On the intrinsic value of information objects and the infosphere

Luciano Floridi
*Dipartimento di Scienze Filosofiche, Università degli Studi di Bari and Faculty of Philosophy, Sub-Faculty of Computation, Programme in Comparative Media Law and Policy, Information Ethics Group, Oxford University (*address for correspondence: Wolfson College, OX2 6UD, Oxford, UK; E-mail: luciano.floridi@philosophy.oxford.ac.uk)*

Abstract. What is the most general common set of attributes that characterises something as intrinsically valuable and hence as subject to some moral respect, and without which something would rightly be considered intrinsically worthless or even positively unworthy and therefore rightly to be disrespected in itself? This paper develops and supports the thesis that the minimal condition of possibility of an entity's least intrinsic value is to be identified with its ontological status as an information object. All entities, even when interpreted as only clusters of information, still have a minimal moral worth *qua* information objects and so may deserve to be respected. The paper is organised into four main sections. Section 1 models moral action as an information system using the object-oriented programming methodology (OOP). Section 2 addresses the question of what role the several components constituting the moral system can have in an ethical analysis. If they can play only an instrumental role, then Computer Ethics (CE) is probably bound to remain at most a practical, field-dependent, applied or professional ethics. However, Computer Ethics can give rise to a macroethical approach, namely Information Ethics (IE), if one can show that ethical concern should be extended to include not only human, animal or biological entities, but also information objects. The following two sections show how this minimalist level of analysis can be achieved. Section 3 provides an axiological analysis of information objects. It criticises the Kantian approach to the concept of intrinsic value and shows that it can be improved by using the methodology introduced in the first section. The solution of the Kantian problem prompts the reformulation of the key question concerning the moral worth of an entity: what is the intrinsic value of x *qua* an object constituted by its *inherited* attributes? In answering this question, it is argued that entities can share different observable properties depending on the level of abstraction adopted, and that it is still possible to speak of moral value even at the highest level of ontological abstraction represented by the informational analysis. Section 4 develops a minimalist axiology based on the concept of information object. It further supports IE's position by addressing five objections that may undermine its acceptability.

Key words: axiology, computer ethics, dignity, entropy, information ethics, information object, infosphere, intrinsic value, Kant, Kingdom of Ends, Kingdom of Nature, ontology, object-oriented programming, respect

1. Introduction: An object-oriented model of moral action

This section introduces the technical concepts and terminology necessary to develop an informational approach to the concept of moral patient. The reader acquainted with the Object Oriented Programming (OOP) methodology[1] may wish to move directly to section two.

The first task is to analyse a moral action as a dynamic system arising out of the interaction of seven principal components: (1) the agent, (2) the patient, (3) their interactions, (4) the agent's general frame of information, (5) the factual information concerning the situation that is at least partly available to the agent, (6) the general environment in which the agent and the patient are located, and finally (7) the specific situation in which the interaction occurs. The second task is to show how this dynamic system can be modelled in terms of an information system by using the OOP methodology.

Any action, whether morally loaded or not, has the logical structure of a variably interactive process relating one or more sources (depending on whether one is working within a multiagent context), the agent a, with one or more destinations, the patient p. The agent initiates the process and the patient reacts more or less interactively to it.[2] Once a and b are interpreted, their analysis depends on the level of abstraction

[1] This article follows the standard terminology and conceptual apparatus provided by Rumbaugh 1991. On conceptual modelling of informational systems see also Flynn and Diaz Fragoso (1996), Veryard (1992) and Boman et al. (1997).

[2] The terms 'agent' and 'patient' are standard in Ethics and therefore will be maintained in this paper; however, it is essential to stress the interactive nature of the process and hence

(LoA) adopted and the corresponding set of observables available at that level.[3] Suppose, for example, that we interpret p as Mary (p = Mary). Depending on the LoA and the corresponding set of observables, p = Mary can be analysed as the unique individual person called Mary, as a woman, as a human being, as an animal, as a form of life, as a physical body and so forth. The higher the LoA, the more impoverished is the set of observables, and the more extended is the scope of the analysis. As the Turing Test shows, 'erasing' observables raises the LoA, until it becomes impossible to discriminate between two inputs sources. If Mary is analysed as a human being, more observables could lead one to analyse Mary at a lower LoA as a woman, and less observables could lead one to analyse Mary at a higher LoA as an animal.

At the LoA provided by an informational analysis (LoA^{i}), both a and p are information objects. In our example, this means that p = Mary is analysed as an information object that interacts and shares a number of properties with other information objects, like a digital customer profile. It does not mean that a or p are necessarily *only* information objects.

The OOP approach provides a very flexible and powerful methodology with which to clarify and make precise the concept of 'information object' as an entity constituted by a bundle of properties, to use a Humean expression. Before introducing it, an example may help to outline the basic idea.

Consider a pawn in a chess game. Its identity is not determined by its contingent properties as a physical body, including its shape and colour. Rather, a pawn is a set of data (properties like white or black and its strategic position on the board) and three 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, and it can be promoted to any piece except a king when it reaches the opposite side of the board. For a good player, the actual piece is only a placeholder. The real pawn is an 'information object'. It is not a material thing but a mental entity, to put it in Berkeley's terms. The physical placeholder can be replaced by a cork without any semantic loss at the LoA required by the game.

Let us now turn to a more theoretical analysis. OOP is a revolutionary method of programming that has radically changed the approach to software development. Historically, a program was viewed as an algorithmic procedure that takes input data, processes it, and produces output data. The difficulty was then represented by the elaboration of the algorithmic process. OOP has shifted the focus from the logic procedures required to manipulate the objects to the objects that need to be manipulated. In OOP, data structures (cf. the pawn's property of being white) and their behaviour (programming code, cf. the pawn's power to capture pieces only by moving diagonally forward) are packaged together as information objects. Objects are then grouped in a hierarchy of classes (e.g., pawns), with each class inheriting characteristics from the class above it (e.g., all pieces but the king can be captured, so every pawn can be captured). 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. The routines or logic sequences that can manipulate the objects are called *methods*. A method is a particular implementation of an operation, i.e., an action or transformation that an object performs or is subject to by a certain class. Objects communicate with each other through well-defined interfaces called *messages*. Examples of objects range from the buttons and scroll bars in a window to human beings like Mary (described by name, address, and so forth), from stock-exchange shares to buildings and pawns. This ontological concept should not be confused with the purely syntactical and quantitative concepts of information available in information and computation theory, or with the semantic approach popular in the philosophy of language, in the philosophy of mind and in cognitive science. Henceforth, 'information object' and its cognate terms will be used in the OOP sense just introduced. Small caps (LIKE THIS) will hep to indicate that the object in question is an information object.

Let us now return to the informational modelling of a and p. When a and p are analysed as information objects at LoA^{i} this means that they are considered and treated as discrete, self-contained, encapsulated[4] packages containing (i) the appropriate data structures, which constitute the nature of the entity in question: state of the object, its unique identity, and attributes; and (ii) a collection of operations, functions, or

the fact that the patient is hardly ever a passive receiver of an action. The unidirectional, bivalent, causal model is often far too simplistic. A better way to qualify the patient in connection with the agent would be to refer to it as the 'reagent'.

[3] Dijkstra (1968) and Parnas (1972) are two classic papers introducing the concept of LoA, Medvidovic (1996) provides a review.

[4] Encapsulation or information hiding is the technique of keeping together data structures and the methods (class-implemented operations), which act on them in such a way that the package's internal structure can be accessed only by means of the approved package routines. External aspects of an object, which are accessible to other objects, are thus separated from the internal implementation details of the object itself, which remain hidden from other objects.

procedures (methods), which are activated (invoked) by various interactions or stimuli, namely messages received from other objects or changes within itself, and correspondingly define how the object behaves or reacts to them. Both *a* and *p* are sufficiently permanent (continuant) information objects. They can be simple or complex systems constituted by less complex information objects.

The moral action itself can now be modelled as an information process, i.e., a series of messages (M), invoked by *a*, that brings about a transformation of states *directly* (more on this qualification shortly) affecting *p*, which may variously respond to M with changes and/or other messages, depending on how M is interpreted by *p*'s methods.

So the first three information components of our system are *a*, *p* and M. The fourth component is the personal or subjective frame of information within which the agent operates. This shell,[5] which is really an integral part of *a*'s nature but that it is useful to treat separately, is the information frame that encapsulates the subjective world of information of the agent (*a*'s subjective *infosphere*, see below). It is constituted by internally dynamic and interactive records (modules) of *a*'s moral values, prejudices, past patterns of behaviour, attitudes, likes and dislikes, phobias, emotional inclinations, moral beliefs acquired through education, past ethical evaluations, memories of moral experiences (e.g., of similar situations in which she acted as a witness) or of other moral actions performed in the past, and so forth. In short, it represents the ethical and epistemic conceptualising interface between *a* and the environment. The shell, although it embodies aspects of the agent's life, is constantly evolving through time, may contain shared or imported elements from other agent's shells, may be epistemically only partly accessible to *a* herself and in practice only partly under the control of *a*'s will. Nevertheless, it contributes substantially to the shaping of *a*'s behaviour, by screening *a* from the direct impact of the information environment, filtering and regulating *a*'s access to, and hence highlighting and interpreting the relevant aspects of, the factual information concerning the specific moral situation in which the agent is involved in space and time.

The factual information concerning the moral situation represents the fifth dynamic component of the system. It is the only element in the model that remains unmodified when the LoA changes. We still speak of factual information even at the lower LoA, where there are sufficient observables to analyse both *a* and *p* not just as two information objects but also as two persons, for example. For this reason, the majority of ethical theories are ready to recognise factual information as playing an instrumental role in any moral action. Socratic positions explain the existence of evil in terms of ignorance. According to Warnock 1971, for example, lack of information is one of the main factors that cause 'things to go badly'. More 'weakly', it is common to assume that an action with a potential moral value can be treated as actually moral or immoral only insofar as its source *a* is, among other things, conscious (*a* is aware of *a*'s actions), sufficiently free (*a* is rationally autonomous in the Kantian sense, and can intentionally bring about, stop or modify the course of action in question, at least partly, depending on the situation), reasonable (*a* is intelligent in Mill's sense, i.e., has some capacity to forecast the consequences of *a*'s actions) and informed. Traditional ethical theories share the view that a moral action and its corresponding evaluation can take place in a state of only relative scarcity of freedom and information and that there is no morality in a state of total determinism or ignorance (cf. animal behaviour).

We now come to the sixth component. At LoA^{i}, *a* does or does not implement, and hence variously controls and adjusts, *a*'s autonomous and informed behaviour in a dynamic interaction with the general environment in which *a* is located, e.g., a given culture, society, family situation, financial status, group of individuals, set of working conditions, and so forth. The same holds true for *p*. In Floridi (1999a, b), this informational environment has been described as the *infosphere*. It is a context constituted by the whole system of information objects, including all agents and patients, messages, their attributes and mutual relations.

The specific region of the infosphere in space and time within which the moral action takes place represents the last component of the system, namely the moral situation. Borrowing a term from robotics, this information microworld can be defined as the *envelope*[6] of the moral action.

To summarise, here is the complete list of information components:

1. *a* = moral agent
2. *p* = moral patient
3. M = moral action, constructed as an interactive information process
4. *shell* = a's personal world of information

[5] The term comes from the operating system architecture vocabulary, not from OOP. It is the portion of the operating system that defines the interface between the operating system and its users.

[6] The 'envelope' of a robot is the working environment within which it operates or, more precisely, the volume of space encompassing the maximum designed movements of all the robot's parts.

5. *factual information* = information about the moral situation
6. *infosphere* = the general environment
7. *envelope* = the moral situation

Two comments are now in order. First, when the message is a reflective process or a process with a feedback effect, a may be identical with, or treated as, one of the p. I shall come back to this important point in section 3. Second, it is hardly ever the case that a message affects only a discrete set of well-specified patients p. It is convenient to limit our attention to a simplified dynamic model, and this is why I specified 'directly' above, but one needs to remember that a message functions like a vector, with a given direction and a discrete force, not as a binary switch. Once the message has been released, its direct and indirect effects almost immediately cease to be under the control of its source a, while their life extends in time and space, in the form of a gradually decreasing continuum. Using another OOP concept, we can then speak of the *propagation* of an operation, which starts at some initial OBJECT and flows from OBJECT to OBJECT through association links in the system and according to possibly specifiable rules. During the propagation, the vector may change both in direction and in force. Clearly, a message affects not just the immediate target of the process but also the envelope – hence a as well, a's shell and the factual information – and finally the whole infosphere. Think of an abused child who, as an adult, becomes an abuser. In principle, all seven components may be treated as patients. We shall see in section 4.3 that a negative axiology (a theory of intrinsic unworthiness) requires a more adequate conception of what kind of entity may count as a patient.

2. The role of information in ethics

We are now ready to phrase the foundationalist problem in CE in terms of LoA^{i}. Does LoA^{i} provide an additional perspective that can further expand the ethical discourse so as to include the world of morally significant phenomena involving information objects? Or does LoA^{i} represent a threshold beyond which nothing of moral significance really happens? Does looking at reality through the highly abstract lens of information analysis improve our ethical understanding or is it an ethically pointless (when not misleading) exercise?

In Floridi (1999a), it is argued that if information objects can have at most only an instrumental value then CE is likely to remain at most a Microethics, i.e., a practical, field-dependent, applied or professional ethics (Gotterbarn 1991, 1992, 2001; Langford 1995), which plays only an ancillary role with respect to other Macroethics, i.e., theoretical, field-independent, applicable ethics such as Deontologism or Consequentialism (Johnson 2000). This is because Macroethics attempt to establish not just the necessary and sufficient conditions of adequacy for the occurrence of a moral action, e.g., its information input, but, more importantly, what ought to be the very nature of the action in question, why a certain action would be morally right or wrong, what ought to be done in a given moral situation, and what the duties, the 'oughts' and 'ought nots' of a moral agent may be.

Still in Floridi (1999a), it is argued that IE, as the foundation of CE, can develop a Macroethical approach. To do so, IE needs to be able to show that the agent-related *behaviour* and the patient-related *status* of information objects *qua* information objects can be morally significant, over and above the instrumental function that may be attributed to them by other ethical approaches, and hence that they can contribute to determining, normatively, ethical duties and legally enforceable rights.

IE's claim consists of two theses. The first thesis states that information objects *qua* information objects can be moral agents. This means not just *analysing* an interpreted a as an information object (e.g., a = Mary) – this is elementary, as it requires only the adoption of the right LoA – but rather showing that a can be correctly *interpreted* as an information object (e.g., that an artificial agent, like a piece of software, can play the role of a moral agent) at the usual LoA adopted by other ethical theories, that is at the LoA where p = Mary is analysed as a human being. The thesis has been discussed and defended in Floridi and Sanders (2001a, b). It will not be re-addressed here.

The second thesis states that information objects *qua* information objects can have an intrinsic moral value, although possibly quite minimal, and hence that they can be moral patients, subject to some equally minimal degree of moral respect – *a disinterested, appreciative and careful attention* (Hepburn 1984) – again at the ordinary LoA where a = Mary is analysed as a human being. The task of the rest of this article is to clarify and support this second thesis. Since the strategy is complex, it may be worth outlining it at the outset.

The issue is approached top-down, starting from the discussion of the unproblematic case in which the patient is an ordinary human being, who is recognised to have intrinsic moral worth. At this low LoA, one of the best philosophical positions available, namely Kant's, suggests that *only* rational beings have an intrinsic moral worth. The objection is that the Kantian position is not fully satisfactory, and needs to be modified. This is shown by gradually impoverishing the

ontological status of p = Mary (for a discussion of this 'argument from marginal cases' see Baird Callicott 1980). By eliminating more and more of the properties enjoyed by Mary, the LoA is raised until the stage is reached at which, on the one hand, one would still like to be able to understand why p = Mary may still enjoy some degree of intrinsic moral value and hence be subject to some level of moral respect, even if Mary is reduced to a mere brainless entity, but on the other hand the Kantian analysis is unable to provide a satisfactory answer. At this point, two arguments support the attribution of an intrinsic moral value to information objects. The first, positive argument consists in showing that an information-object-oriented approach can successfully deal with the problem left unsolved by Kant. The second, negative argument consists in dismantling not only the Kantian position but also any other position that adopts some other LoA higher than the Kantian-anthropocentric one but still lower than LoA^i, like a biocentric LoA. Showing that both an anthropocentric and a biocentric axiology are unsatisfactory is a crucial step, since it re-opens the fundamental problem of what entities can qualify as centres of some moral worth, allows one to approach it afresh, and shifts at least part of the burden of proof on IE's critics. If ordinary human beings are not the only entities enjoying some form of moral respect, what else qualifies? Only sentient beings? Only biological systems? What justifies including some entities and excluding others? Suppose we replace an anthropocentric approach with a biocentric one. Why biocentrism and not ontocentrism? Why can biological life and its preservation be considered morally relevant phenomena in themselves, independently of human interests, but not existence and its protection? In many contexts,[7] it is perfectly reasonable to exercise moral respect towards inanimate entities *per se*, independently of any human interest; could it not be just a matter of ethical sensibility? It seems that any attempt to exclude non-living entities is based on some specific, low LoA and its corresponding observables but that this is an arbitrary choice. In the scale of beings, there may be no good reasons to stop anywhere else but at the bottom: 'all things in the biosphere have an equal right to live and blossom' (Naess 1973). There seems to be no good reason not to adopt a higher and more inclusive, ontocentric LoA.

In Floridi (1999a), the previous two arguments are paralleled by two other lines of reasoning, one metatheoretical the other historical. Although the reader is referred to that longer discussion for further details, it may be useful to summarise them here, since they bear further support in favour of IE's position.[8]

The metatheoretical argument has already made a brief appearance. Enlarging the conception of what can count as a centre of moral respect has the advantage of enabling one to make sense of the innovative nature of CE and to deal more satisfactorily with the original character of some of its moral issues by approaching them from a theoretically strong perspective.

The historical argument is connected with the negative one. Through time, ethics has steadily moved from a narrow to a gradually more inclusive concept of what can count as a centre of moral worth, from the citizen to the biosphere (Nash 1989). The emergence of the infosphere, as the new environment in which human beings spend much of their lives, explains the need to enlarge further the conception of what can qualify as a moral patient. IE represents the most recent development in this ecumenical trend, a Platonist environmentalism without a biocentric bias, as it were. More than fifty years ago, Leopold defined land ethic as something that 'changes the role of *Homo sapiens* from conqueror of the land-community to plain member and citizen of it. It implies respect for his fellow-members, and also respect for the community as such. The land ethic simply enlarges the boundaries of the community to include soils, waters, plants, and animals, or collectively: the land' (Leopold 1949, p. 403). The time has come to translate environmental ethics into terms of infosphere and information objects.[9]

[7] See for example the 'Principles of Archaeological Ethics' adopted by the Society for American Archaeology, http://www.saa.org/Aboutsaa/Ethics/prethic.html, *The International Journal of Cultural Property*, or the ICOM (International Council of Museums) Code of Professional Ethics http://www.icom.org/ethics.html. In many ethical codes for librarians and other library employees adopted by national library or librarians associations or implemented by government agencies (http://www.faife.dk/ethics/codes.htm), 'information objects' are considered to have a moral value and deserve respect. For example, the Italian Library Association (AIB) has endorsed a 'Librarian's Code of Conduct: Fundamental Principles' (http://www.faife.dk/ethics/aibcode.htm) that is divided into 3 sections, 'Duties toward the User', 'Duties toward the Profession' and 'Duties toward Documents and Information', where it is stated that '3.1 The librarian undertakes to promote the enhancement and preservation of documents and information'.

[8] As specified in Floridi (1999a), these arguments are 'intellectual' not 'strategic' (Norton 1989): they are addressed to the philosophically minded interlocutor, not to the reluctant policy-maker, who will more easily (or perhaps just with less difficulty) be convinced by reasonings centred on human interests.

[9] Mark Rowlands, for example, has recently proposed an interesting approach to environmental ethics in terms of naturalization of *semantic* information. According to him, 'There is value in the environment. This value consists in a certain sort of information, information that exists in the relation between affordances of the environment and their indices. This infor-

Once the intrinsic moral worth of an information object has been introduced as a viable solution to the problem left unsolved by the Kantian approach, two more tasks lie ahead. One is to show that IE's thesis is coherent. The other is to show that the main objections against it can be answered. Both tasks are undertaken in section 4. Their successful fulfilment further reinforces IE's position.

3. An axiological analysis of information

The status of *a* and *p*, and the possible modifications in the nature of both, brought about by the information process M, are the axiological elements that play a decisive role in the normative assessment of a moral action. In what follows, the analysis is restricted to *p* only (we shall return to the analysis of M in section 4.3), for three reasons.

First, the problem is whether an entity x *qua* information object – i.e., not insofar as it is a specific type of entity like Mary – can have any intrinsic moral value that could contribute to regulating a moral action affecting it. Since it is usually assumed that any entity x that can act as a moral *a* can also qualify as a moral *p* but not vice versa – e.g., it is generally assumed that animals can at most be moral *p* but not moral *a* (Rosenfeld 1995) – it is better to focus on the informational nature of an object as a possible patient.

Second, whenever the action in question is found to be either reflexive (e.g., suicide) or retroactive (e.g., moral vices acquired through the repetition of actions that are not in themselves necessarily to be deprecated from a moral point of view) the model allows the extension to x = *a* of any conclusion reached about x = *p*.

Third, by discussing the moral worth of an information object as a *p* in the most universal and abstract terms, it is possible to extrapolate from the specific nature and position taken up in a given envelope by a component of the system, and generalise the conclusions reached about *p* so as to include any possible information element that may in principle be affected by the behaviour of *a* and hence qualify as a patient of a moral process. Thus, other envelopes, the infosphere itself and the methods can be considered patients of *a*'s actions, in a way that will become fully clear once a negative axiology is developed, in section 4.3.

mation exists independently of [...] sentient creatures. [...] The information is *there*. It is in the world. What makes this information value, however, is the fact that it is valued by valuing creatures [because of evolutionary reasons], or that it would be valued by valuing creatures if there were any around' (Rowlands 2000, p. 153).

Once the analysis is restricted to x = *p*, the question to be addressed is: is there any degree of unconditional (i.e., neither instrumental nor emotional) and intrinsic worth in the nature of a *p* = information object that should determine, constrain, guide or shape *a*'s moral actions? That is, does an information object as a *p* have an intrinsic moral value that, *ceteris paribus*, could contribute to the moral configuration of *a*'s behaviour? Insofar as *p* has some intrinsic value, it contributes to the configuration of a moral action by requiring (we shall analyse later what is implied in this 'communication', see K.2 below) that *a* recognises such value in a special intentional way, that is by having *respect* for it. Now *a*'s respect for *p*'s intrinsic value consists in two things: the appreciation of *p*'s specific worth and the corresponding, uncoerced, arguably overridable disposition to treat *p* appropriately, if possible, according to this acknowledged worth. Objects capable of intentional states can have respect for *p*'s intrinsic value and hence act as moral agents at least to this extent, but are they also the only entities that can have an intrinsic value as patients?

3.1. *A critique of Kantian axiology*

According to Kant,[10] the previous question can be answered with a firm 'yes'. Either x can rightly function only as a means to an end other than itself, in which case it has an *instrumental* or *emotional* value (economic value); or x qualifies also as an end in itself, in which case it has an *intrinsic*, *moral* value *qua* x and it is valued and respected for its own sake.[11] Kant argues that anything can have an instrumental value for the sake of something else, but that only human beings, as rationally autonomous agents, can also have an intrinsic and absolute moral value, which he calls *dignity*. This is because only rationally autonomous agents, understood as 'good wills', can be the source of moral goodness, thanks to their rational and free actions. The Kantian dichotomy, intrinsic vs. instrumental value, has at least three major consequences:

(K.1) the dichotomy justifies the coextension of (i) the class of entities that enjoy moral value, (ii) the class of entities capable of moral respect,

[10] In the *Groundwork of the Metaphysics of Morals*, p. 84 (henceforth *Groundwork*, published in Kant 1996) Kant writes: 'In the kingdom of ends, everything has either a price or a dignity. What has a price can be replaced by something else as its equivalent; what on the other hand is raised above all price and therefore admits of no equivalent has a dignity.'

[11] 'Intrinsic value' is often recognised to be an ambiguous expression (cf. for example Benson 2000). It can mean 'non-instrumental value', as in Kant and in this paper (see also note 8 above), or 'inherent value' that is a value that something enjoys independently of the existence of any evaluating source.

and (iii) the class of entities that deserve to be morally respected. In Kant, the only type of entity that has moral value is the same type of entity that may correctly qualify as a moral patient and that may in principle act as a moral agent.

(K.2) the dichotomy solves the communication problem between *a* and *p*. Thanks to K.1, *a* is immediately acquainted with the moral value of *p*, and hence can respect it, because both entities belong to the same type of class, namely Kant's 'Kingdom of Ends'.[12] We shall see that, since IE rejects K.1, it cannot rely on the solution provided by K.2.

(K.3) the dichotomy implies that an entity's moral value is a kind of unconditional and incomparable worth. Either x has an instrumental value, subject to degrees, economically significant but morally irrelevant, or x has an unconditional and intrinsic value, which is morally relevant but absolute, and cannot rightly be subject to economic assessment.

The Kantian dichotomy is questionable because K.3 clashes with some basic intuitions and fails to take into account two important distinctions.

It seems reasonable to assume that different entities may have different degrees of relative value that can constrain *a*'s behaviour, without necessarily having an instrumental value, i.e., a value relative to human feelings, impulses or inclinations, as Kant would phrase it. Likewise, it seems intuitively acceptable that life, biological organisms or the absence of pain in sentient beings can all have a great deal of moral value and deserve a corresponding amount of respect. For example, one may argue that a human being who is even *inherently* (i.e., not just contingently, e.g., because of unlucky circumstances that may change) incapable of any intentional, rational and free behaviour still has some moral value, no matter how humble, which deserves to be respected, although not necessarily only for instrumental or emotional reasons. More generally, the default position seems to be that only rational beings are capable of respect but, contrary to what Kant suggests in (K.1), 'having an absolute moral value (dignity)', 'being capable of respect' and 'being intrinsically respectable' do not range over the same class of entities. Rational beings are capable of various degrees of respect, to which there seem to correspond various degrees of moral value. Kant is probably right in arguing that 'good wills' definitely deserve more respect than other entities, because they are among the conditions of possibility of what is morally good, but it requires some positive argument to show that 'good wills' do not constitute only a subclass of the entities that may have a morally significant claim on the agent, as entities subject to some respect.

All this is *prima facie* reasonable and represents a serious challenge for the Kantian dichotomy. Kant seems unduly to restrict the sense of 'relative value' to meaning only 'contingent worth depending on the agent's interest',[13] so that if x can be rightly and appropriately used only as a means, then x has no absolute value (and this follows), has only a relative value (and this also follows), but this can only mean that x's value has no moral nature whatsoever, because x's value is to be interpreted as depending only on the instrumental or emotional interest of the agent, which is a clear *non sequitur*, if one rejects the very controversial equation just spelt out.

3.2. *An OOP approach to axiology*

According to Kant, not only do the Kingdom of Ends and the Kingdom of Nature remain largely separate and independent, but the former becomes a closed system, which is allowed to rule over the latter without there being even the possibility *in principle* of the latter providing some constraints.[14] Two distinctions can help to improve Kant's anthropocentric axiology.[15]

Let us agree with Kant that there are different ways and degrees in which an entity may have some instrumental value. When the value in question is neither instrumental nor just emotional,[16] one can first distinguish between extrinsic and intrinsic value and, correspondingly, between two types of respect. An entity x has extrinsic value when it is respected as y. For example, a piece of cloth may be respected as a flag, a person may be respected as a police officer,

[12] See for example *Groundwork*, p. 85.

[13] *Groundwork*, p. 79.

[14] *Groundwork*, p. 73 ('act as if the maxim of your actions were to become by your will a universal law of nature'), see also pp. 86–88. On p. 86 Kant writes: 'all maxims from one's own lawgiving are to harmonise with a possible kingdom of ends as with a kingdom of nature', but on p. 88 it seems that only God as a single sovereign can bring together the kingdom of ends with the kingdom of nature.

[15] It is interesting to note that the four examples used by Kant to illustrate the application of the 'Law of Nature' formulation of the imperative ('act as if the maxim of your actions were to become by your will a universal law of nature') in *Groundwork*, pp. 73–75 are all 'anthropocentric' and concern only duties to oneself or to others, so when Kant speaks of the 'Formula of Humanity' version of the imperative in Groundwork, p. 80 ('So act that you use humanity, whether in your person or in the person of any other, always at the same time as an end, never merely as a means'), he employs the same four anthropocentric examples.

[16] This is Kant's 'fancy price', see *Groundwork*, p. 84.

or a practice may be respected as a cult. This sense of relative respect is associated with a sense of value that is no longer instrumental or emotional and may be called *symbolic*. Symbolic value is still utterly contingent, may be acquired or lost, and can be increased as well as reduced. In brief, it is utterly extrinsic.

In order to capture in full the fact that x has moral value in itself, a value that belongs to x in all circumstances (necessarily), not under certain conditions, and is not subject to modification unless x ceases to exist as x, one needs to consider the case in which x deserves to be respected not just symbolically, as something else, but *qua* x. It is here that the analysis must depart from the Kantian position more radically and introduce a second distinction.

The moral value of an entity is based on its ontology. What the entity is determines the degree of moral value it enjoys, if any, whether and how it deserves to be respected and hence what kind of moral claims it can have on the agent. For Kant, x's intrinsic value is indissolubly linked with x's essential nature only as a certain *type* of entity. Thus, an individual, e.g., Mary, has moral value (Kant's dignity) not as a specific person, but only insofar as she is a token of the general type 'free and rational human being', i.e., a member of the 'Kingdom of Ends'. In respecting p = Mary, the agent is not primarily or directly respecting the specific, unique and contingent individual *qua* herself, but rather the universal type she instantiates.[17] So the Kantian analysis fails to distinguish between two separate senses in which the nature of x determines x's moral value. It is a crucial oversight.

The two senses can be clarified by relying on the OOP methodology introduced in the first section. The specific nature (essence) of an object x consists in certain attributes that x could not have lacked from the start except by never having come into being as x, and cannot lose except by ceasing to exist as x. This essence is a factually indissoluble, although logically distinguishable, combination of x's *local* and *inherited* attributes. For example, if PERSON is the descendant object, and LIVING ORGANISM is the ancestor object, we may say that 'freedom' is an essential and local attribute of PERSON, that is a new property, not previously implemented in any of the ancestor objects, while 'sentient animal' is an essential attribute inherited by PERSON from its ancestor object LIVING ORGANISM. Suppose now that an object x has an intrinsic value. It is correct to say, with Kant, that x's intrinsic value depends on its essence, but it is also important to specify that this essence, and the corresponding intrinsic value, are both aggregates, i.e., they are the result of a specific combination of local and inherited attributes. What difference does this make? In the example, one can respect Mary because of her local attribute 'free agent', which is part of her essence. Her essence also includes that of being a 'living organism capable of feelings of pain and pleasure'. Let us refer to the former as Mary's F attribute and to the latter as Mary's L attribute, and let us simplify matters by saying that Mary inherits L from her ancestor object called ANIMAL. Suppose now that Mary is radically and definitely deprived of her local attribute F, that is, let us imagine that she becomes *inherently* incapable of any free and intentional behaviour, e.g., because she is born brain-dead, so that the absence of certain observables corresponds to a real ontological feature. What would be the Kantian position with respect to her moral value? There seem to be only four alternatives. None of them is fully satisfactory and this leads to the adoption of a different approach.

A radical solution would be to 'bite the bullet' and argue that

(A) Mary lacks the necessary attribute F, so she can have no justified claim to moral respect. Citizenship of the Kingdom of Ends is a necessary and sufficient condition but it can be lost and, without it, there are no moral rights.

Of course, (A) is logically acceptable, but its unpleasant consequences inevitably clash with some of the most elementary moral intuitions. According to (A), for example, one could freely dispose of Mary's organs without being subject to any moral assessment.

If one wishes to maintain that Mary still deserves to be respected despite the lack of F, one may try to argue, still with Kant, that

(B) Mary still possesses moral value as a type, as an entity that *somehow* still enjoys the local attribute F, because *in principle*, though not in practice, she is still a member of the class 'free agents'.

(B) tries to rationalise the *prima facie* justified request that Mary may still possess some moral value, and hence deserve to be respected, by working on a rather problematic interpretation – as something 'absent-yet-still-present' – of the set of properties necessary to qualify as a rational being. The trouble with (B) is that it becomes soon unclear what it means to have 'somehow' and 'in principle' a certain type of attribute, unless by this we wish to refer either to (B.1) a logical possibility or to (B.2) the object's potentiality, the actual possibility being unavailable by hypothesis.

Consider (B.1). The new criterion – respect any x of which it would be a contradiction to say that it could not qualify as a 'free agent' – becomes too vague,

[17] *Groundwork*, p. 84: 'Hence morality, and humanity insofar as it is capable of morality, is that which alone has dignity'.

because it is also logically possible that a chicken could behave as a free agent.

Consider now (B.2). This is compatible with Kant's ontology. The problem is that, by saying that Mary may still have the attribute F *potentially*, one may mean that

(B.2.a) although born brain-dead, Mary is still morally respectable because she is potentially free by nature, and this is the case because she is a human being. This 'potentially free' feature of her nature cannot be taken away merely because some factor (malformation, accident, disease, drugs, etc.) is in fact preventing her from 'actualising' that potential. The potential can exist 'unactualised' and yet consist of more than mere logical possibility, as her lost freedom is something she could have in a way a chicken never could.

(B.2.a) would allow the Kantian philosopher to solve the axiological problem, if it were not for the fact that, as it stands, it is confronted by two substantial problems.

The first problem is that (B.2.a) begs the question. In the counterexample, Mary does not *happen* to lack the attribute F *momentarily* or just *accidentally*, e.g., because she is under the temporary effect of a drug. If this were the case, (B.2.a) would be correct, but there would be no interesting challenge for the Kantian axiology anyway. Rather, it is assumed that Mary has been *essentially* or *inherently* deprived of her attribute F. She is not and will never be capable of any free behaviour, for example because she is born irreversibly brain-dead. There is no LoA at which Mary enjoys attribute F. In OOP terms, the attribute F has been erased from the description of the information object labelled 'Mary'. A supporter of (B.2.a) could reply that Mary's F attribute cannot be taken away by a contingent event, e.g., a car accident. Yet, this is simply false (second problem). Although essential by hypothesis, a potential attribute is not necessarily a permanent feature of an object and, contrary to what (B.2.a) seems to suggest, it may be removed, even within an Aristotelian ontology. This is an intrinsic feature of the potentiality/actuality distinction, which was originally developed to provide an explanation of change and transformation. A potentiality is an individual's capacity to acquire a certain new state, and this capacity can disappear if the attribute becomes actual, or if the conditions of possibility of the actualisation of the potential attribute are irreversibly removed. If the potentiality of being x is a necessary attribute to qualify as y, this does not mean that whatever is y cannot lose the attribute x, but only that, if y loses x, then y becomes something else different from y. To illustrate the point with a more Aristotelian example: a healthy man has the potentiality of becoming a marathon runner, but once he has become one, this means that he has changed into something else and has lost the potentiality of becoming a marathon runner in favour of the actuality of such potentiality. Likewise, if a healthy man loses his legs, he no longer enjoys the potentiality of becoming a marathon runner. When the potential attribute belongs to the essence of the object, its removal implies the re-categorisation of the individual in a different class, but this is precisely the problem confronting us at the moment: whether a person born brain-dead, who may not count any longer as an ordinarily rational human being, may still be entitled to some moral respect even if the only entities entitled to moral respect are rational beings.

(B.2.a) does not provide a satisfactory answer, but it does contain a valuable point. We have seen that it is not true that, if the attribute F is *practically* not actualisable, F is therefore utterly lost and can be regained only as a logical possibility. Yet this is not the issue addressed by our counterexample, in which the attribute F becomes *essentially* not actualisable. What must be conceded to (B.2.a), however, is that there still is a considerable difference between saying that a chicken *could* be free and that Mary, who is brain-damaged, as a human being *had the potentiality* of being free. The difference would be blurred by a mere reference to a logical possibility, but can be captured by a counterfactual analysis, which leads us to reformulate (B.2.a) thus:

(B.2.b) to claim that Mary is potentially free is to claim that, under normal circumstances, Mary would have not been deprived of F and so she would have been morally respectable.

(B.2.b) is qualitatively (naturalness) and quantitatively (probability) stronger than (B.1). This is obvious if we try to replace 'Mary' with 'chicken' in (B.2.b). (B.2.b) is also more stringent than (A). Nevertheless, it can at most support a 'counterfactual respect', which is still too weak for solving the axiological issue raised by the counterexample. Had Mary had the attribute F (had circumstances been different) she would have been the object of moral respect. This is all one can argue on the ground of (B.2.b). Since Mary does lack the attribute F, however, the counterfactual analysis leaves us with the possibility of being fully justified in showing no moral respect for her. We are not denying that, in another possible world, she would have deserved some respect. We are recalling that, given the present circumstances, she is not 'eligible'.

A Kantian axiology fails to accommodate the counterexample satisfactorily because it is unable to

clarify, in a convincing way, why Mary should be morally respected only on the ground of what she actually lacks by definition and irreversibly in the first place. This discloses a general problem affecting Kant's and other similar approaches. When Kant speaks of moral respect, he has in mind, primarily, perfectly rational agents and only derivatively human beings seen as fallen creatures. In his deontological ethics, a person is morally respectable only in an indirect sense, insofar as she or he implements the properties necessary and sufficient to qualify as a rational being. If the person in question does satisfy such conditions, this hides the fact that, in respecting her, one is really asked to respect not the individual but a class, to which the individual person, however, does not have to belong necessarily. If the person no longer satisfies such conditions, it becomes clear that she was being respected only because she was partaking of the special properties of the class of rational beings.

The solution of the problem requires a shift in perspective. It is hard to see how one could explain and justify any form of respect towards Mary based on some local attribute that, *ex hypothesi*, does not exist. A completely different alternative consists in arguing that Mary still has some form of moral value as an entity that enjoys the inherited attribute L, at a higher LoA. One may no longer express towards Mary exactly the same respect one would have towards a free agent, but one could still feel compelled to respect her at least as a living organism capable of feelings. This alternative looks for the minimal, not the maximal conditions of moral worth, and appears more reasonable and in line with our common sense. It is the one favoured by IE, which now argues for a more decisive step in the same direction.

Once the distinction between local and inherited attributes is introduced, asking what the intrinsic value of x *qua* x is means to ask three different questions.

(1) What is the intrinsic value of x *qua* this specific entity constituted by this specific aggregate of *local and inherited* attributes?

A full answer to (1) can be provided only by combining the two senses in which x has an intrinsic value according to (2) and (3) below. A theory that concentrates only on (1) is a theory of individual moral value, i.e. of the intrinsic value that x possesses in itself as a specific individual, not just as an instantiation of a type. Note that x may be either a single entity (Mary) or a whole class (Women), so the theory does not have to be nominalist.

(2) What is the intrinsic value of x *qua* an entity constituted by its *local* attributes?

Since Kant's concepts of essence, type-token and class membership cut across our concepts of inheritance[18] and aggregate of local and inherited attributes, none of the three questions is exactly the question addressed by Kant, yet (2) is probably the one that comes closest to the Kantian approach, where the local attributes are interpreted as the essential properties of the class of all human beings. A theory that concentrates on (2) may develop a maximalist axiology like Kant's, according to which there is only a restricted selection of local attribute – e.g., intentionality, self-determination, and rationality – that qualify an object as having moral value. Kant is right in arguing that this special object, defined as a 'rational being' or 'good will', is the one that has the highest moral value (dignity) and hence deserves absolute respect. Nonetheless, he is wrong in assuming that this is the only sense in which it is possible to speak of moral worth and respect because one could also ask the following question:

(3) What is the intrinsic value of x *qua* an entity constituted by its *inherited* attributes?

By progressively raising the LoA, one can answer this question by referring to the nature of the entity in question as an information object. We have seen that in the case of the pawn, this is really what matters most. In the case of Mary, the local attributes are far more important, yet this is not a good reason to conclude that, if Mary is reduced to an information object, e.g., as an agent in a virtual context or as an entry in a database, then this information object is devoid of any moral value and can be rightly vandalised, exploited, degraded, or carelessly manipulated irrespectively of any moral concern and constraint. As we shall see, an entity x can be respected at different LoA, including the level at which x is only an information object. Thus, in Mary's case, for example, IE argues that:

(C) if Mary qualifies as a living organism, biocentric ethical concerns apply. Suppose, however, that Mary does not qualify as a living organism any longer. Her corpse still enjoys a degree of intrinsic moral worth because of its nature as an information object and as such it can still exercise a corresponding claim to moral respect.

As Apollo in the last book of the Iliad reminds us, not even Achilles has the moral right to 'outrage the sense-

[18] In OOP, inheritance is the sharing of attributes and operations among classes based on an 'is-a-kind-of', hierarchical relationship between objects. An object is the ancestor object of another, which inherits its attributes and methods. An object may have more than one ancestor (multiple inheritance), may share an ancestor with other objects (shared inheritance) and inheritance may be dynamic (ancestors can be added, deleted or changed through time).

less clay'. Hector's body deserves a minimal level of moral respect.

An axiology that concentrates on (3) can be pluralist or minimalist. A pluralist axiology finds in a selection of inherited attributes – such as intelligence, sensations or biological life – the ontological source of the intrinsic value of an entity, and therefore assigns to a wide variety of entities, namely all those that inherit one or more of these attributes, some moral value and hence a corresponding claim to *a*'s respect. Of course, the moral value in question cannot be absolute, since the theory accepts more than one inherited attribute as comparable, when not competing. It is likely, however, that there may develop a hierarchy of inherited attributes and of priorities in moral standing, and hence a minimalist theory.

A minimalist axiology does not necessarily have to be monist but is not pluralist in the sense that it does not admit that there may be more than one, incomparable and non-equivalent, minimal degree of value. It accepts only one set of inherited attributes as the minimal condition of possibility of intrinsic worth and, as a result, assigns to all the objects that inherit these attributes a corresponding, minimal degree of absolute moral value, in the following sense. Here 'absolute' still means not relative, as in the Kantian 'question'. However, in (2) or more generally in Kant's axiology, the intrinsic value of an entity is incomparable because it is unique, in the sense that there are no other types of moral value, and hence, *a fortiori*, it can not be increased or overridden on the basis of considerations involving other levels or degrees of moral value. On the contrary, here the minimal intrinsic worth of an entity is incomparable because it is unique in the sense that it can be reduced no further, it is necessarily shared, universally, by all entities that may have any intrinsic value at all, and it deserves to be respected by default yet only *ceteris paribus*, that is to say, it can be overridden in view of considerations involving other degrees of moral value at lower LoA. Entities are more or less morally respectable, and we shall see in a moment that an action too is less respectable the more 'entropy' it generates (what 'entropy' means in this context will be explained shortly).

3.3. *'IE's minimalist axiology*

To the question 'What entities have moral value and hence deserve respect?' Two types of answers are now possible, one maximalist or Kantian, and the other minimalist, depending on what we mean by 'moral value'. Minimalist theories of intrinsic worth have tried to identify in various ways the inherited attributes, i.e., the minimal condition of possibility of the lowest possible degree of intrinsic worth, without which an entity becomes intrinsically worthless, and hence deserves no moral respect. Investigations have led researchers to move from more restricted to more inclusive, anthropocentric criteria and then further on towards biocentric criteria. As the most recent stage in this dialectical development, IE maintains that even biocentric analyses of the inherited attributes are still biased and too restricted in scope. As Deep Ecologists argue, inanimate things too can have an intrinsic value. In 1968, Lynn White asked: 'Do people have ethical obligations toward rocks? ... To almost all Americans, still saturated with ideas historically dominant in Christianity ... the question makes no sense at all. If the time comes when to any considerable group of us such a question is no longer ridiculous, we may be on the verge of a change of value structures that will make possible measures to cope with the growing ecologic crisis. One hopes that there is enough time left.' Today, there are geologists' codes of ethics stating, for example, '(9) Don't disfigure rock surfaces with brightly painted numbers, symbols or clusters of core-holes' (http://www.bbc.co.uk/education/rocks/code.shtml) for apparently no other reason than a basic sense of respect for the environment in all its forms. Indeed, even ideal, intangible or intellectual objects can have a minimal degree of moral value, no matter how humble, and so be entitled to some respect. UNESCO recognises this in its protection of 'masterpieces of the oral and intangible heritage of humanity' (http://www.unesco.org/culture/heritage/intangible/). What lies behind these examples is the view that if x can be *p*, then x's nature can be taken into consideration by *a*, and contribute to shaping *a*'s action, no matter how minimally. The minimal criterion for qualifying as an object that as a *p* may rightly claim some degree of respect, is more general than any biocentric reference to the object's attributes as a biological or living entity. What, then, is the most general possible common set of attributes which characterises something as intrinsically valuable and an object of respect, and without which something would rightly be considered intrinsically worthless (not just instrumentally useless or emotionally insignificant) or even positively unworthy and therefore rightly to be disrespected in itself? The least biased and most fundamental solution is to identify the minimal condition of possibility of an entity's least intrinsic worth with its nature as an information object. The information nature of an entity x that may, in principle, act as a patient *p* of a moral action is the lowest threshold of inherited attributes that constitutes its minimal intrinsic worth, which in turn may deserve to be respected by the agent. Alternatively, to put it more concisely, being

an information object *qua* information object is the minimal condition of possibility of moral worth and hence of normative respect. This is the central axiological thesis of any future Information Ethics that will emerge as a Macroethics, to use another typical Kantian phrase.

4. Five objections

We have seen that several arguments support the adoption of IE, yet this does not mean that IE's position is uncontroversial. In this final section, five possible objections are discussed. Answering them will help to make IE more acceptable to those who are not yet convinced of its merits.

4.1. *The need for an ontology*

The first objection concerns the development of a user-oriented information ontology that might help CE to deal with ICT-related moral issues. According to IE, the least (i.e., not further reducible), unconditional (i.e., neither instrumental nor emotional), intrinsic (i.e., belonging to its inherited essence in the OOP sense) and absolute (as clarified above) worth of any entity x, which in principle may fulfil the role of *p*, consists in x's nature *qua* information object and in the very fact of being a possible patient of *a*'s action. On the one hand, the effect of x's role as *p* is completely exhausted in inducing *a*'s respect. On the other hand, understanding in detail how *p*'s moral value, interpreted as an information object x, can contribute to the configuration of *a*'s action in some specific circumstances seems to require an information ontology, namely a theory of the intrinsic attributes of an information object and their integrity, understood as unimpaired and uncorrupted unity and persistence[19] across time. So, what is the objection here?

If the objection is that the need of an ontology affects only IE, it is obviously mistaken. Every Macroethics is based on a specific ontology. Aristotle's, Kant's, Mill's and Environmentalist theories, to mention only four examples that privilege the human or biological nature of *p* as the ground of *p*'s worth, are all based on specific anthropological, psychological, physiological or biological theories.

If the objection is that IE would find developing an information ontology an impossible task, again it is mistaken. One of the main reasons to adopt OOP as a modelling methodology is precisely because it provides the kind of theoretically powerful approach needed to develop successfully an information ontology that is not ethically pre-loaded or biased.

If the objection is that IE needs to provide its own ontology in order to avoid being normatively empty, it is still mistaken. By suggesting that information objects may require respect even if they do not share human or biological properties, IE provides a general frame for moral evaluation, not a list of commandments or detailed prescriptions (compare this to the 'emptiness' of deontological approaches). In Floridi (1999a, 2001a, b) this frame has been built in terms of ethical stewardship of the information environment, the infosphere. It may be worth recalling here the four universal laws against *information entropy* – that is the *destruction*, *pollution* and *depletion* (marked reduction in quantity, content, quality and value) of information objects:

1. information entropy ought not to be caused in the infosphere
2. information entropy ought to be prevented in the infosphere
3. information entropy ought to be removed from the infosphere
4. the infosphere ought to be protected, extended, improved, enriched and enhanced.

Probably the right way of reading the objection is as a reminder that much work still needs to be done to develop IE in full. This is correct and I shall say a bit more about it in section 4.5.

4.2. *How can an information object have 'a good of its own'?*

This objection is based on Taylor (1981) and (1986). Here is the outline:

(i) an entity x is subject to moral respect if and only if x has an intrinsic value
(ii) x has an intrinsic value if and only if
  (ii.a) x 'has a good of its own', that is x can be benefited or harmed; and
  (ii.b) x's flourishing is a good thing
(iii) biological entities (Taylor's 'teleological centres of life'), including non-sentient beings, satisfy (ii.a) and (ii.b)
(iv) it follows that biological entities have an intrinsic value[20] and hence are subject to moral respect
(v) non-biological entities, including information objects, fail to satisfy (ii.a) and therefore (ii.b)

[19] Adapting another OOP concept, persistence can here be defined as the property of any object that outlives the process that generates it.

[20] This is to be understood in perfectionist terms, following Sumner 1996.

(vi) it follows that non-biological entities do not have any intrinsic value and are not subject to moral respect.

The argument is *designed to promote an enlargement* of the domain of entities subject to moral respect, so as to include animals and plants (argument *ad includendum*). It does so by means of condition (ii), which is basically an instruction to adopt a higher LoA than the anthropocentric one. As for the rest of the Kantian axiological frame, the argument strives to keep everything unchanged. In particular, 'intrinsic value' and 'moral respect' are treated as binary phenomena, which can either occur or fail to do so, but have no degrees. Judged in terms of its goal, the argument *ad includendum* may seem reasonable and convincing. All its weakness emerges in (v) and (vi), when the argument *has the effect of excluding* what cannot and should not be subject to moral respect (argument *ad excludendum*).[21]

Regarding (v), anyone endorsing the argument must also accept that a company, a party, or a form of government can all satisfy both (ii.a) and (ii.b), and hence that premise (v) is unjustified. Recall that the argument is meant to show that some non-sentient beings also qualify as morally respectable. What (v) should state is that non-teleological entities fail to satisfy (ii.a) and (ii.b). But now, what are we supposed to conclude about artificial systems like software agents in cyberspace, which are endowed with teleological capacities? From a strictly biocentric perspective, the argument is too permissive.

Regarding (vi), the argument purports to show that anything whose ontological status is either 'higher' or 'lower' than that of a biological entity must inevitably be excluded from moral considerations concerning its intrinsic value and respectability. This is probably wrong. If God exists (and this is a conditional statement), God certainly does qualify as an entity with intrinsic moral value, deserving to be respected. And yet, God cannot be benefited or harmed, at least not in the teleological sense required by the argument. God cannot flourish either, so according to the argument God has no intrinsic value and is not morally respectable. A less stringent but similar case can be made for physical objects like the two giant Buddha statues near Bamiyan. According to the argument, they had no intrinsic value and did not qualify for any degree of moral respect.

One can always bite the bullet, but it seems that something has gone badly wrong with the argument. The fact is that condition (ii) is too strong and rather *ad hoc*. In order to defend the moral respectability of biological entities, it introduces an unnecessarily strict *teleological bias*, which requires x to have the capacities to interact with the environment, to go through a cycle of various developmental states and to pursue goals for its own good. Now, adding a robust dose of teleologism certainly does the trick and (ii) succeeds in enlarging the domain of morally respectable entities, but the approach is too strong and backfires. The enlargement is obtained at the expense of non-biological entities that one may not have any reason to exclude in principle. This is an unreasonable cost once we realise that the argument is at the same time very ecumenical when it comes to a variety of teleological systems, including artificial and social agents.

To fix the argument one needs to invert the relation between x having an intrinsic value and x having a good of its own. If x has a good of its own and x's flourishing is a good thing, then x has an intrinsic value, not vice versa, and certainly not 'if and only if'. This inversion requires a re-consideration of the teleological component in (ii). The proper LoA is not represented by the analysis of what x dynamically strives to be, but by the properties that x has as an entity, even statically. Therefore, the correct terminology to express this point should not be biocentrically biased in the first place. After all, the harm/benefit pair is only a biocentric and teleological kind of the more general pair damage/enhancement. Here is how the argument should be revised:

(i) an entity x is subject to moral respect if and only if x has an intrinsic value
(ii) x has an intrinsic value if and only if (ii.a) x 'has a good of its own', that is x can be enhanced or damaged; and (ii.b) x's existence as x is a good thing
(iii) all existing entities, including information objects, satisfy (ii.a) and (ii.b)
(iv) it follows that all existing entities have some intrinsic value and are subject to some moral respect.

The new version is no longer a biocentric objection against IE but actually an ontocentric argument in its favour. It now fosters moral respect not only for a spider, but also for God (if God exists), for the two Buddha statues, for Mary's corpse and for a database.

Clearing condition (ii) of its biological and teleological bias has at least three consequences. The first two are favourable and show that IE is perfectly coherent with strands of environmental ethics that defend a non-biocentric approach (see for example, Hepburn 1984). First, the original argument implicitly assumes that the true value-bearers are only

[21] For an environmentalist position that accepts the argument *ad includendum* but rejects the argument *ad excludendum* see Rolston III (1985).

biological individuals, not systems (imagine a whole valley taken as an ecosystem), so moral respect is paid to individuals and only derivatively (instrumentally) to systems encompassing them. In the new version, the argument defends the intrinsic value and moral respectability of systems as well as individuals. Second, since we now consider the whole domain of existing entities as being subject to some degree of moral respect, it would be unreasonable to assume that they all qualify for exactly the same kind of absolute respect. A biocentric ethics can still adopt a one-dimensional view of value and respect. Once the Kantian scheme collapses, it must be replaced by a non-absolutist, multidimensional approach. Things have various degrees of intrinsic value and hence demand various levels of moral respect, from the low-level represented by an overridable, disinterested, appreciative and careful attention for the properties of an information object like a customer profile to the high-level, absolute respect for human dignity.

The last consequence is that now the argument is purely *ad includendum*. As such, it may be just too inclusive and turn into a counterargument. The latter could take two forms.

First, one may be reluctant to endorse an 'ontocentric outlook on nature', to adapt Taylor's phrase, because the idea that any entity may enjoy at least a minimal level of moral status may be hard to swallow. Isn't IE unbearably supererogatory? The replies to this protest can be various. One should recall the recurrent qualifications 'overridable' and 'ceteris paribus' and the crucial importance of what have been called 'levels of abstraction' at which a moral situation is analysed. Environmental ethics accepts culling as a moral practice and does not indicate as one's duty the provision of a vegetarian diet to wild carnivores. IE is equally reasonable: fighting information entropy is the general moral law to be followed, not an impossible and ridiculous struggle against thermodynamics, or the ultimate benchmark for any moral evaluation, as if human beings had to be treated as mere numbers. 'Respect information objects for their own sake, if you can', this is the injunction. We need to adopt an ethics of stewardship towards the infosphere; is this really too demanding or unwise? Perhaps we should think twice: is it actually easier to accept the idea that all non-biological entities have no intrinsic value whatsoever? Perhaps we should consider that the ethical game may be more opaque, subtle and difficult to play than humanity has so far wished to acknowledge. Perhaps we could be less pessimistic: human sensitivity has already improved quite radically in the past, and may improve further. Perhaps we should just be cautious: given how fallible we are, it may be better to be too inclusive than discriminative. In each of these answers, one needs to remember that IE is meant to address a context, CE problems, where agents are above all *creators* not just users of the surrounding 'nature', and this new situation brings with it 'divine' responsibilities that may require a special theoretical effort.

Second, one may object that the argument fails to account for the existence of the morally unworthy in general and of evil in particular. Is there anything that actually does not qualify as intrinsically valuable even in the most minimal sense? At the moment, we are missing a revised version of conditions (v) and (vi). This objection is more substantial than the former and deserves its own separate treatment.

4.3. *What happened to evil?*

An axiology that accorded some positive degree of intrinsic worth and hence of moral respectability to literally anything would be of very little value in itself, because, in so doing, it would clearly fail to make sense of a whole sphere of moral facts and the commonly acknowledged presence of worthless and unworthy patients. If IE hopes to be treated as a Macroethics, it must be able to provide a negative axiology as well (for an extended discussion of the concept of 'artificial evil' see Floridi and Sanders 2001a, b).

There seems to be no specific verb in English that fully conveys exactly and only the opposite of 'respect', so let us treat 'irrespect' as meaning simply 'lack of both respect and disrespect'. By 'disrespect' and its cognate words one can then refer to the morally justified and active form of 'anti-respect' towards an 'unrespectable' x, which consists in not causing x, preventing x, removing x, or modifying x so that it is no longer to be disrespected. If something is intrinsically worthless, then it is simply unrespectable, and it is morally indifferent whether *a* respects it as a *p*. If something is intrinsically unworthy, then it is positively to be disrespected inasmuch as it has a certain degree of 'indignity', and not only is it morally wrong if *a* shows respect for it as *p*, but morally right if *a* shows a corresponding degree of disrespect for it, in the technical sense introduced above.

Now, according to IE, something is intrinsically worthless, lacks any moral value and cannot be a centre of moral respect if and only if it does not have even the minimal status of information object. But the only meaningful sense in which it is possible to speak of a 'something' that fails to qualify as an information entity is by speaking of an object that is intrinsically impossible, i.e., a logical contradiction in itself. There are an infinite number of inconsistent objects, but since anything may be predicated of any inconsistent object, there is only one object-type that qualifies as intrinsi-

cally worthless and unrespectable. Let us call it C. C represents the zero degree in our scale of moral worth. It indicates the precise sense in which LoA^i is the highest level of abstraction.

Below C, we find anything that has some possible degree of intrinsic unworthiness and is correspondingly to be disrespected. Information objects can at worst be worthless, never unworthy. Does this mean that the class of unworthy elements is empty? Obviously not. Actions can also be patients and, insofar as they have an information nature as messages, it is possible to apply to them what has been said above about the intrinsic worth of information objects. However, while objects can at worst be intrinsically worthless, messages can also be unworthy and deserve to be disrespected. Messages are not only information objects in themselves but also processes that affect other information objects either positively or negatively. Let us call messages that respect and take adequate care of the welfare of *p* 'positive messages', and messages that do not respect or take adequate care of the welfare of *p* 'negative messages'. Negative messages are unworthy and hence deserve to be disrespected inasmuch as they 'maltreat' their patients. A message that 'maltreats' *p* is a message that does not respect *p*'s information nature, i.e., a message that increases information entropy (in the sense specified above). It is never morally right to show respect for a negative message and *a* has a duty to be comparatively disrespectful towards an unworthy message and to try not to cause, but rather to prevent or remove information entropy.

Messages, but not objects, can rightly deserve to be disrespected as intrinsically unworthy. In more metaphysical language, any process that denies existence, *insofar* as it denies existence, deserves no respect (note that it may still deserve respect for other, overriding reasons), but anything that is, *insofar* as it is, deserves some respect *qua* entity. Ultimate and absolute evil as an object has no moral value at all, and is simply unrespectable because it is an instance of C, in other words it is logically impossible, for it would have to be an object without the status of information object. From an OOP perspective, there can be evil only in terms of negative messages, that is morally bad actions. These are intrinsically more or less to be disrespected, and ought not to be caused, but prevented, removed or modified in such a way as to become no longer evil. The degree of disrespect that *a* ought to show towards a negative message is proportionate to the degree of its unworthiness.

In a possible infosphere in which there were no changes whatsoever there would be no evil. This is the IE version of the Platonic thesis concerning the goodness of being. It clarifies the sense in which something can be extrinsically disrespected: an agent that activates a negative and hence unworthy message is indirectly and contingently deserving of disrespect, but only as a source of M, hence extrinsically.

The extension of the concept of intrinsic worth to any x *qua* information object is now paralleled by the extension of the concept of intrinsic unworthiness to any message *qua* negative process and source of entropy. Messages do not need to be intentional to be unworthy and hence deserving of disrespect, so not every natural process deserves to be respected for the simple fact that it is natural. We live in an improvable infosphere, where moral agents have a duty to apply their ethical stewardship. Their essential capacity to implement positive messages and disrespect negative ones is precisely what makes them the objects with the highest moral value (dignity).

4.4. *Is there a communication problem?*

We saw in (K.2) that, when there is no asymmetry between *a* and *p*, in principle *a* should encounter no conceptual difficulties in recognising *p*'s moral value, and hence in behaving respectfully. Both entities belong to the same class, share the same essential nature and hence the same kind of moral value. The process of communication between *p*'s essence, *p*'s moral value, *a*'s respect for *p*'s moral value and M's adequacy to both *a*'s respect and *p*'s moral value is granted by a principle of reflective respect, whereby the agent can recognise in the patient a member of the same ontological community, a sort of 'alter-ego', and thus easily extend to *p* all the considerations of moral worthiness and requirements of adequate respect that *a* would expect to be rightly applied to *a* itself. This reflective respect is at the root of the Golden rule: *a* can adequately regulate *a*'s actions towards *p* in a way which is already morally successful even if *a* considers only (perhaps just empathically if not rationally) how *a* would like to be treated if *a* were in *p*'s position.

The principle of reflective respect cannot easily be exported when there is an asymmetry in the nature of *a* and *p*. Human self-respect and personal interest in one's own well-being provide some guidelines on how to behave towards *p* that are less and less intuitive the more *p* is ontologically distant from *a*. Simplifying, some reflective respect can still be at work when one is dealing with an animal, but much less so when a tree or a mountain is in question (Leopold 1949, 'Thinking like a Mountain'), and reflective respect becomes truly problematic when the reality one is dealing with is not biological, like a unique database, a patient that, according to IE, can still enjoy some minimal moral value *per se* because of its status as an information object. The risk is to fall into some

form of naïve anthropomorphism. What seems to be required, on *a*'s side, is a 'transpersonal identification', as environmental ethicists like to say. This 'infophilic' or information-friendly attitude is rather more abstract and less spontaneous than commonsensical or empathic feelings. To be able to expand 'the ever-widening circle of ecological consciousness' (Nash 1989) and to appreciate what *a* has in common with *p* when *p* is an information object, *a* should try to transcend *a*'s own particular nature, recognise *a*'s own minimal status as an information object as well, and then extend the respect – which *a* would expect any other agent to pay to *a* as an information object – to any other information object that may be the patient of *a*'s actions. All this requires a change in ethical sensibility. If over-simplified, the perspective can easily be made absurd or ridiculous. Of course, IE does not argue that destroying an old copy of a database is a moral crime in itself. This is just too silly. IE argues that destroying a unique database can be morally evaluated at different levels of abstraction, and that most macroethics work at the low level represented by anthropocentric or biocentric interests, and are perfectly justified in doing so, but that there is also a higher, more minimalist level at which all entities share a lowest common denominator, their nature as information objects, and that this level too can contribute to our ethical understanding. This means that when any other level of analysis is irrelevant, IE's high LoA is still sufficient to provide the agent with some minimal normative perspective. Putnam's twin earth mental experiment can help to clarify the point. Suppose there is a perfect copy of the world; call it twin earth. Suppose that our world and twin earth differ only in this: the unique database is destroyed in our world, but is left intact in twin earth. There is absolutely no other difference. IE accepts the view that twin earth would be a slightly, perhaps very slightly, but still recognisably a better place just because it would be an ontologically richer place. The principle of ontic uniformity grants that the agent *a* acknowledges *a*'s membership of the infosphere and so recognises the inherited attributes *a* shares with all other information components of the infosphere as the ontological ground of their common minimal moral value. The principle of ontic solidarity grants that, by default, the agent *a* will treat all elements of the infosphere, including *a*, as having at least a minimal, overridable moral value *qua* information objects. The moral attitude promoted by IE that emerges from the two principles can be defined, with a play on words, as an 'object-oriented' attitude. In environmental circles this is discussed in terms of ontological or cosmological identification with Being (Fox 1990).

4.5. *Right but irrelevant?*

Someone convinced of the coherence of IE's position could still move the following objection. The problem about IE is not the theory, but its practical irrelevance: IE is too abstract, in the technical sense that its LoA is too high. Recall that IE fully endorses the view that attributing moral worth to information objects provides only a minimalist approach, always overridable in view of moral concerns formulated by other macroethical analyses at lower LoA. Since in everyday life and in ordinary moral decisions there will always be overriding moral concerns, isn't IE completely irrelevant, even if it is right?

The objection raises an important point, as we shall see at the end of the section, but it is largely unjustified. It is simply false that there are always contrasting and overriding ethical concerns (Benn 1998). Ethical theories do not necessarily have to disagree and hence compete with each other in their conclusions. In many cases, they are complementary and can enrich each other. This holds true for IE as well. Moreover, IE has its own special field of application, CE, and other theories seem to have had difficulties in adapting to this new area,[22] so in this sense too there may not be overriding concerns. IE calls our attention to problems that will become increasingly important the more dephysicalised and digitalised our environment becomes. In a society that calls itself 'the information society' it is vital to develop an ethical theory that has the conceptual resources to take into account the status of information objects. IE is an 'architectural' ethics, an ethics addressed not only to the users but also to the creators and designers of the infosphere (Floridi and Sanders 2003). Human beings have evolved as the most successful manipulators and exploiters of nature. Past macroethics have long recognised this fact and tried to cope with its consequences normatively. But human history is also the history of the *ontic divide*, a history of projects and constructions, of detachment from and rejection of the physical world, of replacement of the natural by a human-made environment. *Eco* means 'home', and the infosphere is the new 'home' that is being constructed for future generations. It is the fast-growing environment that human beings, as information objects, are going to share with other non-biological information objects. Clearly, an ethical approach to information ecology is badly needed. IE strive to provide a good, unbiased platform from which to educate not only computer science and ICT students but also the citizens of an information society. The new generations will need a mature sense of ethical respon-

[22] See papers on the uniqueness debate, Johnson (1999), Maner (1999), Floridi and Sanders (2002), Tavani (2000, 2001, 2002).

sibility and stewardship of the whole environment, both biological and informational, to foster responsible care of it rather than despoliation or mere exploitation.

I said that ultimately the objection does raise an important point. IE's goal is to fill an 'ethical vacuum' brought to light by the ICT revolution, to paraphrase Moor (1985). The objection reminds us that IE will prove its value only if its applications bear fruit. This is the work that needs to be done in the near future.[23]

Bibliography

J. Baird Callicott. Animal Liberation: A Triangular Affair. *Environmental Ethics*, 2: 311–338, 1980, rep. with a new Preface in Elliot (1995).

P. Benn. *Ethics*. UCL Press, London, 1998.

J. Benson. *Environmental Ethics*. Routledge, London and New York, 2000.

M. Boman et al. *Conceptual Modelling*. Prentice Hall, London, 1997.

E.W. Dijkstra. Structure of the 'THE'-Multiprogramming System. *Communications of the ACM*, 11.5: 341–346, 1968, http://www.acm.org/classics/mar96/

R. Elliot (ed.). *Environmental Ethics*. Oxford University Press, Oxford, 1995.

L. Floridi. Information Ethics: On the Philosophical Foundation of Computer Ethics. *Ethics and Information Technology*, (1.1): 37–56, 1999a. Preprint from http://www.wolfson.ox.ac.uk/∼floridi/papers.htm

L. Floridi. *Philosophy and Computing*. Routledge, London and New York, 1999b.

L. Floridi (ed.). *Etica & Politica*, special issue on *Computer Ethics*, 2, 1999c, http://www.univ.trieste.it/∼dipfilo/etica_e_politica/1999_2/homepage.html

L. Floridi. Information Ethics: An Environmental Approach to the Digital Divide, paper presented as invited expert to the UNESCO World Commission on the Ethics of Scientific Knowledge and Technology (COMEST), First Meeting of the Sub-Commission on the Ethics of the Information Society (UNESCO, Paris, June 18–19, 2001). Preprint from http://www.wolfson.ox.ac.uk/∼floridi/papers.htm

L. Floridi. Ethics in the Infosphere. *The Philosophers' Magazine*, 6: 18–19, 2001b. Preprint from http://www.wolfson.ox.ac.uk/∼floridi/papers.htm

L. Floridi and J.W. Sanders. Artificial Evil and the Foundation of Computer Ethics. *Ethics and Information Technology*, (3.1), 55–66, 2001a. Preprint from http://www.wolfson.ox.ac.uk/*sim*floridi/papers.htm

L. Floridi and J.W. Sanders. On the Morality of Artificial Agents. *CEPE 2001, Computer Ethics: Philosophical Enquiry* (Lancaster University, 14–16 December, 2001), forthcoming in A. Marturano and L. Introna, editors, *Ethics of Virtualities – Essays on the limits of the bio-power technologies*, to be published for the series Culture Machine. Athlone Press, London, 2001b. Preprint from http://www.wolfson.ox.ac.uk/∼floridi/papers.htm

L. Floridi and J.W. Sanders. Mapping the Foundationalist Debate in Computer Ethics. *Ethics and Information Technology*, (4.1), 1–9, 2002. Preprint from http://www.wolfson.ox.ac.uk/∼floridi/papers.htm

L. Floridi and J.W. Sanders. Internet Ethics: the Constructionist Values of Homo Poieticus', forthcoming in Robert Cavalier, editor, *The Impact of the Internet on Our Moral Lives*. SUNY, New York, 2003.

D.J. Flynn and O. Diaz Fragoso. *Information Modelling – An International Perspective*. Prentice Hall, London, 1996.

W. Fox. *Towards a Transpersonal Ecology – Developing New Foundations for Environmentalism*. Shambhala, Boston, London, 1990.

D.W. Gotterbarn. Computer Ethics: Responsibility Regained, first published in the National Forum, rep. in *Business Legal and Ethical Issues*, 1991, Australian Computer Society August 1993 and in Johnson and Nissenbaum 1995, http://www-cs.etsu-tn.edu/gotterbarn/artpp1.htm,

D.W. Gotterbarn. The Use and Abuse of Computer Ethics, special ethics issue of *The Journal of Systems and Software*, 17.1, 1992, http://www.southernct.edu/organizations/rccs/resources/teaching/teaching_mono/gotterbarn02/gotterbarn02_intro.html

D.W. Gotterbarn. Software Engineering Ethics. In J. Marciniak, editor, *Encyclopedia of Software Engineering*, 2nd ed. Wiley-Interscience, New York, 2001.

R. Hepburn. *Wonder and Other Essays*. Edinburgh University Press, Edinburgh, 1984.

D.G. Johnson. Sorting Out the Uniqueness of Computer-Ethical Issues, in Floridi (1999c), http://www.univ.trieste.it/∼dipfilo/etica_e_politica/1999_2/homepage.html

D.G. Johnson. *Computer Ethics*, 3 rd ed. Prentice Hall, Englewood Cliffs, NJ, 2000.

D.G. Johnson and H. Nissenbaum (editors). *Computers, Ethics, and Social Values*. Prentice Hall, Englewood Cliffs, NJ, 1995.

I. Kant. *Practical Philosophy*. In M.J. Gregor, editor and translator, with introduction by A.W. Wood. Cambridge University Press, Cambridge, 1996.

D. Langford. *Practical Computer Ethics*. McGraw-Hill, London, 1995.

A. Leopold. *The Sand County Almanac*. Oxford University Press, New York, 1949.

W. Maner. Is Computer Ethics Unique? in Floridi (1999b), http://www.univ.trieste.it/∼dipfilo/etica_e_politica/1999_2/homepage.html

N. Medvidovic et al. Formal Modeling of Software Architectures at Multiple Levels of Abstraction. *Proceedings of the California Software Symposium*, pages 28–40. April

[23] A first version of this paper was given at *Computer Ethics: Philosophical Enquiry* (CEPE'98) in Association with the ACM SIG on Computers and Society, London School of Economics and Political Science, London, 14–15 December 1998. A revised version was given at the Third University of Rome, as part of the course for the master's degree, in 1999. On each occasion, I received useful feedback. Richard Keshen and Roger Crisp made a number of valuable suggestions on two other drafts of this paper. Rondo Keele provided constructive criticisms about a previous analysis of the example discussed in 3.2. Finally, I am grateful to the two anonymous referees for their detailed comments and many useful suggestions on how to improve the text, and to Paul Oldfield for his copyediting.

17, 1996, Los Angeles, CA http://www.ics.uci.edu/pub/c2/papers/ADL-CSS96-MTW.ps and http://www.ics.uci.edu/pub/c2/papers/ADL-CSS96-MTW.pdf

J.H. Moor. What is Computer Ethics? *Metaphilosophy*, 16.4, 266–275, 1985.

A. Naess. The Shallow and the Deep, Long-Range Ecology Movement. *Inquiry*, 16, 95–100, 1973.

R.F. Nash. *The Rights of Nature*. The University of Wisconsin Press, Madison, Wisconsin, 1989.

B.G. Norton. The Cultural Approach to Conservation Biology. In D. Western and M.C. Pears, editors, *Conservation in the Twenty-first Century*. Oxford University Press, New York, Oxford, 1989.

D.C. Parnas. On the Criteria to be Used in Decomposing Systems into Modules. *Communications of the ACM*, 15.12, 1053–1058, 1972, http://www.acm.org/classics/may96/

H. Rolston III. Duties to Endangered Species. *BioScience*, 35, 718–726, 1985, rep. in Elliot (1995).

R. Rosenfeld. Can Animals Be Evil?: Kekes' Character-Morality, the Hard Reaction to Evil, and Animals. *Between the Species*, 11.1–2, 33–38, 1995.

W. Rowlands. *The Environmental Crisis – Understanding the Value of Nature*. St. Martin's Press, New York, 2000.

J. Rumbaugh et al. *Object-Oriented Modeling and Design*. Prentice Hall, Englewood Cliffs, NJ, 1991.

L.W. Sumner. *Welfare Happiness and Ethics*. Oxford University Press, New York – Oxford, 1996.

H.T. Tavani. The Uniqueness Debate in Computer Ethics: What Exactly is at Issue, and Why Does It Matter? *Ethics and Information Technology* (4.1), 37–54, 2002.

H.T. Tavani. Computer Ethics: Current Perspectives and Resources. *APA Newsletters on Philosophy and Computers*. Spring, 99.2, 2000, http://www.apa.udel.edu/apa/publications/newsletters/v99n2/computers/feature-tavani.asp

H.T. Tavani. The State of Computer Ethics as A Philosophical Field of Inquiry: Some Contemporary Perspectives, Future Projections, and Current Resources. *Ethics and Information Technology*, 3.2, 97–108, 2001.

P.W. Taylor. The Ethics of Respect for Nature. *Environmental Ethics* 3.3, 1981, rep. in M.E. Zimmerman, editor, *Environmental Philosophy*. Prentice Hall, Englewood Cliffs, NJ, 1993.

P.W. Taylor. *Respect for Nature – A Theory of Environmental Ethics*. Princeton University Press, Princeton NJ.

R. Veryard. *Information Modelling – Practical Guidance*. Prentice Hall, London, 1992.

G.J. Warnock. *The Object of Morality*. Methuen, London, 1971.

Lynn White Jr. The Historical Roots of Our Ecological Crisis. *Science*, 155, 1203–1207, 1967.