and more valuable epistemic commodity.

Most would certainly agree with the commonsense notion that knowledge is something more than information. Furthermore the view is largely uncontroversial and widespread within our technology driven information based society. Such fashionable assumptions however often

offer up a veritable philosophical orchard, ripe for the picking. Whilst the *information continuum*, as a kind of handy model may well provide a practical and convenient scaffold upon which we attempt to build our knowledge and information based systems it does not detract us from a genuine philosophical worry about the precise nature of the relationship between its three constituents. In this chapter I primarily focus on the right hand side of the continuum, the side that assumes an apparent asymmetry whereby information is alleged to precede knowledge. The big picture question worthy of consideration and one that we shall keep in mind (this question is first considered in Floridi's (2004, p570) *Problem 13*) is can we formulate an *information first* approach to epistemology, or in other words can (or indeed should) epistemology presuppose a theory of information? Such an account would be on the face of it, external or objective in nature and thereby not reliant upon any analysis of belief or justification, concepts that have proven to be historically problematic with respect to epistemology.

On the other hand what if it turns out to be the case that the asymmetry is wrong, that information does not always precede knowledge, that instead knowledge at times takes on a much more primordial nature. Clearly there would be philosophical implications, for epistemology, information ethics and the philosophy of information. Furthermore there are likely be significant implications for information technology, which bases much of its practice and scholarly effort upon the assumption. The aim of this paper is precisely to investigate this possibility. In short the paper will present an argument that the relationship between information and knowledge is in fact coextensive. The corollary to this is that the assumption of any strict asymmetry between information and knowledge is at best a convenient myth.

Perhaps some readers even if only having had a brief acquaintance with formal philosophy have heard epistemology's old slogan. To the question 'What is knowledge?' the temptation may be to answer, 'Why justified, true belief of course'. It is indeed true that epistemologists have historically

looked toward belief, justification and truth in their attempts to characterise knowledge and although this originally Platonic conception of knowledge still holds sway with many thinkers, there has been a recent exodus from the position.

The first significant challenge to epistemological orthodoxy came in 1963 with a surprisingly brief article entitled "Is Justified True Belief Knowledge?" (Gettier, 1963). Consider the popularly stated tripartite schema *S* knows that *P* if and only if:

(i) *P* is true
(ii) *S* believes that *P*, and
(iii) *S* is justified in believing that *P*.

Gettier's article provided two counterexamples to the schema in order to show how it could not constitute sufficient conditions for the truth of the proposition, S knows that P. Gettier's examples sought to demonstrate the fallibility of justification, that is, if S could plausibly be justified in believing P but P turned out to be merely accidentally true, then any such construal no matter how prima-facie justifiable could not constitute knowledge. According to Dretske (1981), "the truth of what one believes may be quite unrelated to one's grounds (justification) for believing it."

It seemed we needed a new conception of knowledge. One such approach is to take an externalist view of the conditions required to support knowledge claims. The externalist would claim that the required conditions exist outside the bounds of the psychological states of the knowledge claimer (psychological states such as holding a particular belief). And so the externalist might frame an account of knowledge as follows: K knows that s is F is equal to K's belief that s is F in that it is caused in some special kind of way by the relevant external facts (Dretske, 1981., Adams2004). These relevant external facts, or casual factors must exist external to the machinations of the mind of a knower and so constitute appropriate belief-

independent knowledge-yielding conditions required to avoid the Gettier problem.

Knowledge on this account is therefore equated in some particular way with special relevant causal factors in the sense that the knowledge itself is what justifies the belief, more so than the sense in which it gets justified. Beliefs on such an account could be justified if and only if they are justified by prior knowledge. On the externalist account knowledge is reliant upon the genuine reliability of all relevant objective supporting evidence (we might propose 'information' as fulfilling this role) and not upon any deeper metaphysical analysis of belief, truth or justification.

As our starting point we take as given the general definition of information (GDI hereafter), elaborated by Floridi (2004) as "…namely, the view that semantic information can be satisfactorily analysed in terms of well-formed, meaningful, and truthful data. This semantic approach is simple and powerful enough for the task at hand." There are other accounts of information; Shannon's mathematical theory of communication (MTC) for example considers information as an encoded physical entity. Its main concern is with the transmission and receiving of encoded data. The MTC however does not address issues relating to meaning and reference or how semantic information can be related to truth and knowledge. Since these are the issues most central to this essay we shall leave the MTC aside.

In the section that follows we briefly discuss issues related to the data → information side of the continuum as well as introducing 'Russell's Principle', what Gareth Evans terms discriminating knowledge: the capacity we have to pick out individual objects against background conditions. We then shift our focus to the relationship between information and knowledge. We consider the possibility that information might play a role as objective 'evidence' required for a non-doxastic (non-belief based), externalist account of knowledge.

Contraposing an information-first approach to epistemology is Timothy Williamson's (2000) knowledge-first approach. Williamson defends a

principle whereby one's evidence (what in fact looks very similar to a claim that one's relevant information) is equal to one's knowledge. If Williamson is correct, then knowledge, so it would seem, is in fact either a prior requirement such that the knower could claim to be in an informed state, or if evidence is, as I suggest, just the same as information then on Williamson's account knowledge and information are somehow coextensive. We consider this in detail later in the paper. Whatever the answer it surely cannot be the case that both the information-first and the knowledge-first approach to epistemology are true.

In not too dissimilar fashion Gareth Evans (1982) suggests that in order for an informee to understand any information-yielding event, be it a proposition or the perception of an object, photograph or whatever, some prior information must already be held by the informee. Evans (*ibid*) argues the informee evaluates and appreciates the remark (in the case of a proposition)[39] according to the content of the relevant information and its relationship with information already in his possession.

Williamson (2000) reverses conventional epistemology by arguing that knowing is a sui generis mental state, meaning that 'to know' cannot be explained with appeals to any number or combinations of internal psychological states that the knower might be experiencing for example belief or belief combined with justification, nor in combination with external (for example environmental) conditions. Instead Williamson treats knowledge as prime and he uses this position to develop his accounts of evidence and justification. Williamson however was not the first to suggest the belief / knowledge relationship as conceived by the traditionalists was

[39] There of course are other arguments as to how information might be represented, arguments for example that claim evidence or information is not always necessarily propositional in nature, those discussions are beyond the scope of this article. For the time being I am assuming that all evidence, and all information can be adequately represented propositionally, this interpretation of information squares both with the GDI as well as with our popular computational representations of information, for example information stored in databases.

the wrong way around. The operations of the informational system according to Evans are far more primitive, "It is as well to reserve 'belief' for the notion of a far more sophisticated cognitive state: one that is connected with (and, in my opinion, defined in terms of) the notion of judgement, and so, also, connected with the notion of reasons" (1982, p124).

Thus we see in Williamson and Evans an analysis of knowledge that breaks with conventional epistemology whereby belief was traditionally assumed prior to knowledge. Evans on the one hand derives his arguments from a detailed consideration of the role reference plays in knowledge and information. Williamson dispenses with the notion of knowledge analysed in terms of justified true belief in favour of the thesis whereby knowing is understood as an irreducible mental state equated to evidence. Both are non-doxastic, externally oriented approaches, however both need reconciling with the current philosophical notion of semantic information. This paper draws on the ideas of Evans and Williamson in respect of the apparent asymmetry in the information continuum and concludes with a consideration of confirmation holism, which seems to emerge as a result of the discussions relating to Russell's Principle and the coextensiveness of information and knowledge.

6.2 Gareth Evans and "Russell's Principle"

Quine (1969) makes the distinction between the so-called *theory of reference* and the *theory of meaning*. Central concepts belonging to the later include meaning itself, synonymy, significance, analyticity and, entailment. To the prior belong the concepts of truth, denotation and extension. Note here however that Quine makes no mention of the concepts of belief or justification that are typically associated with epistemology, in fact the role of justification is notably absent in Quine's epistemology, knowledge for Quine was clearly a matter of semantics.

Quine goes on to suggest, "given any two fields, it is conceivable that a concept might be compounded of concepts from both fields", and although

Quine believed the potential hybrid concept in this case to be a theory of meaning, for Gareth Evans the notion of *information* emerges out of this murky conjunction we often call semantics. For Evans information plays a role ultimately critical within his theory of reference. [40]

Nevertheless Quine (1969) seems to look in a different direction in respect to theories of reference and meaning, the problems related to ontological commitment loom large - any given existential statements presuppose objects of a given kind, that is, the sentence which follows a quantifier, is true of some objects of that kind.[41] It is our goal here, however, to examine the apparent asymmetry associated with the information continuum and so following Evans (1982) we draw an important line in the sand: we shall ignore questions of ontology.

> I have supposed myself to be working within a scheme of interpretation for the language which fixes the interpretation of, and hence fixes the objects capable of satisfying, its predicates; the questions which I want to discuss arise after these decisions have been made. (Evans, 1962, p327)

Although Evans's decision to ignore questions of ontology seems somewhat arbitrary we find support for putting ontological issues aside in Floridi's (2004) principle of *ontological neutrality.* Under the GDI the definition of a datum is cast as $d = (x \neq y)$ where x and y are two uninterpreted variables representing a difference. A datum thus is a

[40] Compare this with Quine's "observation sentences", *Word and Object*, 1960 [11]

[41] Quine's concern relates to the commitments we might want to make in respect of the actual mind-independent existence or otherwise of the particular entities we might be referencing in propositions. This is less of an issue for everyday objects but somewhat more problematic for abstract or unobservable entities. Nevertheless we avoid getting into a discussion about 'realism', that is, whether or not the informational entities that we wish to talk about exist independently of the mind.

relational entity however the definition leaves underdetermined the kinds of support required for the implementation of the inequality. That is, a datum is not bound by any ontological commitment. Floridi garners support for this position from a variety of thinkers, perhaps most eloquently phrased by Wiener (in Floridi, 2004, p44) "Information is information, not matter or energy. No materialism which does not admit this can survive at the present day."

Consider for a moment the notion of a datum as a relational entity. Inherent in such a notion is the basic assumption that it is impossible for one to make a judgement about any *thing* at all (a datum, entity, object or whatever) unless he *knows which* thing it is that the judgement is about. A relational entity is a distinguishable object, first and foremost distinguishable from other entities. Following Floridi (2004) that *x* is indeed not equal to *y* represents a lack of uniformity between two signs or signals. Thus in order to pick some thing out as a datum *per se* one must be in a position to disambiguate that datum from surrounding background conditions, which in turn constitute data.

Evans (1982) noted as much and called this "Russell's Principle" and there are some basic interpretations to which we might appeal. Firstly, our commonsense usage of the term '*knows which*' does not automatically assume to presuppose that the individual does in fact know which item he is making a judgement about – it maybe that the individual merely perceives a difference. Alternatively information may originate from an object *x* but the content may not refer to the object particularly well just in case some error or problem occurs either at the source or with the information transmission.

On the other hand we could simply argue that none of this matters. If indeed there is some thing *x* that the individual has observed and is consequently considering, then whether *x* is a *P* or *x* is a *Q*, doesn't alter the fact, and quite probably the knowledge, that it is *x* that he is considering. Thus we can hold that two individuals observing the same object *x*, to be in

the same informational state but that they may not be holding the same content.

The knowledge required to make the distinction between *x* and something else *y*, or to put it another way, the knowledge required to pick out a datum following the GDI definition whereby $d = (x \neq y)$ is what Evans (1982) calls *discriminating knowledge* this being a subject's "capacity to distinguish the object of his judgement from all other things". In this case the concept '*knows which*' is merely the ability to notice the difference. Following Floridi (2004, p45), "the GDI is neutral with respect to the identification of data with specific relata", that is the black dot or the white sheet of paper upon which it appears, although individually distinct require each other as relational entities in order to qualify as datum, the ability we have to make this distinction would seem fundamentally critical to any understanding of information. Floridi (2004, p44) quotes Bateson who stresses the point, "what we mean by information – the elementary unit of information is a difference, which makes a difference". A difference is just a discrete state, and the making of a difference just means that the datum is potentially meaningful.

There is more to this point, at the most fundamental level we are talking about the ability to formulate a basic and primitive concept of an object (or a state of affairs). Knowledge for what it is for '*d* is a datum' to be true necessarily, if indeed $d = (x \neq y)$, requires at least some primitive notion of what differentiates *x* from *y*. That is by definition, we require distinguishing knowledge about *x* since *x* is differentiated from other objects (or states of affairs) by the fact that *x* stands in its own right as a 'discrete state' (a datum).

According to Floridi "the most important type of semantic information is *factual information*, which tells the informee something *about* something else" (2004, p40). Implicit within the notion of factual information is the concept of identification or reference, a way of knowing which object is the one in question, summed up by Quine (1969), in "no entity without identity".

For only if one is able to correctly individuate or express distinguishing facts about an object from which the information derives can we assume then that that factual information does indeed say anything meaningful at all.

At this point we are merely flagging (not resolving) the issues raised by Russell's Principle as requiring a thorough reconciliation with the GDI as semantic content. The problem turns on the fact that a person's knowledge of a particular object relies upon a capacity to cite discriminating facts about the particular object from which the information derives. A denial of Russell's Principle would seemingly render a subject unable to identify a datum in the first place due to the inability to comprehend the side of the equation ($x \neq y$) regardless of *d*. One must know there is a difference in order to ascent to as much.

6.3 Can Information and Knowledge be Coextensive?

In *Phenomenology of Perception*, Merleau-Ponty (1962) argues "Empiricism cannot see that we need to *know* what we are looking for, otherwise we would not be looking for it" (italics added). Although Merleau-Ponty[42] no doubt had a much larger agenda in mind, namely to expose the problematic nature of some of philosophy's traditional dilemmas, we take his suggestion as permission to proceed in questioning prevailing *doxa*. In this section thus we explore the possibility of information and knowledge as being at least in some cases coextensive.

Coexstensivity between information and knowledge might be outlined as follows:

a. *K* (Knowledge) is a determinant of *I* (Information) if and only if *K* is functionally dependant upon *I*.

and

[42] A more thorough reconciliation of Merleau-Ponty's work in regard to externalist or a "god's eye view" of information is beyond the scope of this paper, however I thank the anonymous referee for drawing this interesting connection to my attention.

b. *I* is a determinant of *K* if and only if *I* is functionally dependant upon *K*.

In terms of the information-first approach to epistemology and excluding any appeals to *a priori* knowledge, *a* is relatively uncontroversial. I say relatively since further work is certainly required in order to thoroughly establish an information-first epistemology if indeed such an approach turns out to be correct. However more straightforwardly, *a* is intuitively uncontroversial compared with *b* because it takes for granted the popularly held assumption that an asymmetry between information and knowledge holds, whereby information naturally precedes knowledge, as per the continuum hypothesis.

In order to clarify *a*, consider the following brief example. The fact that I know that it is raining, for example, can be typically determined either by all the information considered relevant to such a judgment or by other information in the form of testimony provided to us by for example the meteorological office. My knowledge therefore that it is raining is functionally dependant upon the information germane to the fact. Such approaches appeal to our intuitive sense of information as being prior to knowledge and are typically expressed in what Quine calls observation sentences (1960). Information-based thoughts of this kind (observation sentences) expressed by propositions such as "it is raining", are least susceptible to variation under other informational or epistemic influences, they seem to be basic propositional attitudes.

On the other hand *b* is certainly controversial since it hints at a much more primordial interpretation of knowledge. Nevertheless at this point (Merleau-Ponty's concerns aside) we have identified at least one case whereby b seems to hold. The identification of any particular object or state of affairs at least requires the satisfaction of Russell's Principle. That is my making a judgment about the fact that is raining, regardless of the information I receive requires me to know what it is I am looking for. To explain further, the fact that 'it is raining' and that this is what my

information based thought is focused upon requires a prior knowledge that the raw information I am immediately receiving from the environment (if indeed that is the mode by which I am receiving the information) concerns the fact that it is raining and not some other state of affairs – the judgment requires discriminating knowledge. As methodologically ideal as it maybe we do not operate as detached conscious beings observing and interpreting specific and discrete, epistemologically self-contained brute facts about the world.

We now introduce another case whereby knowledge appears to be in some kind of coextensive relationship with information.

Following Evans (1982) an information-based thought about an object or state of affairs needs to be either a fundamental identification of that object (or state of affairs), for example 'it is raining' implies in some way that I know about rain in some essential way, or alternatively, consist in a knowledge of what it is for an identity proposition involving a fundamental identification to be true. Quite simply the necessary truth of any information-based, knowledge statement must entail the truth of the information in the first place (it couldn't be the case that 'false information' entailed any kind of knowledge), but the truth of such information surely could only be determined with a prior knowledge of what it might take for that information-based thought to be true. This is reiterated by Williamson (2000) "knowing is the most general factive attitude, that which one has to a proposition". One either knows it is raining (or not), or doesn't know, or isn't sure, in the latter two cases the subject does not know at all.

A significant portion of Gareth Evan's 1982 text is committed to explaining demonstrative identification. Although we are constrained from explaining Evan's position in detail we can say that the core idea revolves around the concept of what Evans (1982) calls the information-link between a subject and an object whereby the subject is provided with "information about states and doings of that object over a period of time". Ordinary demonstratives express informational links between objects and subjects, to

the objects they identify via propositions such as ‘that man’, ‘this chair’, thus a demonstrative thought is an information-based thought.

I am not arguing that there are any identification guarantees in ‘demonstrative identification’ even given a clear and present information-link between object and subject. An information-link although necessary will not be sufficient to determine the truth of a demonstrative without some controlling idea or concept of the object in question. As we mentioned earlier, whether x is a P or x is a Q doesn’t alter the fact that x is the origin of the information, yet clearly there is a difference between a basic ability to disambiguate x from other objects in the immediate vicinity and a corresponding ability to assent to the truth of x being either a P or a Q. The latter requires the knowledge of what it is for any particular object to be correctly identified as the relevant object, following Floridi, “for example where a place is, what the time is, whether lunch is ready, or that penguins are birds” (2004, p45). Consider this in contrast with incorrectly identifying a place as some other place, or some other incorrect time, or that a penguin was some other kind of animal.

Continuing with Floridi (2004), “factual information is declarative in nature [and], is satisfactorily interpretable in terms of first-order classic predicate logic”. Furthermore with appeal to Dretske (1981) and Barwise & Seligman (1997), information carries with it the mechanism such that demonstrative identification can be satisfied “a’s being (of type) F carries the information that b is G”. This mechanism clearly represents some kind of confirmation relation between the information-based statement and the object in question.

At this point we ought to highlight an important distinction. We can agree that information exists prior and independent to any mental processing that turns it into an information-based thought or statement (subsequently representing demonstrative identification). This is the case to the extent that ‘non-thinking’ organisms utilise ‘information’ in a behavioral kind of way, the sunflower orients itself towards the sun, single cellular organisms

capable of interpreting environmental information have a survival advantage over those that are less able to interpret environmental information, likewise humans utilise various complex information channels to enable us to accomplish tasks such as riding a bike, swimming or catching a ball. This use of information however seems to be somewhat different to our understanding and previous discussion regarding 'semantic information' in that there is clearly an absence of cognitive relations between subject and object.

On the other hand our wish is to try to establish some connections between information and knowledge, we remind ourselves of the big picture question, the speculative thought as to whether we could have an information-based epistemology. Although we don't wish to muddy the waters in respect of the ontological nature of a pure mind-independent information uninterpreted by any human mind as being somehow the same as our information based thoughts or the kind of information passed from informer to informee in the form of testimony, demonstrative identification or, 'factual information of a declarative nature' we find it difficult to discuss the one without the other. Semantic information is a human concept in that it involves the concepts of reference, identification and meaning, according to Colburn, "to be information requires a thinker" (2000). Such concepts of course are quite different to those of justification and belief and I see no problem in maintaining a non-doxastic, externalist position given such bounds.

Thus we have seen that there are at least two considerations whereby a coextensive relation appears to exist between information and knowledge. The first being the satisfaction of Russell's Principle and the second being an individual's ability to correctly identify objects or states of affairs as consisting in a knowledge of what it would take to make any proposition directly related to the objects or states of affairs in question true. Information then, it seems, carries with it what we might call a confirmation relation. The confirmation relation acts between the (correct) identification

of objects (or states of affairs) and the agent utilising the information-link. Furthermore, confirmation relations themselves are a posteriori (otherwise we would have to run an argument as to how we could identify objects innately) in that they represent an empirical truth between the information-based statement (or thought) and the object to which it refers.

6.4 Confirmation Holism

The remainder of the discussion will focus on what implications *confirmation holism* might have for the philosophy of information in light of the above discussion. Confirmation holism is the premise that all theories and consequently the statements that make up those theories are underdetermined by the data or evidence supporting them (1992). The idea here is that statements or theories can never be tested in isolation. That is, we are always reliant upon other information, or knowledge in order to confirm (or disconfirm) any scientific or factual claims. Furthermore this implies that there will always be competing statements, observations or propositions that may be used to prop up a theory. Thus, in principle it is not possible to fully refute a theory based on the falsity of any individual premise encapsulated by that theory. Let us explain why this is important.

In adopting an external, information-theoretic approach to epistemology, an approach that does not demand any analysis of justification or belief in order to support its claims, we are necessarily committed to *a posteriori* or an experience-dependant account of what confirms what in the world.

The reason for this follows much of the discussion above: we are talking about mind-independent, externally generated, information-links between objects and subjects, that is, factual information of a declarative nature as representative of the confirming evidence required to make knowledge claims. Information is conveyed via perception and transmitted via language (Dummett in Floridi, 2004). The factual component of an information-based statement surely must depend on the empirical confirmatory experience to which it is intimately linked. Furthermore and

perhaps more importantly, the confirming methodology *itself* is also *a posteriori*. This is the case since every declarative factual informationally based proposition is itself reliant upon, firstly the knowledge that $d = (x \neq y)$ or Russell's Principle, and secondly an individuals knowledge of the meaning of any given information based proposition, following Quine (1953) "the meaning of a statement is the method of empirically confirming or infirming it".

The alternative would be that confirming conditions between objects and their corresponding identifying (information-based) statements are somehow *a priori*. Such an approach ends up having to carry the burden of internalism and is thus obligated to appeals to the analyticity of confirming conditions – and the quick route back to a foundationalist analysis of justification and belief. Evans (1982) echoes Quine's (1953) thoughts, in that any information-based thought must be either a fundamental identification of the object in question or require knowledge of what it takes for such a proposition involving a fundamental identification to be true.

Williamson raises the issue of holism as he considers the implausibility that knowledge actually consists in "epistemically self-sufficient nuggets of information" that might exist in isolation from each other (2000). The suggestion is that if information-based statements are not somehow linked to other information-based statements in confirming knowledge claims, then they can't be confirming anything other than themselves. The *what* and *how* of any such informational connections that might exist between all the relevant information required to make knowledge claims is a story for another time; as is an answer to what mental apparatus might be required to establish such connections. Williamson (2000) contributes to the suspicions in conceding to a form of holism in quoting Peirce, "I cannot make a valid probable inference without taking into account whatever knowledge I have (or, at least, whatever occurs to my mind) that bears on the question". If an information-first epistemology relies upon all the relevant information-based propositions pertinent to the particular issue as the necessary

confirmation relations in respect of knowledge, then it is difficult to see how such statements could exist in isolation.

6.5 Conclusion

This chapter has dealt with the relationship between information and knowledge. The main aim has been to critically examine what is often termed the information continuum, that is, the idea that information is prior to knowledge. The information continuum gives rise to Floridi's *Problem 13*, should epistemology be based upon a theory of information? (2004, pg 570). It is primarily this "problem" that the chapter considers.

I have argued that any interpretation of the continuum as being strictly asymetrical, that is, in favour of information being prior to knowledge, is doubtful. While there is nothing wrong with the notion of information and knowledge being on a continuum *per se*, it is more likely that they are co-exstensive relations.

Furthermore, it seems that the consequence of externalism in respect to an *information-first* approach to knowledge is confirmation holism. That is, we rely upon *other* information or knowledge in order to confirm (or disconfirm) independent, information based statements. I have argued that confirmation holism applies equally at this level as it does for entire scientific theories

In conclusion, the possession of information conveyed by perception and readied for transmission *via* language requires prior knowledge on the part of the subject, regarding the object in question, of what it takes for a proposition involving that object to be true. Furthermore it requires the satisfaction of Russell's Principle, or what Evans calls, in agreement with the Floridian definition of a datum expressed as $d = (x \neq y)$, discriminating knowledge. The discussion here does not necessarily rescind the information-based approach to epistemology, what it does do however is raise concerns over our intuitions of any presupposed asymmetry between information and knowledge.

7 Information Ethics and Entropy

This chapter was presented at the 2012, *Australasian Post-Graduate Philosophy Conference*, University of Auckland.

7.1 Introduction

Amongst philosophically inclined ethicists, Computer Ethics (CE) has been accused of being a kind of "pop ethics" (Bynum, 1991), that is, lacking a conceptual and rational foundation, unsystematic, ad hoc, with analysis often performed after the fact. This problem has been called the foundational problem by Floridi & Sanders (2002). The typical analytical approach taken by industry media and various commentators familiar with technology advancements has been to debate technological and computer based ethical problems from the point of view of the "unethicality" of the problem (Parker, 1981), or the perceived harm that is being done. This usually takes the form of a critique of a particular incident involving computers, databases, smart phones, social media, their use and abuse across various private and public social contexts. Many of the ethical and moral problems occurring within an Information and Communication Technologies (ICT) context are not considered prior to their emergence hence the responses are sometimes seen as kneejerk and lacking any ethical foundation. But it isn't all bad. Over time such dramatic stories outlining ethical issues within ICT and offering some form of response have evolved and have helped the development of CE into a somewhat mainstream discipline at least with IT circles. As a result ethical problems related to ICT are indeed resolved and legislation is enacted.

Luciano Floridi however has been concerned with what he terms the foundational problem (Floridi 1999, Floridi & Sanders 2002). While the pop ethics approach provides some sensitisation to ethical problems and as

such can be a useful starting point, at least from a pedagogical perspective, it lacks a methodological basis according to Floridi.

In response to the foundational problem Luciano Floridi and Jeff Sanders (F&S) have developed Information Ethics (IE) (Floridi, 1999, 2002, Floridi and Sanders, 2001, 2002). *Prima facie* this new and novel theory appears to offer some hope in providing a robust platform upon which moral issues involving ICT can be ethically analysed. Nevertheless the theory has proved controversial and is not without its critics. This paper identifies and explores what I consider to be a significant challenge for the theory. The problem relates to Floridi's use of the term entropy which is quite at odds with various other interpretations of the term. Floridi's conception of information entropy appears to be by and large the ethical equivalent of the uses of evil, pain and suffering in more orthodox ethical theories. Thus, informational entropy for Floridi is the loss, degradation, corruption or depletion of information objects. In this paper I shall argue that Floridi's use of the concept of entropy is somewhat confusing but more importantly sets up a troublesome paradox within his Information Ethics (IE). The paper outlines the problems before presenting a possible solution.

The paper shall proceed as follows. In section 1 we briefly outline Information Ethics, its scope and aims. In this section we also explore Floridi's use and meaning of the concept of entropy. Although not intended to be a comprehensive review of IE it should provide us with enough of a background to proceed. Any further detail will be explained as we go. In section 2 we turn to the issues with informational entropy which form the central part of this paper. We identify these issues and consider them in some detail. Section 3 looks at how the problems might be addressed.

7.2 Information Ethics in Short

Floridi describes IE as an ontocentric,[43] patient-oriented, ecological macroethics (1999). In order to understand Floridi's motivation it is worth noting previous popular ethical accounts. For example, biocentric ethical approaches recognised *life* in general as having intrinsic moral worth. Humans on such accounts are not the only moral targets but all life regardless of sentience contains some inherent worth. Thus humans merely form part of an ecosystem within which all species are morally or ethically valued. Similarly ecocentric ethical approaches consider entire ecological communities as being equivalent to ethical communities. Such ethical models move beyond anthropocentric approaches and are intended to compliment and extend traditional ethical theories thus enabling analysis for human behaviours directed towards non-human entities.

Taking a step further than bio or ecocentric approaches IE has an ontological focus, thus is said to be ontocentric. Ontology is concerned with questions as to what entities can be said to exist and how such entities can be categorised. The principal claim of IE is that *being-in-itself*,[44] or bare existence is fundamental and all things in existence can be described as informational objects. Much of the controversy surrounding IE relates to the fact that IE vastly enlarges the idea about what can count as an object worthy of moral consideration. Indeed existence, on the ontocentric account,

---

[43] The use of the term "ontocentric" in relation to information ethics first appears to be in Floridi's 1999 paper, *Information Ethics: On the Philosophical Foundation of Computer Ethics*. The term was a response to popular bio and ecocentric ethical theories of the time given that IE was proposed as an ontologically focused ethical theory. Floridi argued that the most primitive common factor of all existing entities is their information state. That is any form of being can be interpreted as a "coherent body of information". He goes on to argue, "that Computer Ethics is infocentric is tantamount to interpreting it, correctly, as an ontocentric, object-oriented theory." (1999, p43).

[44] Philosophers may be familiar with the term *being-in-itself* from existential philosophy. I like the Sartrean form of the term where *being-in-itself* refers to bare objects in the external world, representing simply basic existence, not conscious, nor active or passive. While Sartre excludes humans from an informational level of abstraction it's not difficult to perceive humans as informational objects in the Sartrean sense.

is considered intrinsically valuable. Thus rather than asking the question central to traditional anthropocentric ethical theories, "what ought I be", or "what ought I do" instead our enquiries ask "what ought to be respected or improved".

This raises an important and distinguishing feature of IE. "What I ought to be" rests upon what kind of people or things are worthy of respect. Thus once we have agreed upon a set of characteristics universally considered worthy of respect, then we can determine what kind of person for example we should aspire to be. Similarly the idea surrounding "what ought I do" rests upon the set of actions we agree that might be respected.[45] IE extends the boundary of what is worthy of respect beyond anthropomorphic considerations to the limits of knowledge and existence – everything, due to its composition and interpretation as an information object is worthy of respect, at least in some minimal and possibly overidable way.

Thus while biocentric approaches view *suffering* as an intrinsically negative concept, the equivalent concept in IE is *entropy*, or the degredation, destruction or corruption of existence. These two core issues, *entropy* and to a lesser extent *intrinsic worth* form the critical part of this essay but first let us develop a fuller picture of Floridi's and Sander's IE.

We note that the central claims of IE are particularly strong. IE is binding for all agents or "patients". A patient need not be a human, in fact a patient is considered by Floridi to be any indeed all entities in existence, from a piece of lint to a galaxy, a great master's painting, tables and chairs, or an overflowing ashtray. All such things count as moral patients and according to IE have intrinsic moral value, albeit at times minimal and overidable (Floridi, 2007, p21). *Prima facie* it seems absurd that respect, let alone intrinsic value could be conferred upon certain objects, thus the obvious controversy arises out of the seemingly counter-intuitive extension

[45] I thank Morgan Luck for pointing this out to me in an earlier draft of this paper.

of ethical consideration bestowed upon and "universally binding on all agents in all places at all times" (Floridi, 1999, p49). How does Floridi address issues arising out of this controversy?

Important in understanding IE and its ontocentric focus is an understanding of an issue fundamental to Floridi's Philosophy of Information that is, his use of the concept of *levels of abstraction* (LoA). All entities (agents, patients) according to IE are analysed as informational objects via an informational LoA. The idea of the informational object is described fully in Floridi's informational ontology where information objects are considered to be discrete, encapsulated clusters of data exhibiting their own properties or attributes, behaviours and functions as well as relations within themselves or with other entities. Floridi's explanation of the informational object and his use of the LoA methodology draws heavily upon object oriented programming and computing theory; this creates some problems which have been recently explored (see McKinlay, 2012).

The levels of abstraction concept is a process borrowed from computing that enables systems to be analysed whereby varying layers of complexity can be hidden from the analysis depending upon required use or purpose. Consider this example from computing, database design to be specific, in order to illustrate the LoA idea. We start off at a high (a design or conceptual) level of abstraction where we might develop a simple graphical model of an information system. The relevant informational objects and their relationships with each other are represented simply as graphical symbols. At a lower LoA this model is transcribed, following a set of agreed rules, into some form of computer code (software), this might be a declarative script style of language (in Database SQL for example) or a fully compiled object oriented or procedural language. The method allows design and development from conceptually simple models through to more complex representations. In computing the lowest (physical or materialist) LoA is generally an arbitrarily complex coded implementation on computer

hardware. Each LoA serves a particular purpose. It is not the case that we must choose between differing levels, nor that any particular level is more valid or closer to reality than another level. It is more the case that at each level a model is implemented by and representative of the ontological components available at that particular level. Following Floridi (2008) by stating the LoA a theory makes explicit its ontological commitments. Certainly there is a degree of correspondence across levels. In our database example for instance each level helps us construct a model at a lower level – the result being a database implemented on computer hardware. The LoA method constrains our view of the system in question to a level that is appropriate for the kind of analysis required or the purpose desired.

By observing entities ontologically and via a level of abstraction that allows us to analyse them as informational objects Floridi argues we lift the status of such objects to agents or patients of actions relating to environmental processes, changes and interactions all described at an informational level (2007, p10). Floridi argues that this informational ontology is a matter of metaphysics, not science and as such an informational view of the universe is on a par with other scientific methods of explanation, chemistry or physics for example. What Floridi is trying to say here is that an informational explanation of the universe is as valid in terms of the explanatory framework it offers as for example a quantum physical explanation. Thus according to IE, all being, all informational objects are conferred intrinsic value by virtue of their mere existence and thus all informational objects deserve the right to flourish within their global environment.

The "global environment" in IE is described as the *infosphere*. That is, the global environment viewed at an informational LoA where everything that exists is analysable as an informational object. Whatever is in the infosphere, according to Floridi (2007, p22) is informational, and whatever is not in the infosphere is something that simply cannot be.

IE ordains as moral agents we have significant responsibilties to the infosphere. This responsibility extends not just to other informational objects which maybe living, human or not but also in terms of our abilities to control, mediate, create, destroy or degrade, exploit, model, and in relation to ourselves, genetically, physiologically, socially and so forth, within natural, virtual or artifical environments. Our duties as moral agents according to IE, is not only to oversee our own development of character and habit but to monitor the well being and flourishing across all spheres of influence included within the entire infosphere. This represents a clear movement away from individual virtues towards global values. (Floridi, 2007, p23)

The determinants for moral judgements in IE are layed out by Floridi as four overriding straightforward principles. Entropy is considered the most fundamental negative intrinsic value to be avoided, similar to *pain* in medical ethics, *evil* or *suffering* in more generalised ethical theories. Floridi's use of the term entropy differs slightly from more physical theories of entropy, indeed for Floridi entropy is itself analysable at an informational LoA. I discuss this in further detail shortly; however, first we outline the patient-oriented ontocentric principles thus completing our brief overview of IE. The ontocentric principles are as follows:

0. Entropy ought not to be caused in the infosphere (null law)
1. Entropy ought to be prevented
2. Entropy ought to be removed
3. The flourishing of informational entities as well as the whole of the infosphere ought to be promoted.

The best moral action according to Floridi is one in which all 4 laws are satisfied. Additionally at a minimal level any action that never generates any entropy in the course of its implementation is commendable. Although

any action that merely satisfies the null law (entropy is not caused) is morally neutral, irrelevant or at least insignificant.

7.3 Previous Important Objections

The above section sufficiently outlines Floridi's Information Ethics for our purposes. From here on any further detail will be discussed as it arises. We keep in mind that the purpose of this macroethical theory is by and large to overcome the foundational issues associated with Computer Ethics thereby providing a conceptual backdrop for the application of ethical principles to ICT but further, IE is meant to offer a significantly robust ethical theory in its own right. Also important to note is the point that Floridi in no way proposes IE as a wholesale replacement of other familiar macroethical theories, rather it is meant to augment and extend the standard ethical theories, equiping us with the conceptual and foundational framework that standard computer ethics approaches lack. However, this goal makes the development of a meta-theory no mean feat. Any foundational theory regardless of discipline will need to be robust enough to withstand vigorous criticism.

Indeed IE has suffered from such criticism. Here I outline two recent and significant criticisms that are a starting point for the rest of this essay. To begin with, Stahl (2008) argues that IE is undermined if not contradicted by a form of relativity. This relativity he argues is inherent in Floridi's use of the concept of levels of abstraction. To explain further, if a particular set of ethical claims is dependent upon analysis from a particular LoA then can we really defend these claims under a universal ethical theory? A problem arises when consideration is given to alternative claims or arguments developed with another LoA as their reference point. Particular sets of claims, Stahl wants to argue, are *relative* to the LoA that they are developed within. This problem is particularly evident when clear contradictions regarding ethical claims arise across differing LoA. Stahl's main concern appears to be in the interpretation of humans at the informational LoA as

"*inforgs*" or informational objects and that this somehow diminishes human dignity. "Ethically this could be seen as problematic because it could suggest that humans are fundamentally the same as everything else, e.g. trees, stones or over flowing ashtrays, and therefore human dignity would cease to exist" (Stahl, 2008, p104).

To Stahl's complaint, surely we are not *just* information objects, Floridi responds:

> We are not just inforgs, but we should not fear to consider ourselves as inforgs. This is the sense in which IE is universal. It is inclusive in the sense that the logician will immediately recognise as part of the extensional meaning of universal quantification: all entities are informational in nature, and IE seeks to address the ethical issues that pertain to all of them. (2008, p191)

Stahl goes on to discuss how we might choose between different LoAs for particular analysis:

> As the choice of the LoA is not determined and given that it determines ontological commitments seems that an agent is free to choose, albeit indirectly, the ontology of a phenomenon. This allows the interpretation that an agent gets to determine which information entities exist (are relevant) and therefore need to be considered in ethical decision making. (Stahl, 2008, p100).

Part of the problem with Floridi's use of the term entropy within the context of his information ethics is that he seems to muddle between levels of abstraction. At times he talks of entropy strictly analysed from an informational level, yet at other times he seems to revert to a pragmatic analysis that seems to be in line with a more physicalist understanding of macro objects. In the pages that follow I discuss this in much more detail but first I will briefly review the method of levels of abstraction.

The LoA chosen is the one fit for the purpose of the analysis. With this I believe Floridi would agree. If for example I am building a model of a database, I choose appropriate tools and constructs that fall within the ontological bounds of the data modelling exercise. The constructs, their use, and the ultimate meaning of the model collectively describe the LoA that is appropriate to use for such purposes. However, another LoA is employed when actually implementing the model, SQL, for example in data modelling. Yet another LoA is in operation when considering the real world entities and relationships that the database ultimately represents. It is however true, certainly within computing frameworks' that each level has relatively clear correspondence relations that connect them, that is, a database programmer will understand the entities and relations modelled at the modelling LoA and be able to translate these quite easily to the SQL implementation LoA even though they differ in purpose and semantically. How relationships work or map across different LoA to an informational LoA in Floridi's work isn't always clear.

Nevertheless I agree with Floridi's universality thesis. I have no problem accepting that we can analyse reality from an informational perspective. This is little different to analysing objects from say the perspective of physics, in doing so we might talk about the mass, electrical charge, dimensions and perhaps the atomic make up of the object. An informational story is equally valid, it is not a question of which LoA is the *correct* one, the correct one is the one related to the kind of analysis one is performing. The question as to whether or not we can "catch" all the information, is however open. As a philosophical problem this issue differs little from problems in the philosophy of science relating to the reliability of science and whether or not it tells the true story of reality. While such questions are interesting they are not really relevant to our discussion here.

Stahl does raise an issue related to an analysis relying upon a particular perspective being relevant (i.e. the particular LoA being employed) and I will draw upon and demonstrate how, in relation to Floridi's concept of

entropy such concerns do seem to pose a problem for IE. The following issue I raise while inspired by Stahl's problem, differs slightly.

It is certainly the case in computing and in other scientific and technological disciplines that we can identify more or less explicit links between different levels of abstraction. That is, we are to be able to reconcile or map across levels of abstraction without significant contradiction. This was made clear in the database example above but it can also be demonstrated across physical, biological and biochemical examples in science. It seems evident to me that should we stumble upon apparent and or significant contradictions between analyses at various levels of abstraction that this must point to some kind of mistake. To be explicit with regard to IE, if something is to be conferred intrinsic worth at the informational level, then some worth ought to be apparent at other levels. For example, Floridi would no doubt argue that the overflowing ashtray analysed at an informational level has some minimal (overideable) intrinsic worth. For certain informational objects at least, intrinsic worth is minimal and ultimately overidable (2002). How we are to decide upon which informational objects are only to be afforded a minimal and overrideable value is not explained by Floridi. I suggest that we are generally informed as to which entities can be passed over via an alternative commonsense everyday LoA, that is, one that generally considers the detritus of human activity for example as having no value whatsoever. Furthermore I consider the possibility later in this paper that such decisions are ultimately consequentialist in nature and that others seem more likely to be resolved with appeal to ecocentric rather than informationally ontocentric arguments. However first I outline another recent objection to Floridi's IE relevant to this work and related to the present discussion.

Brey (2008) raises the question of the validity of conferring intrinsic moral worth upon all informational objects. Once again the Floridian notion of LoA is his starting point Brey argues, the claim that all information objects, that is, everything in the infosphere, including all extant objects,

states of affairs, and constituent states of affairs, have intrinsic value relies upon those entities being analysed at the level of abstraction that supports the claim in the first place. In a sense Floridi is question begging according to Brey. He argues, IE simply claims that all objects *qua* information objects deserve moral respect without adequate explanation as to why this should be the case. Brey goes on to argue that in essence IE collapses to some form of bio-centric ethics.

There are some parallels in my argument to Brey's but what I am most interested in is the apparent contradiction across levels of abstraction with regard to intrinsic worth. Brey's argument states,

> IE tells us that we should be equally protective of human beings and vats of toxic waste, or of any other information object, and that we have an (albeit overridable) duty to contribute to the improvement and flourishing of pieces of lint and human excrement. At best, this suggests that IE gives us very little guidance in making moral choices. At worst, it suggests that IE gives us the wrong kind of guidance. (2008, Criticism, para 7)

Clearly at some practical day to day level (LoA) we could not imagine conferring any intrinsic worth at all upon Brey's colourful examples, the thought *prima facie* is intuitively absurd, yet history does instruct to be wary about jumping too quickly to such conclusions. For example it may have seemed patently absurd to confer moral respect on animals merely a few decades ago. Yet while Floridi's ontocentric principles prescribe at the informational level such objects deserve to be respected, at least minimally as information objects, there is certainly a sense in which we would tend to think this cannot apply to all examples. The question is how we are to decide which objects are worthy of more respect than others. I believe we

can resolve this problem with a careful reconsideration of Floridi's concept of entropy but first a brief summary of the objections I raise.

7.4 Present Issues

The first issue relates to Floridi's use of the concept of Entropy. The ontocentric principles instruct us that we ought to prevent, remove and not cause entropy in the infosphere. However almost all human activity (even analysed at an informational LoA) is entropy causing. Entropy operates in one direction and is irreversible, at a most basic level it is the necessary evil associated with the preservation of our internal order and structure. We had to chop the tree down to burn a fire to keep warm and to cook our food. Wherever physical changes take place entropy is at work.

This being the case we are ultimately faced with making decisions based upon which entropy contributing behaviours are more or less morally culpable. It seems this collapses IE to a form of consequentialism – that is we are faced with making decisions about which entropy causing activities are we able to live with. While all entities *qua* informational objects have some minimal moral worth, there are some cases where the moral worth of particular entities is minimal and overridable. It is quite perplexing as to how decisions might be made purely at an informational LoA, regarding which entities IE can overlook. Such decisions seem to be made not with regard to an object's status as an information object but at some generalised LoA where we make everyday decisions about the status of objects considered of no value, for example dust, empty food packaging, email spam, previously watched TV shows stored on hard drives and so on.

A more problematic issue however I think is this: If IE is to make any real sense, we must be able to reconcile an analysis of information objects at an informational LoA with an analysis at a different LoA that does not result

in a significant entropy attribution contradiction. In other words, if a particular object at an "everyday level of abstraction" (ELoA)[46] is the cause of considerable entropy (take for example a coal fired power station) then it is difficult to reconcile how at an informational LoA such an object might be exonerated. The problem that arises is that IE does not appear to offer guidance as to how we ought to make decisions about which entropy causing behaviours are less morally blameworthy. Certainly at the thermodynamic LoA, a coal fired power station is essentially an entropy creation machine, taking an energy source from the ground and converting that resource into another form of energy (electricity) and heat with a measure of resultant entropy – such processes epitomise all general descriptions of entropy. Yet the information object that is the coal deposit as well as the surrounding ecosystem is clearly being degraded and corrupted as a result of the mining activity. At an informational LoA entropy is being increased in exchange for energy which might be used to create or sustain order and complexity elsewhere.

How are we to decide between which entropy causing actions or behaviours are morally permissible? Given this problem it would seem our ethical decision making process ultimately collapses IE into some version of consequentialism. In the end we must decide which set of consequences based on the decisions we make that we can live with and these decisions are typically made at other non-informational LoA. At the LoA of where entities are analysed purely as informational objects, in the case of a coal fired power station, it's hard to see how informational entropy can apply in a way that is significantly different to an analysis of entropy at the thermodynamic level. If it were to differ significantly then it is somewhat

[46] By "everyday level of abstraction" I simply mean understanding of a system at a general macroscopic level where we would generally all agree upon the observables and the behaviour of that system. Floridi uses the example of an analysis of traffic light. At the LoA which see us all driving around obeying the road code we can all agree upon a LoA that directs our driving behaviour based upon the colour's *red*, *amber* and *green* of traffic lights.

contradictory that moral value can be minimal and overridable for some informational objects since the minimality and overriding conditions seem to be informed from an alternative ELoA, one that would naturally consider some entities worthless – e.g. toxic waste, lint, excrement and the like.

Secondly, related to the first issue and in some senses this represents an extension of Brey's 2008 argument from ontological relativity. Brey's argument from ontological relativity runs as follows:

> For something to possess intrinsic value, it must possess one or more properties that bestow intrinsic value upon it, for example the property of being rational, capable of suffering, or being an information object. Such properties have to be objective and inalienable properties of the object in question, not subjective or contingent ones, because otherwise the assigned value is (at best) extrinsic, that is, resulting from the attribution of contingent roles or subjective meanings to objects. The property of being an information object, I will argue, is not an objective, inalienable property of objects but a contingent or subjective one. (Brey, 2008, p9)

Let's accept Floridi's claim that all things *qua* information objects possess intrinsic value. In doing so however there is still a good case for arguing that some informational objects will simply never qualify as being intrinsically valuable, based upon the fact that they themselves are significant entropy increasing entities or that they are the waste product or the resultant entropic product, the result of some other process. This is certainly the case from a LoA that takes into consideration prior processes and/or activities that were a part of a process from which the entity in question emerged. As such by virtue of their existence within the infosphere and at an informationally analysed LoA, such entities demand elimination, or at least ought to be prevented in the first place. The

Floridian concept of the information object, in many cases, and in particular artifactual cases (ie, with regard to man-made entities) sets up a troublesome paradox. If every extant object is at some fundamental level an information object, which under IE deserves some minimal moral respect, then all entropy causing events or entities both deserve respect and ought to be prevented or eliminated under IEs onto-centric principles, at the same time. Let us now examine these issues in further detail.

7.5 Entropy

In order to clarify his particular conception of entropy Floridi offers an allegory (1999, p54). He asks us to consider a young boy playing in a dumping ground. The boy amuses himself by throwing rocks at abandoned old broken down cars, breaking their windscreens and lights. Of course this is precisely the sort of behaviour that young boys enjoy, yet no doubt some of us may consider the boy's time better served in doing something more constructive. While we might concede such childish amusement is largely victimless – the boy isn't doing anyone any harm, after all the old cars have been dumped, our tendency might be to consider such behaviour *vandalism*, a term loaded of course with considerable negative moral judgement.

Floridi considers different ethical perspectives in light of this case. He dispenses with the Kantian or deontological approach with its ends/means dictum since (he argues) it cannot directly apply where human interaction is not present. The boy is playing alone. Application of the *universal legislator* principle in such situations, Floridi argues, can lead to absurdities. "The boy may agree with Kant... as happens in every game: he is not the only one allowed to break the cars' windscreens in the dumping-ground and anyone else is welcome to take part in the game" (*ibid*).

According to Floridi eco or bio-centric approaches hold no water either since breaking the cars windscreens does little to modify the fact that it is in fact a dumping ground. At the risk of being accused of nitpicking we might forward an argument related to the overall mess to which the boy seems to

be contributing. For example broken glass, or perhaps accidently spilled oil due to the boy's vandal-like behaviour may have a greater long term impact on the immediate environment than if the boy desists in his behaviour.

Consequentialism holds that the *consequences* of action form the basis for moral judgement, thus a morally good action is one that produces a good outcome, or a good consequence. On the consequentialist reading we may be tempted to commend the boy and his amusing game since it is increasing the immediate level of happiness in the world. On the other side of the coin however the righteous adult or parent witnessing the behaviour may be inclined to feel less than happy to learn of her charge's petty amusement thereby negating any consequentialist claims. While Floridi dispenses with the consequentialist approach I believe we can get further mileage with a deeper analysis, particularly related to Floridi's ontocentric principles regarding entropy. I shall come back to this issue a little later.

Virtue Ethics, Floridi contends, appears to be a plausible traditional ethical model that approaches a workable solution. Virtue ethics considers the character of the agent involved (the boy in this case) rather than the consequences or the action itself. The boy's destructive fun would appear on the face of it, certainly to most god-fearing folk morally deprecable. The fact that we find such an act morally deprecable, albeit minimally, lends weight to argument that we might have some concern for the boy's character. Such vandalism, some would believe, could only have negative consequences for the boy's future development. Certainly most normative accounts of the boy's behaviour are likely to consider it minimally anti-social. Herein lies the fallacy Floridi argues, it is precisely because we find the behaviour deprecable that we infer that it will lead to negative consequences (*ibid*).

Finally we come to Information Ethics. IE, Floridi claims, provides an ethical argument as to why the boy's behaviour is "a case of blameworthy vandalism" (*ibid*). The boy, Floridi claims "is not respecting the objects for what they are, and his game is only increasing the level of entropy in the

dumping ground, pointlessly" (*ibid*). That is, under IE, even a dumped broken down old car deserves the right to flourish within the infosphere as an information object.

I do not find this argument entirely convincing, it is not immediately obvious that the boy's behaviour constitutes "blameworthy vandalism". IE just seems to be *telling* us it is wrong without really explaining why. It seems Floridi's claim approaches the level of question begging of which he accuses virtue ethics. Bynum following Wiener argues ever increasing entropy is "the great destroyer that eventually dismantles all patterns and structures" (in Van Den Hoven & Weckert 2008, p17). Thus entropy, operating in one temporal direction eventually dismantles all information objects. There are certainly other arguments that could be made to say why a boy ought not to smash the windows of a discarded car in a dump yard but these arguments seem to be related to the antisocial nature of the behaviour or the good character of the boy. The only thing I think we can blame the boy for is perhaps the acceleration of entropy.

While Floridi's dump yard allegory is easy to understand it is not entirely convincing. Certainly the boy's behaviour on the face of it and at some level exhibits a lack of care or consideration for his surroundings. Floridi argues, "he ought to stop destroying bits of the infosphere and show more respect for what is naturally different from himself and yet similar, as an information entity, to himself" (*ibid*). Yet there are other examples of 'accelerated entropy' that we could cite that I suspect could be claimed legitimate. Take for example a bulldozer driver crushing and totally destroying the car for it to be buried. Would this be considered 'blameworthy vandalism', if not, why not? Disregarding any arguments regarding the intentions of the boy vs. the bulldozer driver the end result is the same, both agents are contributing to information entropy. Surely under IE it is also the responsibility of the bulldozer driver to not destroy bits of the infosphere. It seems Floridi's argument at some level rests upon deep seated virtue or even Kantian based ethical undercurrent. We ought to

respect things as they are because this is the right thing to do. I think there is a way out of this problem which draws upon the ontocentric principles and I will return to this later in this paper.

Furthermore there is another argument that could be developed against the dump yard allegory. Floridi seems to confuse levels of abstraction with regard to entropy here. Certainly at an ELoA as well as a "thermodynamic LoA" to be sure the boy seems to be contributing to entropy (at least accelerating the process) within his surroundings. Entropy is unidirectional, the window can't be put back together again – and then Floridi argues, the boy is showing lack of consideration or care for his surroundings. However at an informational LoA in fact the boy is arguably increasing the informational complexity of his immediate surroundings. Each piece of broken glass, for example, surely from an informational LoA, is individually as much a valued as informational object in its own right as an intact windscreen is. I can only imagine appeals to anthropic levels of abstraction that would value an intact windscreen over the thousands of beautiful, glittering, individual, unique packets of information that represent the shattered windscreen. We could argue informational entropy in this case does not appear to be being degraded or destroyed; indeed we could argue the contrary is the case. It is certainly the case that entropy is increasing but the entropy in question is being caused at another LoA, one that views the degradation, decay and general increase in disorder of macroscopic objects as entropy, but this seems to be the kind of entropy Floridi argues that informational entropy is not.

To understand this further we need to take a look at conceptions of entropy in general, noting that Floridi argues that his use of the term differs considerably from the typical physicist's conception forming a part of the second law of thermodynamics. The physicalist thermodynamic conception of entropy states that within any process in an isolated system, in relation to the physical systems, potential for work entropy always increases. Entropy runs in one direction and is generally irreversible – we cannot put back

together the shattered windscreen. It is the general property relating to heat/energy equilibrium. Thus within the universe, hot things are cooling down, losing energy and cooler things are warming up gaining energy all tending towards a homogenization or equilibrium state. An easy to understand and often used example involves the melting of an ice cube in a glass of water in a warm room. If we consider the glass containing the ice located in the room as a thermodynamic system, as the ice melts the system there is a heat exchange and entropy increases. The ice warms up and melts, the water cools down. By analogy, objects wear out, organisms die, gasses tend to evenly distribute.

There are biological construals of entropy related to thermodynamic entropy. Schrodinger's early biological discussions involving entropy are interesting. In his 1944 text *What is Life?* he suggests that organisms avoid decay and disorder by consuming, eating, drinking and breathing, or assimilating in the case of plants. Lehninger, a biochemist, following Schrodinger, observed the apparent order in cellular organisms and argued that as cells grow and divide the order is compensated for by the corresponding "disorder" they generate within their environment, thus "organisms preserve their internal order by taking from their surroundings free energy, in the form of nutrients or sunlight, and returning to their surroundings an equal amount of energy as heat and entropy." (1993). Clearly the imposition of order is at the corresponding cost of disorder somewhere else. Recently Kaila and Annila (2008) describe evolution in terms of an exchange of energy or a dispersion of energy, and a corresponding increase in entropy towards a stationary state, that is, energy tends to flow toward a state of equilibrium, in the absence of a high freely available energy source. Natural selection, they go on to argue, tends to favour adaptations that lead to faster entropy increases, equivalent to an increased depletion in freely available energy, within the ecosystem.

It is important to note here that clear connections are evident between construals of entropy analysed at a materialist or physicalist LoA with those

at a biological and chemical LoA. There is a general movement towards states of equilibrium at the cost of an increase in entropy elsewhere in the system. Entropy is construed as a part of the natural order of things. From a biological perspective entropy is the trade off for life and order. Organisms that flourish do so at the expense of faster and higher increases in entropy in the surrounding environment as they consume available energy. Floridi however wants to paint informational entropy in a slightly different light. While accepting Wiener's thesis regarding the nature of information, that "information is information, not matter or energy." (1961) it seems reasonable that we should expect to find connections between materialist, chemical or scientific construals of entropy and informational entropy.

Floridi's IE does admit an inverted relation between informational complexity and entropy with regard to exchange of semantic value but there is no acknowledgement of an entropy cost associated with increasing semantic complexity in information objects elsewhere in the infosphere. Floridi explains it like this, "as the infosphere becomes increasingly meaningful and rich in content, the amount of information increases and the amount of entropy decreases, or as entities wear out, the amount of entropy increases and the amount of information decreases." (Floridi, 1999, p44). In a later article Floridi explains entropy as any form of destruction or corruption of informational objects, or "any form of impoverishment of being" (Floridi, 2010, p84). Since every extant object or state of affairs[47] qualifies as an informational object any destruction or corruption of

[47] David Armstrong has developed a comprehensive metaphysical theory of the world consisting of "states of affairs". Armstrong's view is that a state of affairs is the instantiation of a particular property or set of properties in objects. This view is not altogether different to Floridi's informational realism in that the Aristotelian "substance", or Floridi's entity or object is to a large extent unknowable, instead what we can talk about is the information, or what Armstrong would call attributes or universals. In chapter 3 of this work I argue that "states of affairs" constitute finite discrete informational entities in their own right. By extending informational realism in such a way we also accommodate causal processes interpreted at an Informational LoA as legitimate informational entities and thus exposed them to the same ethical scrutiny under IE.

anything at all contributes to entropy and supposedly a corresponding decrease in information.

This seems odd to me. It seems Floridi does not want to attribute any cost associated with increased levels of semantic complexity within the infosphere. We could call this the free lunch argument. All forms of structure and order come at some cost. This applies equally to semantic informational structures. For example, the information conveyed by the rings which form on a tree trunk with the passing of each season come at the expense of that tree taking nutrients from the soil, water, air (carbon dioxide, hydrogen and oxygen) and using sunlight to convert this material into organic material, this process directly contributes to the information conveyed by tree rings.

If we construe structured meaningful information as the opposite to entropy, correspondingly this information is representative of some form of physical order at an alternative LoA, this order (at the physical LoA) is at the expense of some increase in physical entropy somewhere else within the system (the infosphere) and correspondingly at the informational LoA some informational entropy has occurred to bring into being a complex informational object. Just as evolution is a movement toward increased complexity in form, correspondingly at the informational LoA an increase in informational complexity or form is necessarily at the expense of an increase in entropy due to energy consumption somewhere else within the environment. This energy consumption can be represented in informational terms by the degradation or consumption (entropy) of the informational objects that are representative of nutrients, chemicals and so on required to sustain the order of the information object in question.

7.6 Informational Entropy – Ver 2.0

I propose a subtle change in the way we ought to interpret entropy within IE. Analogous to physical (thermodynamic) entropy which is the tendency

toward heat/energy equilibrium, I believe information is conserved in a similar way. On this account informational entropy, similar to thermodynamic entropy is a uni-directional process from organised, structured complex informational objects and states of affairs to increasingly disordered or unstructured states of informationality but equally without loss, that is, information is conserved but transformed. Like thermodynamic entropy, informational entropy is in general a non-preventable process whereby the overall state of informational order in the infosphere is constantly moving towards less ordered states. The destruction of an object does not entail the complete loss or destruction of information, rather a reorganisation of that information. Information generally becomes more distributed, and in some cases more complex[48] but rather than describing the process as a loss of information I think it is better described as information being transformed from one state to another. Informational entropy is thus a form of non-loss decomposition. As an information object is degraded, decomposes, is destroyed or corrupted information is transformed and moves towards less structured, more homogenous and austere informational states.

Consider Floridi's junk yard metaphor. As the boy smashes car windscreens, the totality of information that represents the information object, the car window in Floridi's example, is not lost, however there is a decomposition process, the previously intact windscreen is now represented by thousands of individual shattered discrete fragments of glass, each shattered part exists spatially in relation to the other parts and its surrounding environment and as such contains its own self-contained set of information related to its individual existence. No information is lost, however some degree of entropy has occurred in the sense that the original object has been destroyed but this application of entropy seems to be in the

[48] However any increase in complexity will almost certainly come at a cost. That cost will be an increase in entropy somewhere else within the system.

more traditional physicalist sense of the term. An original object has transformed from an intact state to a state of increased disorder. The process is linear and not reversible. We cannot feasibly put the windscreen back together again. Additionally it seems that only on an anthropomorphic reading could we argue that the intact windscreen is more "meaningful" than the shattered one.

Let's examine in more detail the idea of what Floridi might mean for the infosphere as a whole to become increasingly meaningful and rich in content. The infosphere includes individual informational objects, finite sets of related entities, or in Armstrong's (1998) terminology states of affairs, indeed all extant objects interpreted at an informational LoA. A fundamental problem is this: Virtually all human activity increases the level of entropy in the infosphere at some level of abstraction. This is supported by Kaila and Annila (2008) above, "The principle of increasing entropy, equivalent to the decreasing free energy, is pure and austere but its mechanical manifestations can be complex and intricate." (*ibid*). There is no free lunch. Our existence and more so our flourishing necessary entails an increase in entropy, indeed the more successful we are, the more semantically rich and complex the infosphere becomes, the faster the rate by which entropy increases. As previously discussed this entropy has its informational counterpart, the informational objects being degraded and consumed are the ones that directly support our internal order. Information isn't lost completely, it is transformed into various other forms of information, that which make up our existence, alternatively the complex tools and artifacts we build and require for our existence, this includes the (entropic) waste products that are the result of such degradations and consumption. Furthermore, increasing entropy at any level of abstraction, let's say at concrete materialist, physical or biological LoA represents a corresponding increase in entropy at the informational LoA. I believe this is supported by Kaila and Annila's work; we can draw parallels between increases in informational complexity and evolution. Natural selection

favours adaptations which lead to faster entropy increases and with such adaptations come higher degrees of complexity. Such biological complexity has analogous degrees of informational complexity supporting the correlation between entropy discussed at physical LoA and informational LoA.

We can equally say that increased states of societal order and complexity resulting in increased informational complexity and semantic content require increased energy input – resulting necessarily in increased levels of entropy. Very simple societies utilise low levels of energy and are correspondingly represented with low degrees of individual specialisation and general complexity. (Tainter, 1990). Entropy seen at materialist LoA also has its analogy at an informational LoA. I consider informational entropy as not being the loss of information but instead the transformation of information from structured ordered states to increasingly disordered (but not random) homogenous states tending to what I will call informational equilibrium. The complete end-state of informational equilibrium is simply analogous to the physical thermodynamic equilibrium but interpreted at the informational LoA. If this is not the case then the onus is on the claimant to explain how entropy could be decreasing at one level of abstraction (the informational level) yet increasing at another (physical, chemical or biochemical levels).

Kaila and Annila offer further support for this thesis. "Thermodynamics sheds light on the origin of (genetic) information as a powerful mechanism to increase energy transduction." (2008, p206). This biochemical example clearly points to complex informational states of affairs, in this case evolving genetic information as embodying highly effective entropy increasing mechanisms. Successful organisms are the ones that are able to adapt to and exploit their environment rapidly using available energy thereby accelerating entropy. Under Floridi's ontocentric principles, it seems we are doomed; as highly successful entropy generating organisms we ought to remove ourselves *qua* information objects from the infosphere.

The corollary is this: increases in the semantic richness of content, informational complexity and meaning within the infosphere with regard to objects or states of affairs *qua* information objects at an informational LoA necessarily results in an entropy cost elsewhere. That entropy cost can be seen at the informational LoA but has its analogous interpretation at materialist LoA. All costs can ultimately be sheeted back to an energy conversion cost, a resulting energy loss via dissipated heat, necessarily resulting in an entropy increase toward a state of equilibrium. This has its informational LoA equivalent; all informational costs can be ultimately sheeted back to informational conversion costs.

Let's consider another example. In considering the informational complexity of modern, cohesive, organised, civil society we observe vast intellectual, artistic, physical, cultural, industrial and societal diversity. From an informational perspective such diversity represents significant semantic complexity. Such complexity exhibits close correlations with energy input. Tainter (1990) argues complex socio-political systems require energy for maintenance. Such societies are informationally complex. Access to cheap and abundant energy, largely from fossil fuels has enabled the development of these levels of complexity. Yet this has come at a significant cost. Essentially modern civilisation can be viewed as an entropy machine converting raw materials and resources into waste and heat, no more evident than in the accepted and recognised phenomena of anthropomorphic global warming. Throughout the last couple of hundred years of global industrialism and modernity we have greatly increased semantic complexity within society by turning stored energy into carbon. This equation both from a thermodynamic or energy-theoretic LoA as well as informational LoA sees net entropy increasing. Whether this is manifest in the chopping down of a tree to burn as firewood, or the mining of hydrocarbons to burn as fuel or to be used to manufacture other products, there is necessarily a net energy loss along with the corresponding increase in entropy. Thus while we have increased the complexity of society this has

come at the cost of an overall increase in entropy. While this view is generally speaking from a materialist LoA that admits the laws of thermodynamics, heat exchange and physical entropy into its ontology there is an equally valid analogous view at the informational LoA. Entropy at the thermodynamic LoA must find its corresponding entropy at the informational LoA. The "informational object" that is a vast oil well for example is being degraded in exchange for our comfortable way of life. Following Lehninger above, we (increasingly) take from the environment to preserve order, returning energy, as heat, and entropy and this process is equally analyzable at physical, chemical and informational levels of abstraction.

The previous discussion raises several new issues for IE. I now summarise my main arguments and resulting problems before suggesting methods by which they may be addressed:

- There are more or less explicit correspondences across varying LoA.
- Entropy analysed at a materialist or accepted scientific LoA has its analogy with entropy analysed at an informational LoA.
- Where more or less explicit mappings between LoA are observable, entropy cannot be decreasing at one level and increasing at another.
- Informational entropy is the non-loss decomposition of information towards distributed informational equilibrium.
- Decisions about which informational objects have minimal and overrideable intrinsic worth are generally made at an everyday LoA based upon an everyday judgement about the value or worth of that entity.
- All increases in informational complexity, whether such increases are the result of anthropomorphic activity or biochemical evolutionary processes, or other, necessarily result in an increase in entropy elsewhere in the system. This entropy is equally analyzable at an informational LoA.

Of course this alternative view of informational entropy creates a significant dilemma for IE, in particular IE's ontocentric principles. How are we to decide upon which entropy increasing activities are more (or less) morally blameworthy? The ontocentric principles when applied to the maximisation of our own good don't seem to help us in any significant way since we are almost certain to favour decisions that promote our own well being and flourishing. Thus while, on the ontocentric principles, entropy ought not to be caused in the infosphere, our mere existence, the increasingly complex world we live in, represented by increasing levels of informational order, not to mention continued evolutionary processes presumably still at work, result in corresponding entropic activity somewhere else in the system (infosphere). In consideration of Lehninger (1993), and Kaila and Annila's (2008) arguments, application of IE's ontocentric principles would require us to at least stop all entropy increasing activities yet this is materially impossible, our mere existence is an entropy generating activity. Any increase in semantic complexity, or to use Lehninger's terminology, the generation or maintenance of order that results in a local decreased level of entropy is balanced by an increase in entropy elsewhere within the environment.

At an ELoA it seems considered decisions about which entropy causing activities we can live with would typically be made either with appeal to consequentialist-like eco or biocentric ethical principles. Indeed decisions regarding resource management, ideas about carbon trading, concerns about greenhouse gas emission, other environmental concerns and the like seem to appeal to such ethical principles noting of course that they are balanced against potential gains or general utility benefiting the overall good. However these non-anthropocentric examples amongst no doubt many others do appear to be good candidates for ethical consideration with appeal to IE's ontocentric principles since most of us can agree that eco-systems or

the environment in general certainly qualify as being intrinsically valuable. This approach however doesn't seem to work for all inanimate objects as Brey (2008) points out. One response is that ethical decisions regarding such issues ultimately collapse into some form of consequentialist approach based in bio or ecocentric principles. Indeed this is in part Brey's 2008 argument where he proposes the moral bar with regard to inanimate objects ought to be lowered to the level where such entities are deserving of moral respect rather than being considered intrinsically valuable. While I think Brey's argument has merit I think in order to strengthen IE we need to find a solution which draws directly upon the ontocentric principles and remains true to IE's foundational objectives. I begin to explore this approach below.

7.7 Solving the Entropy Problem in IE

In order to solve this problem, I believe we need to develop an *informational* story about how moral decisions are to be made. It may be no surprise to some that what we end up with looks something like rule-consequentialism yet I would argue that it differs in one important sense. Rule-consequentialism's ultimate ideal is on maximisation of good based upon following an accepted set of rules. In contrast the ideal for IE ought to be informationally based, that is we ought to be able to give an info-theoretic justification for our ethical judgements. At an informational LoA this is possible if we draw upon the ontocentric principles outlined in IE which focus upon the minimisation of informational entropy. Thus moral decisions will be based upon the careful analysis of IE's ontocentric principles taking into consideration the complete set of entropy causing activities that relate to or may be the result of patient oriented moral decisions. This extends the requirement for informational analysis beyond localised increases in semantic complexity which on the face of it do not directly appear to be entropy causing. Let us see how this could work.

We must first recognise that any increase in entropy at a physical/materialist LoA has its counterpart at the informational LoA. I

don't see this as being problematic for IE in anyway, after all as Floridi has pointed out H_2O is conceived of as a chemical object at the chemical LoA, and has as its counterpart "water" at an ELoA (Floridi, 2012, p269). They are essentially the same thing just considered at different LoA. Acceptance of this is critical in that it acknowledges that increases in semantic complexity necessarily come at some (entropic) cost. Further it entails more or less explicit relationships or mappings (or correspondences) between differing levels of abstraction. I think we have seen enough evidence to suggest that this is indeed the case. Certainly to deny it I suggest means that we must accept a version of "entropy" which bears no relation to the general meaning of the concept but more worryingly, it potentially obfuscates any discussion that makes use of the term. We also must agree that almost all activity that results in an increase in complexity including the maintenance of our internal order as living organisms, returns a corresponding amount of energy plus entropy and so any movement to prevent, let alone utterly eliminate entropy is untenable. Indeed Floridi acknowledges that any goal to eliminate the kind of entropy described by the theory of thermodynamics wouldn't simply be controversial but idiotic.

Yet if we claim that informational entropy is vastly different to a materialist reading of entropy we end up with a paradox that requires considerable explanation or worse the *ad hoc* application of IE's principles on a case by case basis based upon our judgement calls about the intrinsic worth (minimal and perhaps overidable) or whether or not we ought to offer moral respect to all manner of information objects and these decisions will almost certainly have to be made at an ELoA. The problem for this approach as previously discussed is that it appears to force IE into making consequentialist based claims about which set of circumstances we can live with and which set we are not prepared to tolerate. Ultimately such judgements as we have seen typically appeal to principles not directly related to IE's ontocentric principles but instead to other norms or prearranged rules depending upon the context of the decision. Consequently

this approach does not appear to solve the foundational problem which Floridi claims IE is intended to address. The corollary is that we must accept as incoherent any claims implied or explicit that 'entropy may be decreasing at an informational LoA while it is increasing at a physical LoA within the same system'. *Prima facie* any increase in semantic complexity will incur an increase in entropy somewhere.

Biological explanations of entropy as we have seen argue that increases in complexity are the result of energy use in exchange for increases in entropy and heat exchange. At an informational LoA I see this as being no different – all increases in informational complexity are the result of energy use (interpreted from an informational LoA). Seeing informational entropy in this light brings it into line with other accepted explanations and also helps establish clear mappings or correspondences between differing LoA. Adopting this approach I believe strengthens the coherency of IE and the method of LoA. It does however create a new problem. It acknowledges that the there will now be a great deal of activity that violates Floridi's four universal laws or ontocentric principles, that is, that information entropy ought not to be caused, be prevented and or removed. We must acknowledge that there is a direct relationship between increasing levels of entropy as semantic complexity increases, whether that is manifest in the evolution, development and order representative of an organism, or in the creation or maintenance of non-sentient being, right through to the informational objects or complex informational states of affairs that are representative of societal complexity, maintenance of our own individual existence and other constituent arbitrarily complex concepts and systems. Thus based upon IE's ontocentric principles how are we now to decide between which entropy causing processes are acceptable and those which are not?

Floridi, in response to Brey's 2008 critique of IE offers us some direction regarding this problem. Recall that IE maintains that "non-sentient entities have some minimal, easily overideable but still intrinsic value" (Floridi,

2008). The burden of proof Floridi suggests is to ask "from a patient-oriented perspective whether there is anything in the universe that is intrinsically and positively worthless ethically and hence rightly disrespectable" (*ibid*). In reference to IE's ontocentric principles we might argue that any informational object or constituent state of affairs that causes or contributes to entropy qualifies as an entity whose intrinsic value is minimal and overridable. Such entities or states of affairs as per the ontocentric principles ought to at least be minimised, if not removed from the infosphere. Although Floridi does not say this explicitly he does say we ought to treat something as intrinsically valuable and hence worthy of moral respect by default, until proven guilty (*ibid*). By acknowledging that there are analogous relationships across varied LoA, particularly with regard to informational entropy, we allow for judicious reasoning at an ELoA regarding the degree of intrinsic value at an onto-informational LoA.

This approach I believe avoids the mistake Floridi accuses Brey of making (2008). While IE begins with the view that existence in all its manifestations including non-sentience is worthy of moral respect, Brey's issue is with what he argues is the "untenable egalitarianism in the valuation of information objects". Brey struggles with the issue of differentiation between objects that he considers might have additional worth beyond their status as information objects (*ibid*). This issue is not as Floridi points out "some daft idea about the intrinsic value of Shakespeare vs. Dan Brown, or chocolate vs. excrement" (2008) but (and although Floridi does not explicitly state this) I would extend the argument, in the objective application of the concept of entropy and IE's universal laws to individual cases. Thus, rather than assigning intrinsic value or moral respect upon objects from sociologically normative perspectives we remain true to IE's ontocentric principles. We ask the question, which are the entities or states of affairs that are entropy causing, or which objects are the result or manifestation of entropic processes? A vat of toxic waste is clearly an example of a by product caused by some entropy causing activity, thus we

ought to give careful consideration to such activities that result in such materials. Correspondingly it stands to reason that such material has a minimal and overidable value. Careful analytical understanding of these questions directs us towards the correct judgement calls within IE.

Floridi might well argue "but that's just what I said". Not quite. Floridi makes it quite clear his concept of entropy is quite separate from materialist conceptions of the term. Semantic complexity is highly valued by Floridi's IE but this value does not take into account the prior entropy equation that allowed for the development of such complexity. The change I propose in the meaning of "informational entropy" within IE brings it into line with other conceptions of entropy thereby theoretically strengthening IE's foundational basis. Floridi simply does not admit to explicit mappings between LoA when it comes to the concept of entropy, yet he seems to have no problem with analogous LoA comparisons with every other entity, water *is* H_20 at the chemical LoA he argues (2012, p269). I want to argue that by viewing entropy at the informational LoA, and acknowledging that all informational complexity incurs an informational entropy cost elsewhere in the infosphere we have the right framework to make judgement calls about which informational objects or entropy causing situations are minimal and overrideable with regard to moral worth. Ethical judgements are always to varying extents tradeoffs. What can we morally live with, and what is morally reprehensible. IE is no different.

We acknowledge that attempts directed towards the wholesale elimination of entropy from all informational objects and systems are unreasonable and untenable. Anthropomorphic and biocentric ethical systems are not immune from such criticisms either; it is simply not possible to eliminate all forms of suffering or pain. The aim of such approaches is to minimise or remove instances of so called "evil" where possible. What the ontocentric principles direct us towards is a more inclusive form of respect for all existing entities, however as Floridi points out (2002, p299) "it would be unreasonable to assume that they all qualify for exactly the same kind of

respect" and "respect information objects for their own sake, if you can" (*ibid*). Where I differ from Floridi is in the application of IE's principles. Floridi seems to revert to an ELoA when making judgement calls about the apparent worth of an informational object. A more objective approach, I believe, is to remain true to the principles themselves and thereby avoid some of the claims that IE collapses into a form of consequentialism. Thus, activities that consist or result in high levels of entropy generation are rightly seen within IE as deserving minimisation or elimination. The consequentialist commitment to maximisation of the good is not necessarily guaranteed. While we acknowledge our own existence as intelligent sentient beings carries a high degree of intrinsic worth and the maintenance of our wellbeing and general existence entails entropy increasing activity, judgement calls between activities and behaviours based upon entropy generation are still able to be made.

This view is highly compatible with eco and biocentric ethical theories but extends to the inclusion of all including non-sentient or living entities as well as being inclusive of systems and states of affairs. In the new version (IE), the argument defends the intrinsic value and moral respectability of systems, non-sentient entities as well as individuals *qua* informational objects. (Floridi, 2002, p300). In consideration of "systems" or the infosphere as a whole, we are now able to balance certain activities or behaviours with levels of entropy generation that they may cause. Our approach is one of stewardship towards the infosphere so for example IE would counsel us against building a coal fired power generation facility in favour of the use of solar or wind power, but only where the entropy equation adds up. IE counsels the boy against smashing the car's windows in the dumping ground due to the resulting unnecessarily caused entropy. This application of IE's principles should not be confused with the laws of thermodynamics although clearly there is a correspondence of sorts across the relevant LoA.

Informational entropy is increased when a coal fired power station burns coal in exchange for heat and electricity. This is because the informational object that is coal as well as the originating coal mine is being degraded resulting in highly dispersed and distributed information. At another level of abstraction, we talk of this as being the creation of carbon dioxide and the generation of hundreds of millions of tonnes of other waste products such as fly ash, toxic gasses, desulphurisation sludge and other toxic entropic by-products. Thus while we acknowledge that generated electricity results in a local increase in the semantic complexity of centres of civilisation, and the maintenance of order for its consumers we also accept that this is in exchange for entropy created elsewhere in the system. This resultant entropy is also analysed at an info-theoretic level of abstraction.

Informational entropy can be illustrated at a digital level. For example, consider an electronic customer profile, while the representation of a customer's identity digitally is certainly worthy of some moral respect as an information object in its own right – "a disinterested, appreciative and careful attention" (Hepburn, 1984, *in* Floridi 2002). A naive reading suggests simply that the corruption or destruction of such objects represents informational entropy and thus should be prevented. Manders-Huit extends this argument by adding, "The question is still very much whether the data subject's experience of himself can be represented adequately in terms of the dynamics of moral persons. A person understands himself as someone who could improve morally, even if he makes no attempts in this direction" (2011, p77). Clearly the electronic profile is relative to the individual it represents, albeit somewhat fixed and certainly bounded. Thus while no information is actually being "lost", informational entropy may be occurring in the sense that the information in being degraded relative to changes in the actual person it seeks to represent. While electronic profiles of course never have the same level of moral value as the individuals they represent, IE's ontocentric principles still apply. We see this in action whereby information technology specialists employ various tools and techniques to protect data

from corruption and loss. In the case where such profiles represent non-existent individuals the value placed upon such information objects is clearly minimal and overrideable.

7.8 Conclusion

This paper has considered in detail Floridi's use of the concept of Entropy within Information Ethics. An argument has been forwarded that Floridi's use of the concept of entropy is problematic in that it appears to allow for situations where entropy can be increasing at a physical LoA while informational entropy may be decreasing at the informational LoA. The corollary to this is that generally speaking any increase in semantic complexity is at the cost of an informational entropy increase somewhere else within the system.

I have argued that this approach results in difficult to resolve paradoxes for IE. The solution I propose is that we reconcile our discussions about entropy to bring them into line with more conventional approaches to the concept and I presented several lines of evidence to support this argument.

With regard to informational entropy and the ontocentric principles, I argued that we need to consider the full informational story. Ethical decisions must be considered in the context of the complete set of entropy causing activities that relate to or may be the result of patient oriented moral decisions.

If we do not take this approach not only does the IE's use of the concept entropy appear to be at odds with orthodox uses of the term but our moral decisions made with appeals to IE seem to collapse into varying forms of consequentialist arguments about the sets of consequences we can tolerate and those we cannot. This is because almost all human activity results in entropy, this is particulary the case in situations where semantic complexity is being increased. While we can tell an informational story about the increase in semantic complexity associated with civilised cultural society,

we can equally talk of the informational entropy associated with the burning of fossil fuels for example in order to sustain that society.

Adopting a truly ontocentric and informational approach allows IE to be even more objective in reasoning about all kinds of moral problems. Furthermore the proposed approach largely avoids appeals on anthropomorphic grounds or, as I have argued, the reliance upon decisions made about which behaviours or actions are morally blameworthy made at an everyday level of abstraction. Instead, the proposed approach appeals directly upon the ontocentric principles to guide our moral compass.

This approach also assists reasoning about which informational objects indeed have minimal and overridable moral value by considering their entropic contribution. Entities that either contribute, or are the resulting by product of entropic processes, generally speaking, accrue little if any moral value. As such we can comfortably assume vats of toxic waste, overflowing ashtrays and the like indeed have minimal (if any) value. In the same sense entities or states of affairs that are highly entropic by nature such as coal fired power stations, the destruction or overuse of resources and energy sources violate the null law (entropy ought not to be caused) and as such would generally not be considered the best moral action.

The approach forward I believe avoids some of the earlier criticisms directed towards IE and rightly applies an informational or ontological framework towards moral decision making. This approach also helps eliminate confusions regarding at which level of abstraction moral decisions are made. By adhering to the ontocentric principles we take a more objective look at moral dilemmas. This approach does not in any way attempt to do away with more orthodox ethical frameworks, IE never had that intention anyway. The idea is to extend the moral base to include all informational objects including non-sentient organisms and entities right through complex ecosystems and societies. My hope with this approach is that we can eliminate the absurdity associated with arguments about attributing intrinsic value to various detritus, waste matter and the like and

consider all entities within a framework based upon the application of IE's ontocentric principles.

8 Summary

Having studied philosophy for many years now I am often asked the question – what actually is philosophy? There are no doubt many possible ways of answering this question however I tend to reply with a relatively canned answer, one that most inquisitors seem generally satisfied with. Philosophy is often about investigating and attempting to explain conceptually hard problems for which currently there appear to be no scientific or mathematical answers. To philosophers this is usually considered metaphysics. I declared at the beginning that I come at these problems from a naturalistic angle, that is, I consider it just as valid that we draw upon our experience of the world as well as the explanations science provides us in helping us decide upon metaphysical puzzles.

However, while this particular work assumes a form of Quinean naturalism as its starting point, I am certainly not arguing that this excludes non-naturalistic avenues of exploration as having validity. For example, I can imagine a fruitful investigation into the ontological status of information from a phenomenological perspective. In his 2004 paper, *Open Problems in the Philosophy of Information*, Floridi suggests, following Searle (1980), that the data-grounding problem, that is, how data can aquire their meaning, is essentially the problem of intrinisic meaning or intentionality. A phenomenological approach to this question may yield better results than the naturalistic angle.

Regardless of our biases most of us go about this endeavour with some trepidation, care and of course with the usual application of logic and reason. Our goal: to shed light upon, clarify, interpret or make some sort of sense out of problems that often have many different possible explanations, depending upon who you ask since it is also true that our philosophical arguments often underdetermine what sometimes turns out to be the case. As I have said my biases lie predominantly with those philosophers that hold science as the yardstick by which knowledge is measured; we are often

labelled naturalists or empiricists. Science we claim is the best measure of what there is, and how we know what there is. I don't intend to offer up any argument for this assumption other than to say many philosophers, scientists and others also believe this to be the case. Perhaps not so coincidently, it turns out that two horns in this claim are central to philosophy; ontology is captured by the question 'what is there?' and epistemology 'how do we know?' It should come as no surprise then that the questions and problems addressed within this thesis are located under these philosophical branches. While these fundamental questions haven't changed in millennia, the ways of considering them do from time to time and this research, I hope, attempts to offer some answers from a new perspective. Nevertheless genuine philosophical problems are always open. We must expect differences of opinion, vigorous debate and honest disagreement. In philosophy we stand squarely faced against a double edged sword. Whilst we expect philosophical arguments to display clarity, rigor and unimpeachable reasoning it is often the case that the stronger the argument or the bolder a theory is the more easily criticised it becomes. This research focuses on a fascinating emerging field in philosophy, the philosophy of information. Philosophy of information is concerned on the one hand with the nature of information but it is also the case that information and computing are providing new ways of understanding many of the more traditional philosophical problems; this thesis deals directly with some of these issues. As Bynum and Moor (1998) suggest, the emergence of the digital or information age is providing us "with a new philosophical prism through which to view philosophical problems" (*ibid*, p1).

In 2004 Luciano Floridi published what I consider to be a watershed article on the philosophy of information. Entitled *Open Problems in the Philosophy of Information* the article introduced 18 distinct open problems concerning the nature of information. It was Floridi's paper that provided the impetus for this body of research. The overarching goal of this work is to contribute toward and clarify the ontological status of information and

consequently its role within Floridi's ontocentric information ethics. While chapter two serves as a primer to the philosophy of information and information ethics what follows is a collection of separately published papers. Two chapters were published as chapters in two separate books and the other three chapters were presented at various conferences and later developed into full papers. As such the work should not be strictly regarded as logically consecutive but rather as being all related to the general research aim of this thesis which is to contribute to the wider philosophical and scholarly understanding summed up best by Floridi's Problem 15 (2004a, p572) *"Wiener's problem: What is the ontological status of information?"* Drawing upon a quote from Barwise and Seligman, Floridi sets up the issue as follows, "If the world were a completely chaotic, unpredictable affair, there would be no information to process. Still, the place of information in the natural world of biological and physical systems is far from clear." (1997, xi). A particularly relevant question concerns the matter of whether information can be characterised as an independent ontological category. Floridi situates this problem with particular reference to Norbert Wiener's suggestion that "information is information, not matter, or energy. No materialism which does not admit this can survive at the present day." (1961, p132).

The *rub* so to speak, or the primary issue which I explore in considerable detail and from various angles (epistemic, causal, explanatory, ontological, object oriented programming perspectives and finally from an ethical perspective) throughout this entire thesis is this: if indeed information does qualify as an independent ontological category then how is it related to the physical or material as well as the mental worlds. This thesis draws upon prior relevant philosophical work much of which to my knowledge has not been utilised in relation to current ideas about the philosophy of information, additionally I try to draw upon empirical cases in deference to Barwise and Seligman's implication but also in line with my declared biases – I believe

in striving for alignment with scientific discourse and theory only serves to strengthen our metaphysical arguments.

Floridi's *informational realism* (2004b, 2008) is of great interest to me. Floridi's ontology presumes the universe to consist of informational objects or entities all dynamically interacting with each other. Floridi draws upon ideas used primarily in computer object oriented programming (OOP) theory to illustrate his theory. Chapter 3 of this thesis (Published in Demir's edited text, *Luciano Floridi's Philosophy of Technology: Critical Reflections*) explores the conceptualisation of objects in the world as information objects. The primary point I make is that attempting to use the ideas in OOP to underwrite a metaphysical theory is a mistake. I do have sympathy for Floridi's metaphysics, that is, just as we might explain and interpret objects, states of affairs, and all manner of simple through complex entities from a physical or materialist level of abstraction, alternatively from biological or chemical ones, or in the case of a set of traffic lights (an example Floridi often uses) an electrical level of abstraction, we might also validly interpret entities as informational objects. However, such an interpretation bears little resemblance to OOP objects which are in their purist sense computational *referents*. That is, they both reference the 'real world' objects they purport to model and in turn they *refer* or 'point' to arbitrarily complex representations on computer disks for which they are indebted for their existence. OOP objects only ever exist at a representational and logical level whereas Floridian informational objects are meant to literally 'be' actual physical objects in the world. To say that real world objects are indeed informational objects then to somehow equate them to OOP objects raises obvious second-order-object problems. Unless by some bizarre logic an informational object is somehow self referential, but such an approach seems to make little sense.

In light of my chapter Floridi in his reply seems to change tack somewhat:

> If the link between informational objects and objects in OOP does not help the reader, the invitation is to drop the comparison immediately. It is only a tool to help make sense of a philosophical thesis, if it is unhelpful you would be better to avoid it. If the way in which I discussed the relation between informational objects and objects in OOP confused other people apart from McKinlay, I am grateful to him for clarifying it. (Floridi in Demir, H., (ed.) 2012, p267).

While I welcome this, I do find it a little puzzling since Floridi in his previous work devotes a considerable amount of time in developing a picture of the information object drawing heavily upon OOP methods and theory. It certainly does seem that he wants to say information objects are just a little bit like Object Oriented Programming objects and my chapter explains why this cannot be the case. The general gist of the chapter argues that this is a fundamental mistake and it seems to some extent Floridi is happy to concede the point. Throughout the remainder of his reply Floridi tries to minimise the link between his informational objects and OOP objects and begins referring to his informational objects as *i-objects* stating that they are structural in nature with reference to ideas discussed in the philosophical metaphysics of structural realism. The structuralist approach is indeed employed by Floridi (see Floridi 2008) and I consider this in more detail in Chapter 4 comparing Floridi's metaphysics to Armstrong's (1989, 1998) equally intriguing 'state of affairs' based metaphysics. I welcome the subtle change of tack by Floridi and others and suggest all further work on informational realism by any philosopher interested in the topic drop references to OOP. OOP as a discipline lacks rigour, is fraught with ambiguity and is absent of any real coherent consensus – the chapter outlines these claims in some detail. Thus any informationally based metaphysical theory will fall considerably short of the kind of philosophical rigour required if it relies upon OOP for explanation.

A final comment I would make considering this chapter and Floridi's instructive response is perhaps just a clarification. It seems that Floridi accuses me (and indeed Quine) of making some kind of mistake as he explains in this passage. "It baffles me why Quine (and perhaps McKinlay as well, but I might be wrong) never acknowledged the constructionist nature of our epistemic interactions with the world. Perhaps some forms of empirical addiction can seriously damage your critical thinking" (2012, p270). I think this in some sense illustrates a point of departure between Floridi, and Quine and I. I open the chapter with a quote from Wittgenstein, "The world is the totality of facts, not of things." (1922, [1.1]). I am primarily arguing that all we can possibly know about the world is inextricably linked to the apparatus that we are uniquely equipped with to sense and interpret our surroundings, this is by and large an empirical experience and as such our explanations are by and large constructions. Following Quine, we understand what a 'Rabbit' is by ostension, not by logical reasoning (1960, p45). When it comes to formulating ideas about unobservable or metaphysical entities I argue that we are primarily informed by our empirical experiences. As such an informational interpretation of the universe, of everything from our observations of excitable cellular interaction, quarks and neutron stars, to our experience of every day kitchen items is a uniquely human method of interpretation. I don't see this as being at odds with Quine's endorsement of empiricism at all. "Whatever evidence there is for science is sensory evidence... and all inculcation of meanings of words must rest ultimately on sensory evidence" (Quine, 1969, p75). I cannot imagine Floridi's wanting to argue that we drop an empiricist thesis of science in favour of some form of *a priori* first philosophy, and I'd certainly have greater hopes for an informational realism than it being labelled merely some kind of Cartesian metaphysic based purely upon clear and distinct ideas (although I'm sure Floridi didn't mean to go this far in his reply). Whilst not explicitly stated, an underlying goal of the chapter is to attempt to find a more empirical basis for informational realism than the one

based upon ideas drawn from object oriented programming theory for my hope was that in doing so we might strengthen the theory.

Having considered in some detail Floridi's fundamental metaphysical entity, the *information object* in Chapter 3, it is time to situate this entity within its metaphysical framework and Chapter 4 addresses this particular issue. We must note that informational realism has a serious metaphysical competitor, that is, David Armstrong's 'States of Affairs' and Chapter 4 offers a critical comparison of these fascinating ontological theories. While there are certainly some similarities and parallels between the theories I doubt Floridi would disagree that Armstrong's metaphysics is considerably more developed than informational realism. This is understandable, informational realism is the new kid on the block and Floridi himself agrees needs further development. Contrastingly Armstrong has spent his life developing his metaphysics – this only means as Floridi states at the end of his 2004 paper, "I hope the reader will be thrilled rather than depressed by the amount of work that lies ahead". Getting into a *who is right/who is wrong* style of argument would simply be premature, instead I utilise Armstrong's work to highlight areas that require development in informational realism. For example no systematic metaphysics would be complete without dealing with a variety of sub-issues, for example, causality, modalities and the intriguing problem of absences (which I separate out and deal with separately as part of chapter 6) amongst other issues. As is often the case in philosophy I conclude this chapter with a set of open problems specific to informational metaphysics where we might direct future research efforts.

A quote from David Armstrong provided the initial inspiration for chapter 5. After reading "Everything that exists makes a difference to the causal powers of something" (Armstrong, 1997, p41) it occurred to me that to tell a complete story about the nature of information one must address issues of causality and chapter 5 attempts to address this in part from the point of view of information's explanatory value. Wesley Salmon, a well

known contemporary writer on scientific explanation spends considerable time in developing a causal theory of explanation. Those familiar with recent work in scientific explanation will notice that the role information plays within the discourse is by and large taken for granted. Peter Railton (1981, p240), within the context of his explanatory theories points out with some concern that while it is typical to speak of sentences or texts conveying information he knows of "no satisfactory account of this familiar and highly general notion" (1981, p. 240). Salmon on the other hand utilises ideas about the transmission of information to develop a causal theory of scientific explanation. This chapter contributes to my overriding goal regarding the nature of information but it also offers what I think represents a unified and information-theoretic approach to scientific explanation, one which addresses Railton's concerns and incorporates Salmon's goal of a coherent causal theory of explanation.

Chapter 6 takes an epistemic turn. In the chapter I consider issues concerning how we interact with the world in order to aquire knowledge, the role information plays and its relationship with knowledge. The chapter begins by exploring an important question within the philosophy of information, that is, can we develop a coherent *information-first* epistemology. The hope is that this may represent a viable departure from doxastic approaches to knowledge which have been substantially challenged by Gettier's 1963 article entitled *Is Justified True Belief Knowledge?*

However, rather than considering the validity of an *information-first* epistemology (which has been considered in detail by Patrick Allo, 2007) the chapter considers assumptions in the question itself. The assumption being that *information* is, in some asymmetric sense, prior to and encapsulated by *knowledge*, it is the validity of this assumption that is called into question by the chapter. Once again the inspiration for the research comes from Floridi's article, this time, *Problem 13: The continuum hypothesis: Should epistemology be based upon a theory of information?* (2004a, p570). The chapter contributes to the underlying goal of this thesis

in clarifying and contributing to an understanding of the nature of information. I draw upon two important thinkers' material, Timothy Williamson with reference to his text *Knowledge and its Limits* (2002) and Gareth Evans's *The Varieties of Reference* (1982) and with reference to some of their arguments I present evidence that suggests that the relationship between information and knowledge appears to be coexstensive. Thus the assumption of a strict asymmetry between information and knowledge where knowledge encapsulates and is something-more-than information I argue is a convenient myth. If my argument is correct then knowledge seems to be not much more than a special kind of information. A very simple formulation of the argument might be something like this; it seems quite natural to say, to *know* something is to be *informed*. In support of my argument, Williamson (2000) defends a principle whereby one's evidence (what in fact looks very similar to a claim that 'one's information') is equal to one's knowledge. If this turns out to be the case, then it indeed does appear that we can claim a coexstensive relationship between information and knowledge. Evans (1982) on the other hand devotes a large part of his text toward analysis of the role reference plays in knowledge and information. Reference and its relationship with information is something of a subplot within this thesis, one which I visit on a couple of occasions throughout this work however this particular chapter looks at this issue in some detail. For Evans at least information plays a critical role within his theory of reference. Information is conveyed via perception and transmitted via language thus the factual component of any information based statement depends upon the empirical confirmatory experience to which it is linked. While my findings in this chapter may seem somewhat modest, the chapter does support my overall goal of developing some empirical support for an informational ontology and as such contributes towards answers to Floridi's problem 15 (2004) regarding the ontological status of information. As for an important question flagged early in the chapter, if it turns out the so called *Information Continuum* is merely a convenient myth, what

implications this might have for information ethics and philosophy of information generally I must leave for future research.

The final chapter of this thesis in keeping with the naturalistic theme running throughout this work is a critique of Floridi's new and novel metaethical theory, Information Ethics. While Floridi's theory relies upon two significant constructs within his informational realism, the *informational object* (discussed critically in detail in chapter 3) as well as another concept drawn from computing theory (also discussed in chapter 3), *levels of abstraction* it is his use of the term *informational entropy* with respect to his metaethics which is considered in detail. Informational entropy according to Floridi is information ethics's equivalent of evil or suffering in more orthodox ethical theories. In information ethics the loss, degradation, corruption or depletion of information objects is considered entropy. Floridi then outlines what he calls the patient oriented ontocentric principles. These normative principles instruct us to not cause entropy, prevent or remove it where possible, and ultimately promote the flourishing of all informational objects. Floridi uses the term entropy quite differently to the way the term is used within other scientific contexts. In fact Floridi distances himself from uses of the term relating to the laws of thermodynamics for instance. In this chapter I argue that this is a mistake. Indeed the chapter is an attempt at reconciling Floridi's use of the term within his information ethics with materialist, physicalist, bio-chemical as well as uses within evolutionary biology. In keeping with the principle of the conservation of energy within thermodynamic theory I apply this to information ethics and entropy in arguing contrary to Floridi that ultimately information is conserved within an isolated system, that is, it is neither created nor removed from the system but instead may change location and form. My reconstrual of the term entropy within information ethics acknowledges that even at an informational level of abstraction, entropy is both an irreversible unidirectional process and that almost all human activity results in increases in entropy. The chapter argues that information, in

keeping with materialist conceptions of entropy, ultimately becomes more austere and distributed. However I argue that this is a non-loss process of informational decomposition. Bringing the concept of informational entropy in line with scientific construals I believe strengthens information ethics, by acknowledging that many processes as well as decisions we make are entropic by their very nature this guides us towards an ethical or moral stance related to informational objects (which might be ecosystems, mountains, human beings or everyday objects) that seeks to minimise entropy causing decisions. This position is also sympathetic with Floridi's attempt to clarify the level of moral value that ought to be assigned to informational objects of the kind that many of us would consider of dubious value. Informational objects that are the end product of some entropic process (Brey's (2008) example included excrement and vats of toxic waste for example) or that are themselves entropy causing (coal fired power stations) have a moral value that may be minimal and overridable. Of course in the case of the coal fired power station, this value is traded against the fact that electricity is necessary in order to preserve order and prevent decay and disorder both at societal levels (where we analyse a society as a complex informational object) as well as individually.

This final chapter I believe contributes some deeper insights into the nature of information not only within the context of information ethics but also to the philosophy of information generally, and with particular reference to the ontological status of information. I defend the claim contrary to Floridi that informational entropy is by and large a process of non-loss informational decomposition.

8.1 The Last Word

During the past decade the philosophy of information as an emerging conceptual branch of philosophy has been gathering momentum. It is clearly multidisciplinary and includes not only investigations into the conceptual nature of information but also covers the elaboration and

application of information-theoretic methodologies to existing philosophical questions (Floridi, 2002). My initial overarching research question and the one which originally motivated this thesis was concerned with the conceptual nature of information:

What is the ontological status of information?

This certainly situates the general focus of this research within the metaphysical realm. My aims were to clarify theoretical and conceptual notions about information that might be brought to bear upon other philosophical problems such as epistemology, scientific explanation, causality and ethics. I discovered, in particular during the final months of this work that while the initial problem was a metaphysical one, my methodology has been by and large naturalistic. In trying to illustrate difficult concepts I drew upon empirical examples from science and general observation. It was initially my supervisor Morgan Luck that enquired: am I doing epistemology or metaphysics? It occurred to me that this potentially sets up a tension. Speculative metaphysics in particular and empiricism are often at odds. To accept one seems to entail dropping the other (see van Fraassen, 2002). However following Chakravartty (2004) I view my somewhat naturalistic method not so much as making any particular claims about reality but rather as a strategy by which the analysis in this work has proceeded. To clarify, naturalism is not so much a factual thesis about the nature of the world, rather it is a stance which defers to natural science with regard to both truth and completeness with regard to the nature of the world (Chakravartty, 2004, p175).

Thus, my endorsment of empiricism is in no way doctrinal (that is, I am not claiming that the only source of knowledge about the world is via experience). In the same sense it seems odd that the metaphysician would postulate a kind of world that didn't bear some correspondance with general experience. If we assume both metaphysics and empiricism are rational

then perhaps an unintended consequence of this work has been to show how application of an epistemic approach can provide support for what is essentially a metaphyscial position.

And so, while my adjustments to informational realism and its corresponding ethics are perhaps modest, this work I hope provides further evidence that Floridi's metaphysics is correct. It just turned out that the best way I could think of to illustrate and elaborate upon these ideas ended up being through empirical examples.

9 References

Adams, F. (2004). Knowledge. in *The Blackwell Guide to The Philosophy of Computing and Information* (Ed. Floridi, Luciano. 228-236). (2004). Blackwell Publishing.

Adriaans, P. (2013). Information. The Stanford Encyclopedia of Philosophy, Edward N. Zalta (ed.), URL=<http://plato.stanford.edu/archives/fall2013/entries/information/>.

Allo, P. (2007*). On Logics and Being Informative: Pluralism, Locality, and Feasibility*. PhD Thesis, Vrije Universiteit Brussel.

Armstrong D.M. (1989). *Universals: An Opinionated Introduction.* Westview Press. (Focus Series)

Armstrong, D. M. (1998). *A World of States of Affairs*. Cambridge: Cambridge University Press.

Armstrong, D. M. (2010). *Sketch for a Systematic Metaphysics*. Oxford: Oxford University Press.

Bar-Hillel, Y., Carnap, R. (1953). Semantic Information. *British Journal for the Philosophy of Science, 4,* 147-157.

Barwise, J. and Seligman, J. (1997). *Information flow: the logic of distributed systems.* Cambridge: Cambridge University Press.

Beardsley, M. (1965). Intrinsic Value. *Journal of Philosophy and Phenomenological Research,* XXVI.

Benacerraf, P. (1973). Mathematical Truth. *The Journal of Philosophy*. Vol LXX, 661-679.

Booch, G. (1994). *Object-Oriented Analysis and Design with Applications*. Addison-Wesley.

Brennan, G., Pettit, P. (2004). Esteem, Identifiability and the Internet. *Analyse and Kritick, 26,* 139-157.

Brey, P. (2000). Method in Computer Ethics: Towards a Multi-level Interdisciplinary Approach. *Ethics and Information Technology, 2*(2), 125-129.

Brey, P. (2008). Do We Have Moral Duties Towards Information Objects?. *Ethics and Information Technology*, 10(1).

Brey, P., Soraker, J.H. (2008). Philosophy of Computing and Information Technology. In D.M. Gabbay, A.W.M. Meijers, J. Woods, & P. Thagard (Eds.), *Philosophy of Technology and Engineering Sciences*: 9 (pp. 1341-1408). North Holland.

Bynum, T. W. (1991). Pop, Para, and Theoretical Computer Ethics. *Proceedings of the National Conference on Computing and Values*.

Bynum T., M. J. (1998). *The Digital Phoenix*. Cambridge: Blackwell.

Bynum, T. W. (2000). A Very Short History of Computer Ethics. In *APA Newsletters on Philosophy and Computers*, http://www.southernct.edu/organisations/rccs/resources/research/introduction/bynum_shrt_hist.html:.

Bynum, T. (2008). Norbert Wienerand the Rise of Information Ethics. In J. van den Hoven and J. Weckert, *Moral Philosophy and Information Technology*. Cambridge: Cambridge University Press.

Bynum, T. (2008). Computer and Information Ethics. The Stanford Encyclopedia of Philosophy (Spring 2011 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/spr2011/entries/ethics-computer/>.

Callicott, B. (1995). Intrinsic Value in Nature: A Metaethical Analysis. *The Electronic Journal of Analytic Philosophy, 3*(Spring (1995).

Carnap, R. (1950). Empiricism, Semantics, and Ontology. *Revue Internationale de Philosophie*, 4, 20-40.

Carr, N. (2004). *Does IT Matter: Information Technology and the Corrosion of Competitive Advantage*. Harvard Business Press.

Chaitin, G.J. (1969). *On the length of programs for computing finite binary sequences: statistical considerations*. J. Assoc. Comput. Mach., 16: 145–159.

Cohen, J. (1993). Freedom of Expression. *Philosophy and Public Affairs, 22*(3), 207-263.

Colburn, T., R. (2000). Information, Thought, and Knowledge. *Proceedings of the 4th World Multiconference on Systematics, Cybernetics, and Informatics*, 10, 467-471. Orlando, Florida:.

Compaine, B. (2001). *The Digital Divide: Facing crisis or creating a myth?* (Compaine, B). Cambridge MA: MIT Press.

Curry, P. (2005). *Ecological Ethics: An Introduction*. Cambridge: Polity Press.

Date, C. J., Darwen, H. (2nd Ed.). (2000). *Foundation for Future Database Systems: The Third Manifesto*. Massachusetts: Addison-Wesley.

D'Agostino, M. (2013). Semantic Information and the Trivialization of Logic: Floridi on the Scandal of Deduction. Information, 4, 33-59.

Demir, H. (2012). *Luciano Floridi's Philosophy of Technology: Critical Reflections*. London: Springer.

Dewey, J. (1939). *Theory of Valuation*. Chicago. IL: University of Chicago Press.

Dijkstra, E. (1968). The Structure of THE Multiprogramming System. *Communications of the ACM, 11*(5), 341-346.

Dillon, R. S. (2009). Respect for Persons, Identity, and Information Technology. *Ethics and Information Technology*, *12*(1), 17-28.

Dowe, P (1992). Wesley Salmon's Process Theory of Causality and the Conserved Quantity Theory. *Philosophy of Science, 59,* 195-216.

Dowe, P. (2000). *Physical Causation*. Cambridge: Cambridge University Press.

Dretske, F. I. (1981). *Knowledge and The Flow of Information*. Cambridge, Mass: MIT Press.

Esfeld, M. (2009). The Modal Nature of Structures in Ontic Structural Realism. *International Studies in the Philosophy of Science, 23*(2), 179-194.

Esfeld, M., Lam, V. (2010). Ontic Structural Realism as a Metaphysics of Objects. In Alisa and Peter Bokulich, *Scientific Structuralism* (Chapter 8). Dordrecht: Springer.

Evans, G. (1982). *Varieties of Reference*. London: Oxford University Press.

Fisher, R.A., (1925). Theory of Statistical Estimation. *Proceedings Cambridge Philosphical Society*, 22(5): 700-725

Floridi, L. (1999). Information Ethics: On the philosophical foundation of computer ethics. *Ethics and Information Technology*, 1(1), 37-56.

Floridi, L. (2002). On the Instrinsic Value of Information Objects and the Infosphere. *Ethics and Information Technology*, 3(4), 287-304.

Floridi, L. (2002). What is the Philosophy of Information. *Metaphilosophy*, 33(1-2), 123-145.

Floridi, L.. (2003). From Data to Semantic Information. *Entropy* (5), 125-145.

Floridi, L. (2004a) Open Problems in the Philosophy of Information. *Metaphilosophy, 35*(4), 554.

Floridi, L.. (2004b) Informational Realism. In G.M. Greco, *IEG Research Report Oxford*: Information Ethics Group.

Floridi, L. ed. (2004c) *The Blackwell Guide to The Philosophy of Computing and Information*. Blackwell Publishing.

Floridi, L. (2006) Information Ethics, its Nature and Scope. *ACM SIGCAS Computers and Society*, 36, (3), 21-36.

Floridi, L.. (2007). Information Ethics, its Nature and Scope. In Jeroen van den Hoven and John Weckert, *Information Technology and Moral Philosophy* (36, 40). Cambridge: Cambridge University Press.

Floridi, L. (2008). Information Ethics: A Reappraisal. *Ethics and Information Technology*, *10*(2-3), 189-204.

Floridi, L. (2008) A Defence of Informational Structural Realism. *Synthese, 161*(2), 219-253.

Floridi, L. (2008). Modern Trends in the Philosophy of Information. In *Philosophy of Information*. Holland: Elsevier.

Floridi, L. (2010). *The Cambridge Handbook of Information and Computer Ethics*. Cambridge: Cambridge University Press.

Floridi, L. (2012). The Road to Philosophy of Information In Demir, H., *Luciano Floridi's Philosophy of Technology: Critical Relections* (245). New York and London, Springer.

Floridi, L. (2012). Reply to McKinlay. In Demir, H., *Luciano Floridi's Philosophy of Technology: Critical Reflections* (267). New York: Springer.

Floridi, L. and Sanders, J.W. (2001) Artificial Evil and the Foundation of Computer Ethics. *Ethics and Information Technology*. 3(1): 55-66.

Floridi, L., Sanders, J. W. (2002). Mapping the Foundationalist Debate in Computer Ethics. *Ethics and Information Technology*, 4, 1-9.

Fodor. J. and Lepore. E. (1992). *Holism: A Shoppers Guide*. Oxford: Blackwell Publishing.

Goodman, N., Quine, W. V. (1947). Steps Toward a Constructive Nominalism. *The Journal of Symbolic Logic*, 12(4), 105-122.

Gorniak-Kocikowska, K. (1996). The Computer Revolution and the Problem of Global Ethics. *Science and Engineering Ethics*, *2*(2).

Gettier, E. L.. (1963). Is Justified True Belief Knowledge?. In *Knowledge and Belief* (Ed. A. Phillip Griffiths, 144-146). Oxford: Oxford University Press. (Originally in Analysis, Vol. 23, Blackwell pp 121-3)

Gotterbarn, D. (1991). Computer Ethics: Responsibility Regained. *The Phi Kappa Phi Journal*, *73*(3), 26-31.

Gotterbarn, D. (2007). The Use and Abuse of Computer Ethics. In Weckert, J, *Computer Ethics* (The International Library of Essays in Public and Professional Ethics, 57). England: Ashgate. (1992, Journal of Systems and Software)

Gotterbarn, D., Rogerson, S. (1998). The Ethics of Software Project Management. In Collste, G, *Ethics in the Age of Information Technology* (278-300). Sweden: Linkoping University Press.

Greeno, J. G. (1970). Evaluation of Statistical Hypotheses Using Information Transmitted. *Philosophy of Science*, *37*, 279-293.

Hart, H. L., Honore, T. (1985). *Causation in the Law*. Oxford: Clarendon Press.

Hayaki, R. (2006). Contingent Objects and the Barcan Formula. *Erkenntnis*, 64, 75-83.

Hempel C. (1962). Deductive-Nomological vs. Statistical Explanation. In *Minnisota Studies in the Philosophy of Science III* (98-169). Minneapolis: University of Minnesota Press.

Hempel, C. (1965). *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*. New York: Free Press.

Hempel, C., Openheim, P., (1948). Studies in the Logic of Explanation. *Philosophy of Science*, 15, 135-175.

Himma, E. (2004). There's something about Mary: The moral value of things *qua* information objects. *Ethics and Information Technology*, 6, 145-159.

Himma, E. (2007). The Information Gap, the Digital Divide and the Obligations of Affluent Nations. *International Review of Information Ethics, 7*.

Hintikka, J. (1973). *Logic, Language Games and Information. Kantian Themes in the Philosophy of Logic*. Oxford: Clarendon Press.

Hunter, L. (Public Image). (1995). Public Image. In Johnson, Deborah., Nissenbaum, Hellen., Computers, Ethics and Social Values (294). Englewood Cliffs, N.J.: Prentice-Hall.

Illari, P. (2012). Why Theories of Causality need Production: An Information-Theoretic Account. *Philosophy and Technology*.

Introna, L. (2005). Disclosive ethics and information technology: Disclosing facial recognition systems.. Ethics and Information Technology, 7(2), 75-86.

Johnson, D. G. (1994). *Computer Ethics* (2nd Ed.). Prentice-Hall.

Johnson, D. G. (1999). Sorting out the Uniqueness of Computer-Ethical Issues. *Etica & Politica / Ethics & Politics*, 2.

Johnson, D. G. (2006). Computer Systems: Moral entities but not moral agents. *Ethics and Information Technology*, 8, 195-204.

Kaila, V. R., Annila, A. (2008). Natural Selection for Least Action. *Proceedings of The Royal Society*, 464, 3055-3070.

Kitcher, P. (1989). Explanatory Unification and the Causal Structure of the World. In *Scientific Explanation* (410-505). Minneapolis: University of Minnesota Press.

Kolmogorov, A.N., (1965). Three Approaches to the Quantitative Definition of Information. *Problems Inform, Transmission*, 1(1): 1-7.

Korsgaard, C. (1983). Two Distinctions in Goodness. *The Philosohical Review, 92*, 169-195.

Kripke, S. (1972). *Naming and Necessity*. Oxford: Basil Blackwell.

Kvasny, L. (2007). The Existential Significance of the Digital Divide for America's Historically Underserved Populations. In Hongladarom, S. & Ess, C., *Information Technology Ethics: Cultural Perspectives* (200-212). Hershey: Idea Group Publishing.

Lauden, L. (1977). *Progress and Its Problems: Towards a Theory of Scientific Growth*. California: University of California Press.

Ladyman, J. (1998). What is Structural Realism. *Studies in the History and Philosophy of Science, 29A*(3), 409-424.

Lehninger (1993). *Principles of Biochemistry* (2nd Ed.). W H Freeman and Co.

Lewis, D. (1973). *Counterfactuals*. Oxford: Oxford University Press.

Lewis, D. (1986) Events. in *Philosophical Papers*, Vol II., pp. 241–269, Oxford: Oxford University Press

McKinlay, S. (2012). The Floridian Notion of the Information Object. In Demir, H., *Luciano Floridi's Philosophy of Technology: Critical Reflections* (223). New York and London: Springer.

Manders-Huits, N. (2010). *Designing for Moral Identity in Information Technology*. Eindhoven: 3TU Centre for Ethics and Technology.

Maner, W. (1996). Unique Ethical Problems in Information Technology. *Science and Engineering Ethics, 2*(2), 137-154.

Martin, J. Odell, J. (1992). *Object Oriented Analysis and Design*. Prentice-Hall Gale.

Mason, R. (1986) Four Ethical Issues of the Information Age, *MIS Quarterly*, 10.pp 19-26

Mathiesen, K. (2008). Censorship and Access to Expression. In Himma, K & Tavani, H, *The Handbook of Information and Computer Ethics* (573-587). New Jersey: John Wiley & Sons Inc.

Merleau-Ponty, M. (1962), *Phenomenology of Perception*, trans, Smith, London: Routledge and Kegan Paul.

Moor, J. H. (1978), Three Myths of Computer Science. *The British Journal for the Philosophy of Science*, 29(3): 213–222

Moor, J. (1985). What is Computer Ethics. *Metaphilosophy*, *16*(4), 266-275.

Moore, G. E. (1903). *Principia Ethica*. Cambridge: Cambridge University Press.

Maxwell, G. (1970). Structural Realism and the Meaning of Theoretical Terms. In Stephen Winokur and Michael Radner, *Analysis of Theories, and Methods of Physics and Psychology* (181-192). Minneapolis: University of Minnesota Press.

Nissenbaum, H. (1998). Protecting Privacy in an Information Age: The Problem of Privacy in Public. In Weckert. J, *Computer Ethics* (The International Library of Essays in Public and Professional Ethics, p331). Hampshire, England: Ashgate.

Oddie, G. (1982). Armstrong on the Eleatic Principle and Abstract Entities. *Philosophical Studies, 41,* 285-295.

Parker, D. (1981). *Ethical Conflicts in Computer Science and Technology*. Arlington, VA: AFIPS Press.

Pascal, F. (2000). *Practical Issues in Database Management: A Reference for the Thinking Practitioner*. New Jersey: Addison-Wesley.

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

Priest, G. (2000). *Logic: A Very Short Introduction*. Oxford: Oxford University Press.

Quine, W. V. (1953). Two Dogmas of Empiricism. In *From a Logical Point of View*(20-46). Harvard: Harvard University Press.

Quine, W. V. (1953). On What There Is. In *From a Logical Point of View* (1). Harvard University Press.

Quine, W. V. (1957). Speaking of Objects. *Proceedings and Addresses of the American Philosophical Association*, 31, 5-22.

Quine, W. V., (1960). *Word and Object*. Cambridge, Massachusetts: MIT Press.

Quine, W. V. (1969). Epistemology Naturalised. In *Ontological Relativity and Other Essays* (75). New York: Columbia University Press.

Quine, W. V. (1992). Structure and Nature. *Journal of Philosophy*, 89(1), 5-9.

Railton, P. (1981). Probability, Explanation, and Information. *Synthese*, 48(2), 233.

Railton, P. (1978). A Deductive-Nomological Model of Probabilistic Explanation. *Philosophy of Science*, 45, 206-226.

Rosen, G. (1994). Objectivity and Modern Idealism. in M. Michael and J. Hawthorne-O'Leary (eds), *Philosophy in Mind*, Dordrecht: Kluwer Academic Publishing.Rumbaugh, J., Blaha, M., Lorensen, W., Eddy, F., Premerlani, W. (1991). *Object-Oriented Modeling and Design*. Pearson Education.

Salmon, W. (1971). Statistical Explanation. In *Statistical Explanation and Statistical Relevance* (29-87). Pittsburgh: University of Pittsburgh Press.

Salmon, W. (1989). *Four Decades of Scientific Explanation*. Minneapolis: University of Minnesota Press.

Salmon, W. (1998). *Causality and Explanation*. Oxford: Oxford University Press.

Schaffer, J. (2004). Causes Need Not be Physically Connected to their Effects: The Case for Negative Causation. In Christopher Hitchcock, *Contemporary Debates in Philosophy of Science*. Blackwell.

Shrodinger (1944). *What is Life?*. Cambridge: Cambridge University Press.

Searle, J.R. (1980) Minds, Brains and Programs. *Behavioural and Brain Sciences*, 3, no. 3:417-57.

Sewell, G., Barker, James, R.. (2001). Neither good, nor bad, but dangerous: Surveillance as an ethical paradox. *Ethics and Information Technology*, 3(3), 183-196.

Shoemaker, D. (2009). Personal Identity and Ethics. In *Stanford Encyclopedia of Philosophy*, http://plato.stanford.edu/entries/identity-ethics/

Stachel, J. (2005). Structure, Individuality and Quantum Gravity. In Rickles, D.P., French, S.R.D., Saatsi, J., *Structural Foundations of Quantum Gravity*. Oxford: Oxford University Press.

Stahl, B. C. (2008). Discourses on Information Ethics: The Claim to Universality. *Ethics and Information Technology, 10*(2-3), 97-108.

Stepanova, E. (2011), *The Role of Information Communication Technologies in the "Arab Spring": Implications Beyond the Region.* Ponars Eurasia Policy Memo No. 159. Ponars Eurasia.

Suber, P. (1988). What Is Software? *Journal of Speculative Philosophy*, 2(2): 89–119

Tavani, H. (2002). The Uniqueness Debate in Computer Ethics: What Exactly is at Issue, and Why Does it Matter. *Ethics and Information Technology, 4*(1), 36-54.

Tainter, J. (2003). *The Collapse of Complex Societies*. New York and Cambridge: Cambridge University Press.

Tavani, H. (2004) Balancing intellectual property rights and the intellectual commons: A Lockean analysis. *Journal of Information Communication and Ethics in Society* (2), S5-S14.

Tavani, H. (2008) Floridi's Ontological Theory of Informational Privacy: Some Implications and Challenges. *Ethics and Information Technology*, 10(2), 155-166.

Turner, R. (2013). The Philosophy of Computer Science, *The Stanford Encyclopedia of Philosophy* (Fall 2013 Edition), Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/fall2013/entries/computer-science/>.

van den Hoven, J. (2008). Information Technology, Privacy, and the Protection of Personal Data. In van den Hoven, J & Weckert, J, (eds.) *Information Technology and Moral Philosophy*. Cambridge University Press.

van Fraassen, B. (1980). *The Scientific Image*. Oxford: Oxford University Press.

van Fraassen, B. (2002). *The Empirical Stance*. London & New Haven: Yale University Press.

van Fraassen, B. (2006). Structure: Its Shadow and Substance. *The British Journal for the Philosophy of Science, 57*, 275-307.

Vedder, A., Wachbroit, R.. (2004). Reliability of information on the Internet: Some distinctions., *Ethics and Information Technology*, 5(4), 211-215.

Weckert, J. (Ed.). (2007) *Computer Ethics*. England: Ashgate.

Wiener N. (1954) *The Human Use of Human Beings: Cybernetics and Society*. Houghton and Mifflin. 2nd Ed. Double Day Anchor.

Wiener, N. (1961). *Cybernetics or Control and Communication in the Animal and the Machine*. Cambridge, MA: MIT Press.

Williams, D. C. (1966). The Elements of Being. In *The Principles of Empirical Realism*. Springfield, Illinois: Charles C. Thomas

Williamson, T. (2002). *Knowledge and its Limits*. Oxford: Oxford University Press.

Wittgenstein, L. (1961). *Tractatus Logico-Philosophicus*. London and New York: Routledge.

Woodward, J. (1987). On and Information-Theoretic Model of Explanation. *Philosophy of Science, 54*, 21-44.

Worrall, J. (1989). Structural Realism: The Best of Both Worlds. *Dialectica, 43*, 99-124.