

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/228198758>

Does Information Have a Moral Worth in Itself?

Article in *SSRN Electronic Journal* · January 1999

DOI: 10.2139/ssrn.144548

CITATIONS

10

READS

59

1 author:



Luciano Floridi

University of Oxford

460 PUBLICATIONS 14,134 CITATIONS

SEE PROFILE

Some of the authors of this publication are also working on these related projects:



Ethics of Data Publication in the Context of Asylum Claims [View project](#)



BITrum: Interdisciplinary Elucidation of the Information Concepts, Theories, Metaphors and Problems [View project](#)

Does information have a moral worth in itself?*

L. Floridi

Wolfson College, Oxford, OX2 6UD, UK

Tel. +44-(0)1865-274137

Fax +44-(0)1865-274125

Email: luciano.floridi@philosophy.ox.ac.uk

Web page: www.wolfson.ox.ac.uk/~floridi

Abstract

The paper provides an axiological analysis of the concepts of respect for information and of information dignity from the vantage point provided by Information Ethics and the conceptual paradigm of object-oriented analysis (OOA). The general perspective adopted is that of an ontocentric approach to the philosophy of information ethics, according to which the latter is an expansion of environmental ethics towards a less biologically biased concept of a 'centre of ethical worth'. The paper attempts to answer the following question: what is the lowest 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 or even positively unworthy and therefore rightly disrespectable in itself? The thesis developed and defended in the paper is that the minimal condition of possibility of an object's least intrinsic worthiness can be identified with its abstract nature as an information entity. Thus, all entities, interpreted as clusters of information, have a minimal moral worth *qua* information objects (i.e. *qua* information), that deserves to be respected. The principles elaborated in the course of the analysis are those of 'reflective respect' (A's respect towards all members of A's class motivated by A's respect towards A not *qua* individual, but *qua* instantiation of a class of entities), 'ontic uniformity' (A's recognition of A's membership to the class of information entities) and 'ontic solidarity' (A's recognition of any information entity's dignity).

Keywords:

axiology, dignity, entropy, information entity, infosphere, Kant, Kingdom of Ends, Kingdom of Nature, moral value, object-oriented analysis, respect

1. INTRODUCTION: AN OBJECT-ORIENTED MODEL OF MORAL ACTION

According to an information ethics (IE) approach, moral phenomena or facts¹ can be structured as information systems arising out of the dynamic interaction of seven principal components. Since in this paper I wish to analyse a fundamental feature of just one specific component, namely the intrinsic moral value of the patient, in what follows I shall sketch only a general model by way of introduction.

To construct a moral fact or phenomenon, we may start by implementing a simplified model of a moral action. Any action, whether morally loaded or not, has the logical structure of a process that relates a set of one or more active sources (depending on whether we are working within a multi-agent context), the agent A, with a set of one or more destinations, the patient P. To clarify the nature of A and P we can borrow the concept of information 'object' from the object-oriented analysis paradigm (OOA).² A and in case P (see below for the case in which P may be a process) can be discrete, self-contained, encapsulated³ packages containing:

- The appropriate data structures, which constitute the nature of the entity in question (state of the object, its unique identity, and attributes).
- A collection of operations, functions or procedures (methods⁴), which are activated (invoked) by various interactions or stimuli, namely messages received from other objects (message passing) or changes within itself, and correspondingly define (implement) how the object behaves or reacts to them.⁵

In Leibnizian and more metaphysical terms, an object can be treated as a sufficiently permanent (a continuant) information monad, a description of the ultimate primal component of all beings. The moral action itself can be constructed as an information process, i.e. a series of messages (M), generated and activated by A, that brings about a transformation of states directly affecting P, which may variously respond to M with changes and/or other messages, depending on how M is interpreted by P's methods.

A, P and M are the first three information components of our system.⁶ The fourth component is represented by the personal or subjective frame of information within which the agent operates. This shell,⁷ which is really an integral part of A's nature but which it is useful to treat separately, is the information frame that encapsulates the subjective infosphere (see below). It is constituted by internally dynamic and interactive records (modules) of the agent's moral values, prejudices, past patterns of behaviour, attitudes, likes and dislikes, phobias, moral beliefs acquired via education, past ethical evaluations, memories of moral experiences (e.g. of similar situations in which he acted as a witness) or of other moral actions performed in the past, expectations, and so forth. In short, it represents the ethical and epistemic conceptualising interface between the agent and the environment. It is not the task of this paper to specify in detail the very complex and dynamic nature of this subjective information frame. Suffice to say that the agent's shell, although it works as a crystallisation of aspects of the agent's life, is constantly evolving through time, may contain shared or imported elements from other agent's shells, may be only partly accessible to the agent himself and only partly under the control of his will, and yet it contributes substantially to the shaping of his behaviour by

screening him from the direct impact of the information environment, filtering and regulating his access to, and hence highlighting and interpreting the relevant aspects of, the factual information concerning the specific moral situation in space and time in which A is involved. This factual information represents the fifth dynamic component of the system and is perhaps the only element in our model that all ethical theories are ready to recognise as playing an instrumental role in any moral action. 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 his 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, depending on the situation), reasonable (A is intelligent in Mill's sense, i.e. has some capacity to forecast the consequences of his actions) and factually informed. Without some factual information about the moral situation in question (factual information is sufficient; although one may speak of knowledge, understood as a combination of factual information and understanding, this would be too great a requirement) no ethical evaluation is possible on the basis of the agent's intentional responsibility. In other words, an intentional action is morally right or wrong, and it is possible to evaluate it as such and hence to assess the responsibility of its source, only if the latter is free and sufficiently informed (any moral action and corresponding evaluation can take place in a state of only relative scarcity of information) and, vice versa, there is no morality in a state of total determinism and/or ignorance (cf. animal behaviour).⁸ Note that, as a corollary, it follows that some apparent moral dilemmas can be resolved by showing that the action in question no longer qualifies as being subject to an ethical evaluation, either because there is not sufficient information or because there is not sufficient freedom in the system (most importantly, if an agent is forced to commit a morally wrong action, his last resort is to refuse to exist as an agent and commit suicide; if even this option is beyond the agent's power—as one may imagine in some version of the case of Jim and the 20 Indians, for example⁹—then the agent becomes a mere instrument in the hands of his torturer and his moral responsibility vanishes. Of course, the existential problem remains of how the agent can make sense of his violently 'instrumental' treatment, his 'moral death' as it were).

It remains to introduce the last two information components. The agent does or does not implement, and hence variously controls and adjusts, his autonomous and informed behaviour in a dynamic interaction with the general environment in which he is located (e.g. a given culture, society, historical epoch, family situation, financial status, group of individuals, working conditions, etc.). In another context,¹⁰ I have described this constantly evolving universe as the infosphere, the environment constituted by the whole network 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, namely the moral situation, represents the last component of the system. Borrowing a term from robotics, we may define this information microworld as the envelope of the moral action.¹¹

Now that all seven components have been introduced, three comments are in order. Firstly, when the message is a reflective process or a process that also has a

feedback effect, A may be identical with, or treated as, one of the P. I shall come back to this point in a moment. Secondly, it is hardly ever the case that a message affects only a discrete set of well-specified patients P. In what follows, it will be convenient to limit our attention to a simplified dynamic model, and this is why I have used the word 'directly' above, but we also need to keep constantly in mind that a message functions like an information vector or wave, with a given direction and 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 always extends in time and space, in the form of a gradually decreasing continuum. Using another OOA 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 the agent himself, his shell and the factual information) and finally the whole infosphere (think of an abused child who, as an adult, becomes an abuser), which may all be treated as patients. This introduces a third comment on another significant feature of our model: there is no difficulty in always treating A as an object, and in the following analysis of intrinsic worthiness it is also appropriate always to treat P as an object, but we shall see that a negative axiology (a theory of intrinsic unworthiness) requires a more adequate conception of what kind of entity may count as a patient, which may also include a message. To summarise, here is a graphic presentation of the IE-OO model:

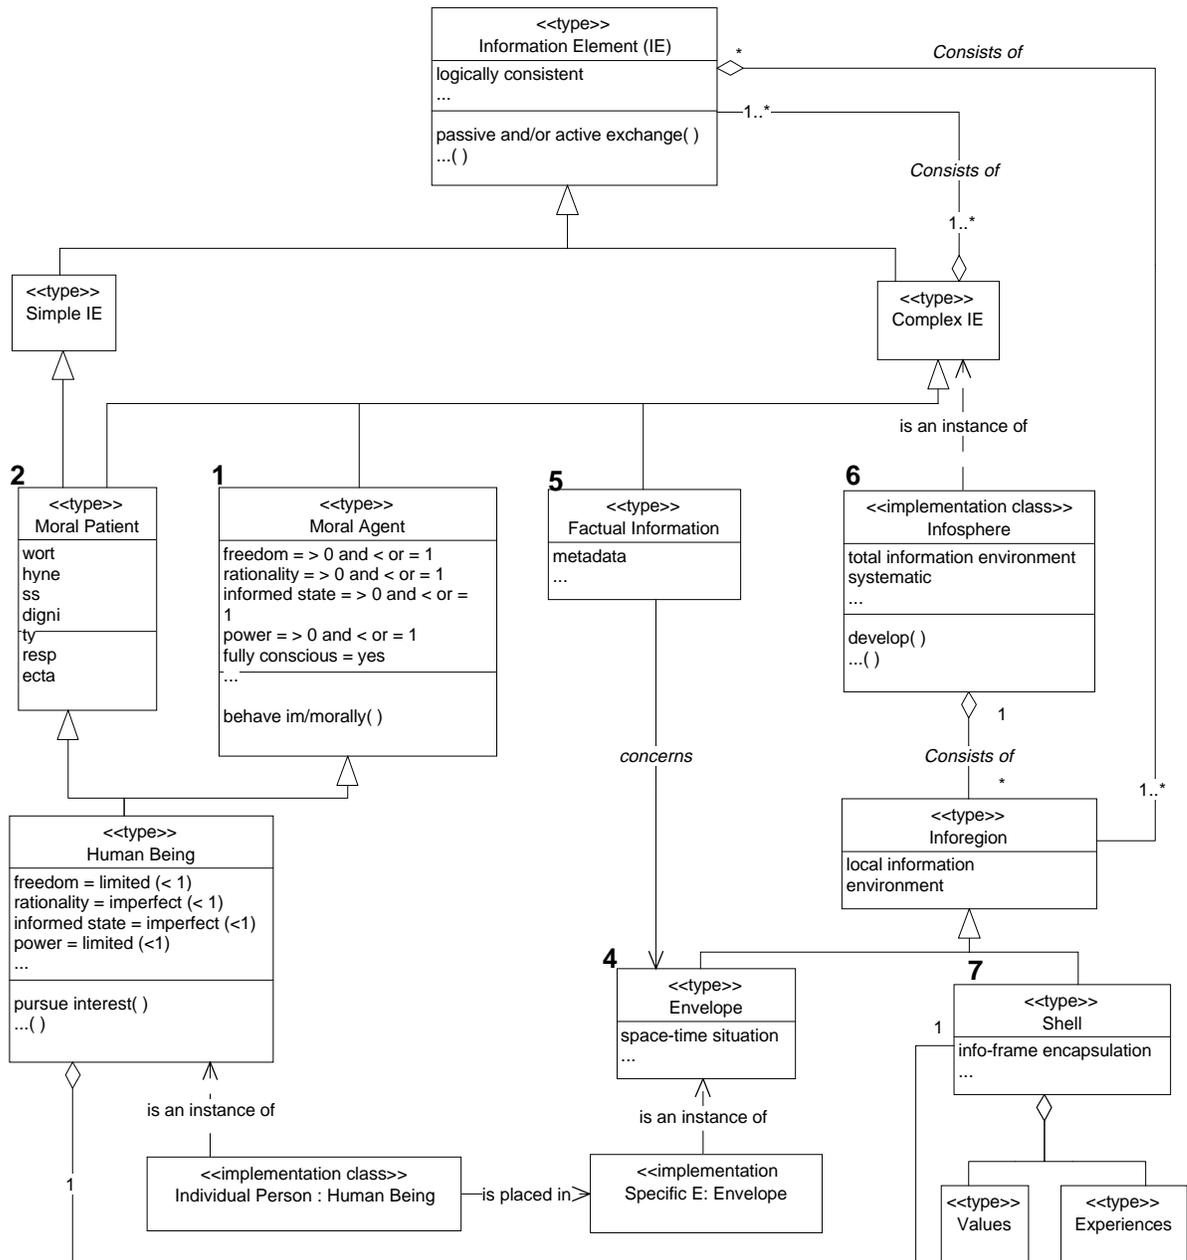


Figure 1. An IE-OO model of moral action: static structure (class) diagram.
 A class diagram is a graphic description of the static structure of a model: it describes its components (especially classes), their internal structures and their mutual relationships. It does not show temporal information.

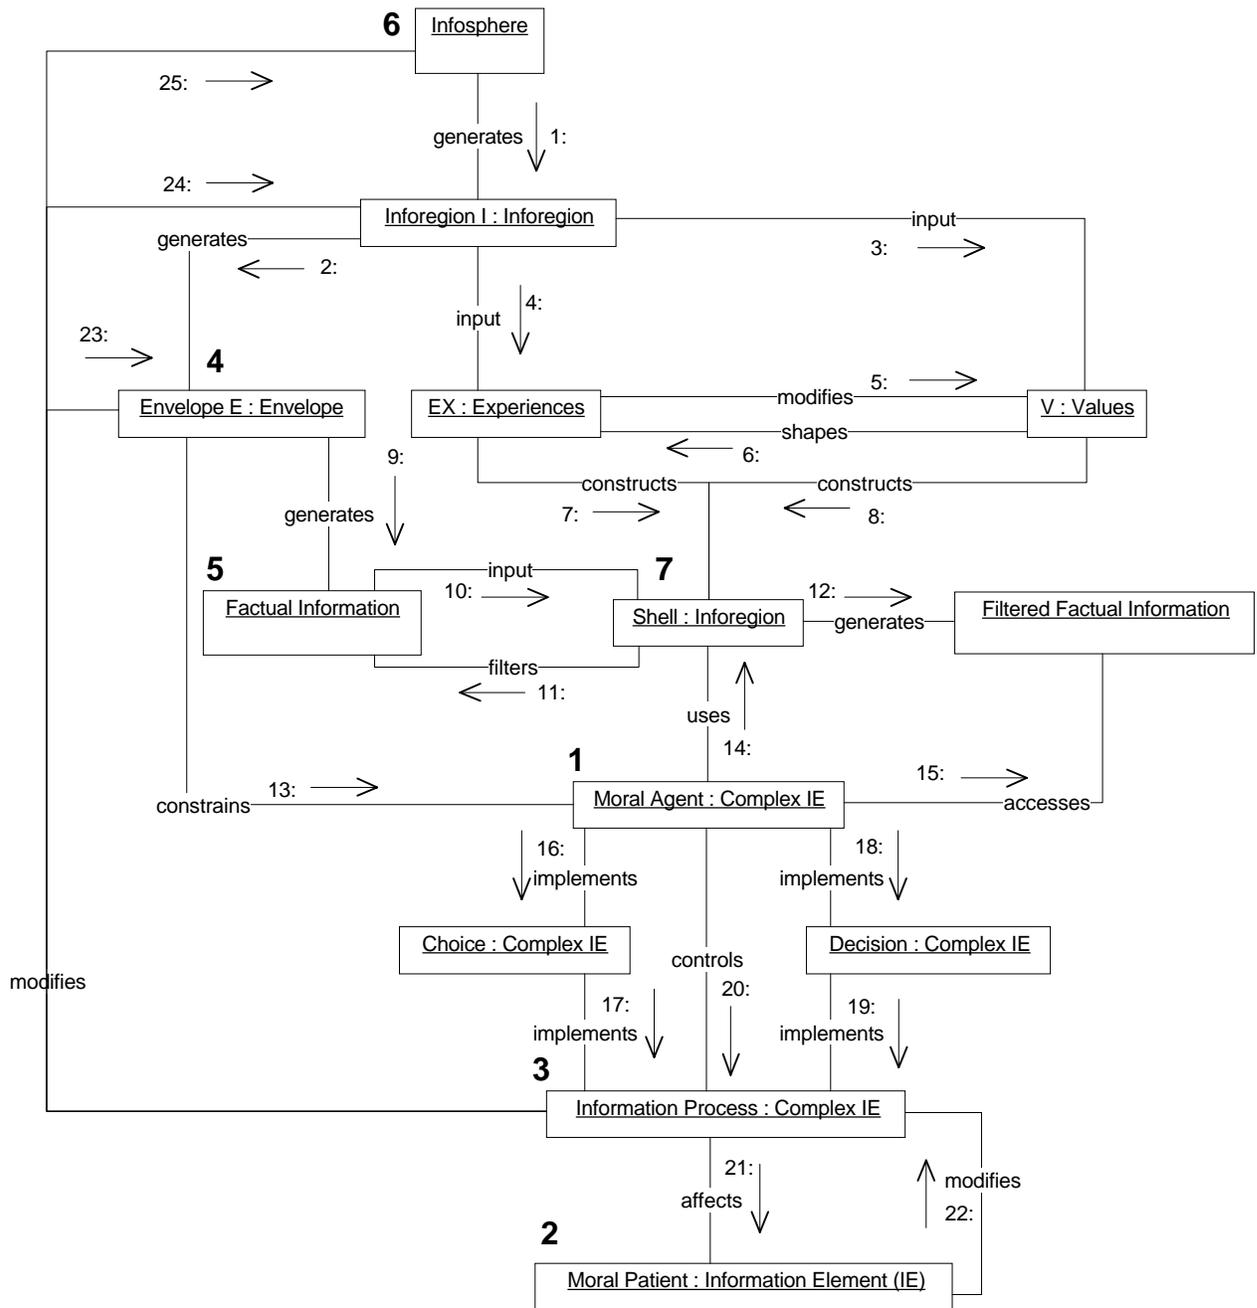


Figure 2. An IE-OO model of moral action: dynamic interaction (collaboration) diagram.

A collaboration diagram is a graphic description of the functional and dynamic interactions within a model: it describes its interaction relations, organised around the components of the system. The sequence of messages exchanged between the objects is logical, not temporal.

2. THE ROLE OF INFORMATION IN ETHICS

We are now in a position to refine our general question, concerning the value and importance of information in the study of moral facts, by investigating more precisely and accurately the specific value and importance that each of the seven components may have within the information system, depending on its role and contribution to the moral action. For example, it would be interesting to develop a detailed typology of the various envelopes (decision problems, moral choices, cases of unintentional disregard or lack of attention, evaluation of third-party actions, control and guide of on-going processes, etc.) within which a course of action with a moral value may develop and, on this ground, to build an information-based theory of moral vagueness (casuistry and the presence of moral latitude in human actions). Likewise, we may investigate whether a moral process has a canonical architecture. What interests us here, however, is a more crucial problem for the foundation of IE. Even a cursory analysis is sufficient to make it clear that, if all components in the system can have only an instrumental value in connection with the moral action, that is, if the set {1, ..., 7} can accurately be described, on the whole, as providing only the (typology of) information input required by any moral action, then IE may, at most, play only an ancillary role with respect to other Macroethics (theoretical, field-independent, applicable ethics such as Deontologism or Consequentialism). This is because the latter 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', the 'dos' and 'don'ts' of a moral agent may be. Thus, the crucial problem for the foundation of IE is that either information itself, in some form or role, can be recognised to have some moral standing—that is an intrinsic moral value over and above its recognised instrumental function, which ought to contribute to determining, normatively, the agent's moral behaviour and hence also have a role in the analysis of legally enforceable rights and ethical duties—or any IE can be at most a Microethics (a practical, field-dependent, applied or professional ethics).

3. AN AXIOLOGICAL ANALYSIS OF INFORMATION

I have already argued elsewhere¹² in favour of the macroethical nature of IE and hence the intrinsic moral value of information. The thesis can be briefly restated here by analysing the seven components already mentioned. It seems clear that 3, 4, 5, 6 and 7 have an immediately instrumental value in connection with the moral action to be performed by the agent. However, the status of the agent and the patient, and the possible modifications in the nature of both, brought about by the information process, cannot be reduced to a mere instrumental factor. These are axiological elements that clearly play a decisive role in the normative assessment of a moral phenomenon, and this is the point on which we now need to concentrate.

At this stage of our analysis, we can already restrict our concern to the fate of P. It would be pointless to pursue two paths, one concerning the agent and the other the patient of the ethical process, because:

- As we shall see in a moment, what we need to consider is whether an information entity in its most general format—not insofar as it is a specific being, e.g. a human being—can have an intrinsic value that ought to regulate a moral action affecting it. Since any object that can act as a moral A can also qualify as a moral P but not vice versa (e.g. animals can be moral P but not moral A) it is better to concentrate on the information nature of an object as a possible P rather than as a possible A.
- As I remarked above, whenever the action in question is found to be either reflexive (e.g. suicide, where A = P) or retroactive (e.g. moral vices acquired through the repetition of actions which are not in themselves necessarily to be deprecated from a moral point of view) the model allows us to extend to A any observation we shall make about P.

By discussing in the most universal and abstract terms the moral worth of an object as a P, not only can we extend our observations to any object insofar as it may also behave as an A, but we can also extrapolate from the specific nature and position taken up in a given envelope by a component of the system, and hence generalise our conclusions so as to include any possible information element that may in principle be affected by the behaviour of a moral agent, and hence qualify as a patient of a moral process, thus including other envelopes, the infosphere itself and the methods, the latter in a way which will become fully clear once we have also developed a negative axiology.

Is there any unconditional (i.e. neither instrumental nor emotional) and intrinsic worthiness, in the nature of an entity that may work as an interpretation of P, that may play a significant role in determining, constraining, guiding or shaping A's moral actions? That is, does an object as a P have a dignity that ought to contribute to the moral configuration of A's message? In so far as a patient object has some dignity, it contributes to the configuration of the message by morally demanding (we shall analyse later what is implied in this 'communication', see K2 below) that A recognises such dignity in a special intentional way, i.e. by having respect for it. A's respect for P's dignity consists both in the appreciation of P's specific worthiness and in the corresponding, uncoerced disposition to treat P appropriately, according to this acknowledged worthiness. Obviously, only objects capable of intentional states can have respect for P's intrinsic value and hence act as moral agents, but are they also the only entities that can have an intrinsic value?

3.1. A Critique of Kantian Axiology

According to a typically Kantian axiology,¹³ there are only two ways in which an entity x may have a value:

- Either x can rightly function only as a means to an end other than itself, and therefore it has an instrumental or emotional value (economic value);
- or x necessarily qualifies also as an end in itself, and therefore it has an intrinsic, moral value *qua* x.

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 value, and hence dignity, for only rationally autonomous agents understood as 'good will' can be the source of moral goodness, thanks to their rational and free actions. This Kantian dichotomy, intrinsic vs. instrumental worthiness:

- K1 Justifies the coextension of (i) the class of entities which have dignity, (ii) the class of entities capable of moral respect, and (iii) the class of entities that deserve to be morally respected. In Kant, the only type of entity that has dignity is the same type of entity that may correctly qualify as a moral patient and that may in principle act as a moral agent.
- K2 Solves the communication problem between A and P: thanks to K1 the agent is acquainted with the dignity of the patient, and hence can respect it, because both entities belong to the same type of class, Kant's 'Kingdom of Ends'.¹⁴ We shall see that, since IE rejects K1, it needs to find a more developed solution to the communication problem than K2.
- K3 Implies that an entity's dignity is a kind of unconditional and incomparable worthiness: 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 turns out to be controversial. Not only does K3 clash with some of our basic intuitions, but K1 fails to take into account two important distinctions.

We commonly assume that different entities may have different degrees of relative worthiness (dogs, for example, may be credited with a lower moral dignity than human beings) that ought at least to constrain the agent's methods, without necessarily having an instrumental worthiness—i.e. a value relative to our feelings, impulses or inclinations, as Kant would phrase it—that may orient them. Likewise, we intuitively accept the view 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. Many would agree, for example, 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 dignity, no matter how humble, that deserves to be respected, although not for instrumental or emotional reasons. More generally, we commonly assume that only rational beings are capable of respect but, contrary to what Kant suggests in K1, we do not treat 'having an absolute dignity', 'being capable of respect' and 'being intrinsically respectable' as symmetric properties. Rational beings are capable of various degrees of respect, to which there seem to correspond various degrees of dignity, and although Kant is right in arguing that 'good will' definitely deserve more respect than other entities, because they are the conditions of possibility of what is morally good, most people would maintain that they only constitute a subclass of the entities that may have a morally significant claim on the agent, as entities subject to some respect.

These intuitions are *prima facie* reasonable, and challenge the Kantian dichotomy. The fundamental problem in Kant's axiology may well have its roots in his ontology, which does not admit of a degree concept of existence, and hence of moral worth, on epistemological grounds (cf. Kant's rejection of the ontological proof and his binary concept of existence), and in the corresponding, conceptual framework represented by the ethical dichotomy between autonomy and heteronomy. The latter seems to lead Kant unduly to restrict the sense of 'relative worth' to meaning only 'contingent worth depending on the agent's interest',¹⁵ so that if x can be rightly and appropriately used only as a means, then x has no absolute worth (and this follows), has only a relative worth (and this also follows), but this can only mean that x's worth has no moral nature whatsoever, because x's worth is now to be interpreted as depending only on the instrumental or emotional interest of the agent, which is a clear *non sequitur*, if we drop the conceptual identity introduced before. The result of the Kantian analysis is that not only do the Kingdom of Ends and the Kingdom of Nature remain largely separate and independent, but the former becomes an information-closed system that is allowed to rule over the latter¹⁶ without there being even the theoretical possibility of the latter providing some indication or constraints.

3.2. An Object-Oriented Approach to Axiology

In order to improve Kant's anthropocentric axiology¹⁷ for our present purposes, it is useful to introduce two distinctions. Let us agree with Kant that there are different ways and degrees in which an entity may have some instrumental worthiness. When the worthiness in question is neither instrumental, nor just emotionally attached to it by a human subject,¹⁸ we 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, or a practice may be respected as a cult. This sense of relative respect is associated with a sense of worthiness which is no longer instrumental or emotional and may be called semantic. Semantic worthiness is still utterly contingent, may be acquired or lost, can be increased as well as reduced. In brief, it is wholly extrinsic. The entity x (e.g. a police officer) may stop being y (e.g. he retires) and hence lose the right to be extrinsically respected (the same person is no longer respected as a police officer), while still remaining essentially the same entity (the specific person in question). In order to capture in full the fact that x has dignity in itself, which belongs to x in all circumstances (necessarily), not under certain conditions, and is not subject to modification unless x ceases to exist, we need to consider the case in which x deserves to be respected not just semantically, as something else, but *qua* x. It is here that we need to depart from the Kantian analysis more radically and introduce a second distinction.

The dignity of an entity is based on its ontology, i.e. what the entity in question is determines what kind of moral worthiness 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. Since Kant adopts a kind of Aristotelian ontology, for him the intrinsic value

of x is indissolubly linked with its essential nature as a certain type of entity. Thus, for Kant an individual, say Lisa, has dignity not as a specific person, but only insofar as she is a token of the general type 'free and rational human being', a member of the 'Kingdom of Ends', in Kantian terminology. In respecting Lisa, the agent is not primarily or directly respecting the specific, unique and contingent individual *qua* herself, but rather the universal type she instantiates.¹⁹ Therefore, the Kantian analysis does not distinguish between two separate senses in which the essential nature of x determines its dignity *qua* x. These two senses can best be clarified by relying on our OOA concepts. The specific nature or essence of an object (whether this is a type or token does not matter now) consists in certain attributes which the object x could not have lacked from the start except by never having come into being as object x, and cannot lose except by ceasing to exist as object x. This essence is a factually indissoluble, yet logically distinguishable combination of (at least some of) its local and inherited attributes. For example, if O2 = Person = descendant object, and O1 = Living Organism = ancestor object, we may say that 'freedom' is an essential and local attribute of Person (i.e. 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. Let us assume now that the object O has an intrinsic value. It is correct to say, with Kant, that O'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 essential attributes.

For example, let us suppose that we respect the object called 'Lisa' 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 Lisa's F attribute and to the latter as Lisa's L attribute, and let us also simplify matters by saying that Lisa inherits L from her ancestor object called 'Animal'. Suppose now that Lisa is radically and definitely deprived of her local attribute F—that is, let us imagine that she becomes *inherently* (see above) incapable of any free and intentional behaviour, e.g. because she is brain-dead—what would be the Kantian position with respect to her dignity? The alternatives seem to be only four. None of them is fully satisfactory.

An extremist solution would be to 'bite the bullet' and argue that

A. Lisa lacks the necessary attribute F, so she can have no justified claim to moral respect. Citizenship to 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 our most elementary moral intuitions. According to (A), for example, we could freely dispose of Lisa's organs without being faced by any moral problem.

If we wish to maintain that, despite the lack of F, Lisa still deserves to be respected, as most people think she should be, we may try to argue, still with Kant, that

B. Lisa still possesses a type-dignity, as an entity that somehow still enjoys the local attribute F, because *in principle*, though not in practice, she still is a member of the class ‘free agents’.

(B) tries to rationalise the *prima facie* justified request that Lisa may still possess some dignity, 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 ‘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 out of the question by hypothesis.

If (B.1) is the case, then the 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, because it is also logically possible that a dog may behave as a free agent.

If (B.2) is the case—as one may argue on the basis of the previous reference to Kant’s Aristotelian ontology—then, by saying that Lisa may still have the attribute F potentially, we may mean that

a. although now brain-dead, Lisa 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 point of 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 dog never could.

(a) would allow our 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.

First, (a) begs the question. In constructing the counterexample, we have not assumed that Lisa *happens* 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, (a) would be correct, but there would be no interesting challenge for the Kantian axiology any way. Rather, we have assumed that Lisa 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 irreversibly brain-dead. The unsatisfactory reply to this first problem, given in (a), is that Lisa’s essential F attribute cannot be taken away by a contingent event, e.g. a car accident. This is simply false (second problem). Although essential by hypothesis, a potential attribute is not a necessarily permanent feature of an object and, contrary to what (a) states, there is no problem in assuming that it may be removed, not even within an Aristotelian ontology (see above the definition of what counts as the essence of an object). This is intrinsic in the conceptual nature 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 either temporarily or definitely, if the attribute becomes actual, or definitely, if the conditions of possibility of the actualisation of the potential attribute are irreversibly removed, as this is the case in

our example. 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 marathoner, but once he has become one, this means that he has changed into something else and has lost the potentiality of becoming a marathoner in favour of the actuality of such potentiality; likewise, if he loses his legs, he no longer enjoys the potentiality of becoming a marathoner. In some cases, 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 facing us at the moment: whether a brain-dead person, who may not count any longer as an ordinary human being, may still be entitled to some moral respect even if the only entities entitled to moral respect are rational beings.

(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* unactualisable, *F* is therefore utterly lost and regainable only as a logical possibility. Yet this is not the issue addressed by our counterexample, in which the attribute *F* becomes *essentially* unactualisable. What must be conceded to (a), however, is that there still is a difference between saying that a dog *could* be free and that Lisa, 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 (a) thus:

b. to say that Lisa is potentially free is to say that, under normal circumstances, Lisa would have not been deprived of *F* and so she would have been morally respectable.

(b) provides a position, with respect to Lisa's possibility of having *F*, that is qualitatively (naturalness) and quantitatively (probability) stronger than (B.1). This is obvious if we try to replace 'Lisa' with 'dog' in (b). (b) is also more stringent than (A), but is it sufficiently strong to solve the axiological problem? The answer is still in the negative. (b) can at most support a 'counterfactual respect', which is still too weak for the purpose of solving the axiological issue raised by our counterexample: had Lisa 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). Since Lisa 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 many possible worlds, she would have deserved some respect, we are stating that, in the present actual world, she is not 'eligible'.

A Kantian axiology fails to accommodate our counterexample satisfactorily because it is unable to clarify, in a convincing way, why Lisa should be morally respected only on the ground of what, in our counterexample, she actually lacks irreversibly. This discloses a general problem affecting the Kantian approach: for Kant, a person is morally respectable only in an indirect sense. 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 'top-down' deontological ethics, a person is respectable only 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 a class to which the individual person does not have to belong necessarily. If the person does not satisfy such conditions, it becomes immediately clear that the reason why she was being respected was only that she was partaking of the ‘holy’ properties of the class of rational beings.

A completely different alternative would be to argue that Lisa still has an object-dignity as an entity that enjoys the inherited attribute L. We may no longer express towards Lisa exactly the same respect we would have towards a free agent, but we would still feel forced to respect her at least as a living organism capable of feelings. This bottom-up alternative looks for the minimal, not the maximal conditions of moral worthiness, appears more reasonable and in line with our common sense, and is the one favoured by IE.

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

1. What is the intrinsic value of O *qua* this specific object 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 O has an intrinsic value according to (2) and (3) below. A theory that concentrates only on (1) is a theory of individual dignity, i.e. of the intrinsic value that O possesses in itself as a specific individual, not just as an instantiation of a type, as in a Kantian axiology. Note that O may be either a single entity (Lisa) or a whole class (Women), so the theory can be either nominalist or universalist.

2. What is the intrinsic value of O *qua* an object constituted by its local attributes?

Since Kant’s concepts of essence, type-token and class membership cut across our concepts of inheritance²⁰ 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 (Kantism), according to which there is only a restricted selection of local attributes, e.g. intentionality, self-determination, and rationality, that qualify an object as having dignity. We have seen that Kant is right in arguing that this special object, defined as a ‘rational being’ or ‘good will’, is the one that has the highest dignity and hence deserves absolute respect. However, we have also remarked that Kant may be wrong in assuming that this is the only sense in which it is possible to speak of moral worthiness and respect. The search for an alternative approach, more in line with our basic intuitions, brings us to the third question:

3. What is the intrinsic value of O *qua* an object constituted by its inherited attributes?

An axiology that concentrates on (3) can be pluralist or minimalist. A pluralist axiology indicates in a selection of inherited attributes—such as intelligence, feelings, biological life or life *tout court*—the ontological source of the

intrinsic worthiness of an entity, and therefore assigns to a wide variety of objects, namely all those that inherit one or more such attributes, some dignity and hence a corresponding claim to A's respect. Of course, the dignity 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 (which does not necessarily have to be monist, quite the contrary, but that is not pluralist in the sense that it does not admit that there may be more than one, incomparable, minimal degree of worthiness) accepts only one set of inherited attributes as the minimal condition of possibility of intrinsic worthiness and, as a result, assigns to all the objects that inherit these attributes a corresponding, minimal degree of absolute dignity, in the following sense. Here 'absolute' still means not relative, as in the Kantian 'question'; but while 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 dignity, and hence, *a fortiori*, it can not be increased or overridden on the basis of considerations involving other levels or degrees of dignity, here the minimal intrinsic worthiness 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 objects that may have any intrinsic value at all, and it deserves to be respected by default, although it can still be overridden in view of considerations involving other levels or degrees of dignity. Objects are more or less morally respectable, and we shall see in a moment that a message too is less respectable the more entropy it generates. To clarify this with a simplified example: if one agrees that potential human life has a lower dignity than actual human life, respect for the former can be overridden by respect for the latter when discussing the moral status of contraception.

4. IE'S MINIMALIST AXIOLOGY

Clearly, to the question 'What entities have absolute dignity 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 'absolute dignity'. Minimalist theories of intrinsic worthiness have tried to identify in various ways the inherited attributes, i.e. the absolutely minimal condition of possibility of the lowest possible degree of intrinsic worthiness, without which an entity becomes intrinsically worthless, and hence deserves no moral respect. Investigations have led researchers to move from more restricted to more flexible, anthropocentric criteria and then further on towards biocentric criteria. As the last stage in this dialectical development, IE maintains that even biocentric analyses of the inherited attributes are still morally biased and too restricted in scope: as Computer Ethics and Deep Ecology convincingly argue, inanimate objects too can have their intrinsic value, and indeed even ideal or intellectual objects seem to have a minimal degree of dignity, no matter how humble, and so are entitled to some respect. If they can be P, then they ought to be taken into respectful consideration by A and contribute to formatting A's messages, no matter how minimally. It

follows that the minimal criterion for qualifying as an object that may rightly claim some degree of respect must be lower than any biocentric reference to the object's attributes as a living entity. What, then, is the lowest 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 disrespectable in itself? The least biased and most fundamental solution is to identify the minimal condition of possibility of an object's least intrinsic worthiness with its abstract nature as an information entity.

The very information nature of an object, that may in principle act as a patient of a moral action, is the lowest threshold of inherited attributes that constitutes its minimal intrinsic worthiness, which in turn deserves to be respected by the agent. Or, to put it more concisely, being an information entity *qua* information entity is the minimal condition of possibility of moral dignity 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. The thesis is controversial, rich in consequences and far from being devoid of difficulties. Here, I shall indicate only three of its problems and then briefly discuss the two most urgent.

According to IE, the least (i.e. not reducible), unconditional (i.e. neither instrumental nor emotional), intrinsic (i.e. belonging to its inherited essence in the OOA sense) and absolute (see above) worthiness of any entity, that in principle may fulfil the role of a patient object P, consists in its nature *qua* information entity and in the very fact of being a possible patient of A's message. On the one hand, the effect of the latter aspect of P's minimal dignity (the effect of its role as patient) is completely exhausted in inducing A's respect. On the other hand, understanding in detail how the dignity of P, as an information entity, can contribute to the configuration of A's message requires an information ontology (a theory of the intrinsic attributes of information, such as integrity understood as unimpaired and uncorrupted unity and persistence²¹ across time) exactly in the same sense in which other minimalist axiologies, that privilege the human or biological nature of P as the ground of P's worthiness, are based on specific anthropological, psychological, physiological or biological theories. Unfortunately, there is no space to develop here even a sketch of an information ontology. Now that we have established that the minimal dignity of any patient consists in its information nature, we need to concentrate on two more immediate and urgent problems: the communication problem and the problem of evil.

4.1. The Communication Problem

We saw in K2 that, when there is no asymmetry between A and P, in principle the agent should encounter no conceptual difficulties in recognising the patient's dignity, and hence in behaving respectfully. Both entities belong to the same class, share the same essential nature and hence the same kind of absolute dignity, and the process of communication between P's essence, P's dignity, A's respect for P's

dignity and M's adequacy to both A's respect and P's dignity is granted by a principle of reflective respect whereby the agent can recognise in the patient an alter-ego, a member of the same ontological community, and thus easily extend to P all the considerations of moral worthiness and requirements of adequate respect that he would expect to be rightly applied to himself. This reflective respect is at the root of the Golden rule: A can adequately regulate his actions towards P in a way which is already morally successful even if he limits himself to considering, perhaps empathically, how he would like to be treated if he 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 little indication of how to behave when P is an animal, a tree, a painting, a grain of sand, someone's diary, a mathematical theory, a social custom, a corpse, a way of greeting and so forth, all 'elements' whose minimal dignity lies in their information status. What seems to be required, on A's side, is a rather more complex and less natural attitude than an empathic effort. A should transcend his own specific nature, recognise his minimal status as an information object himself and hence extend the respect, that he would expect any other agent to pay to himself as an information element, to any other information element that may be the patient of his actions. The former principle of ontic uniformity grants that the agent acknowledges his own membership of the infosphere and so recognises the inherited attributes he shares with all other information components of the infosphere as the ontological ground of their common minimal dignity. The latter principle of ontic solidarity grants that the agent will treat all elements of the infosphere, including himself, as having at least a minimal dignity *qua* information. 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.

4.2. The Problem of Evil: IE's Negative Axiology

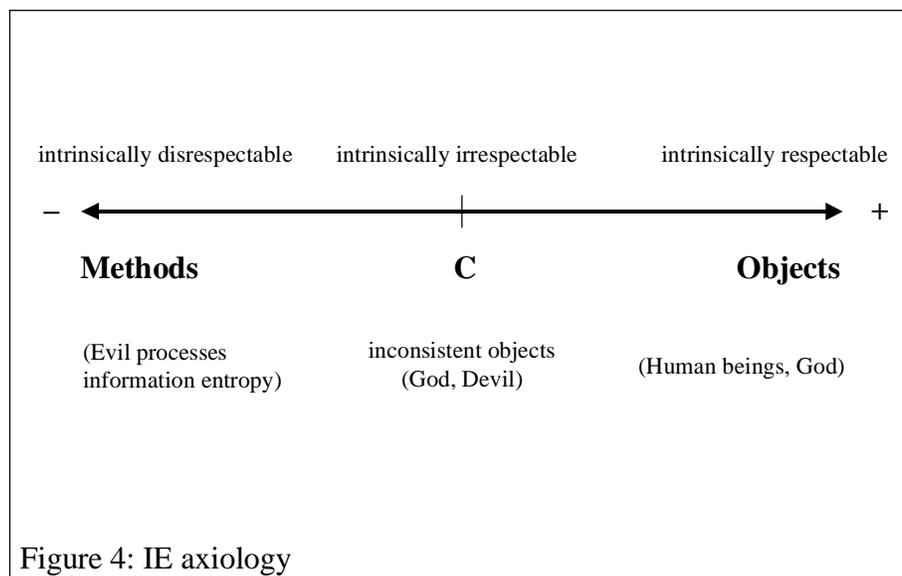
An axiology that accorded some positive degree of intrinsic worthiness and hence of moral respectability to literally anything would be of very little value in itself, for in so doing it would clearly fail to make sense of a whole sphere of moral facts, and the commonly acknowledged presence of worthless and unworthy patients. Thus, it would provide ultimate evidence of its own unreasonableness and unacceptability, not unlike a theory of knowledge that could provide no room for the possibility of human error, or a theory of scientific experiment that made falsification inconceivable. If IE hopes to be treated as a Macroethics, it needs to provide a negative axiology as well.

Let us first introduce some conceptual clarifications. There seems to be no specific verb in English that conveys exactly and only the opposite of 'respect' in full, so let us treat 'irrespect' as simply meaning 'lack of both respect and disrespect'. By 'disrespect' and its cognate words we 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

irrespectable, and it is morally indifferent whether or not 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 it is morally wrong if A shows respect for it as P, but it is morally right if A shows a corresponding degree of disrespect for it, in the technical sense introduced above.

What is intrinsically worthless or unworthy, according to IE? Something is intrinsically worthless, lacks any dignity and cannot be a centre of moral respect if and only if it does not have even the minimal status of information. 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. The temptation of talking in purely negative terms of a similar object must be resisted, since even the apparent lack of explicit information qualifies as the presence of some information, whereas a contradiction posits and erases itself at the same time, thanks to its self-referential nature. There is an infinite number of inconsistent objects, but since each of them contains the affirmation and negation of at least one attribute, and since this inconsistency entails that anything may be predicated of the object in question, following Leibniz's law (identity of indiscernibles) it must be concluded that there is only one object-type that qualifies as intrinsically worthless and irrespectable. Let us call it C. C represents the zero degree in our scale of dignity, and its importance can be clarified by an example. Whether one endorses the ontological argument or not, the following corollary is usually accepted as uncontroversial: given its attributes, namely a collection of all possible perfections to the highest degree, the special object called God either exists necessarily or does not exist necessarily, but cannot exist or fail to exist contingently (either God exists in all possible worlds or in none, but it makes no sense to say that he exists only in some). If God exists necessarily, then he also has the highest degree of dignity and deserves A's highest degree of respect as P. If God does not exist necessarily, then he qualifies as a token of C, he is worthless and it is morally indifferent whether or not A respects him as P. However, it can never be appropriate to say that it is morally wrong for A to show respect for God, or that it is morally right for A to show some degree of disrespect for God. This is IE minimalist version of Pascal's wager.

I have argued above in favour of a degree conception of intrinsic worthiness. The same approach can now be extended to intrinsic unworthiness. Below C, we find anything that has some possible degree of intrinsic unworthiness and is correspondingly to be disrespected. We have just seen that objects can at worst be worthless, never unworthy. Does this mean that the class of unworthy elements is empty? Obviously not. The whole infosphere is made of objects (with their relations, shells, methods and attributes) and messages. Messages can also be patients of A's actions and, insofar as they have an information nature, we can apply to them what has been said above about the intrinsic worthiness of 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 entities in themselves but also processes that affect 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 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 entropy.²² 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 prevent or remove entropy.

5. CONCLUSION: SOME CONSEQUENCES OF IE’S AXIOLOGY

We can now consider a number of consequences of IE’s positive and negative axiology. According to IE, messages, but no 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, but anything that is, insofar as it is, deserves some respect *qua* entity). We saw that the object God is either respectable or irrespectable, but cannot be disrespectable. Likewise, something evil is either irrespectable or disrespectable, but cannot be respectable. Ultimate and absolute evil as an object, e.g. the Devil, however, has no dignity at all, and is simply irrespectable 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 (note the analogy here with the negative version of the ontological argument: the Devil, as the ultimate and absolute hypostatisation of evil nature, necessarily has all the possible limits and defects to the highest degree, including the highest degree of non-existence, namely necessary non-existence, and therefore necessarily does not exist). There can be evil in the infosphere only in terms of negative messages. 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). This also 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. Peter can be extrinsically deserving of disrespect as the perpetrator of a horrible crime, but he cannot be intrinsically deserving of disrespect *qua* the specific entity he necessarily is. No human being can be a Devil.

Finally, the extension of the concept of intrinsic worthiness to any information entity *qua* information 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 (cf. for example the problem of blood transfusion). We live in an improvable infosphere and intentional, rational and free agents have a duty to improve it. Their essential capacity to implement positive messages and disrespect negative ones is precisely what makes them the objects with the highest dignity.

* I am grateful to Richard Keshen, for his thoughtful comments and suggestions on a previous version of this paper, and to Rondo Keele for his valuable criticism of a previous analysis of the example discussed in 3.2.

¹ I am not assuming here the controversial hypothesis that there may be facts or phenomena that naturally qualify as being in themselves moral or ethical. By using the expression 'moral facts or phenomena' I only wish to refer to all those facts or phenomena which ordinarily are evaluated as moral or immoral.

² I hope to be able to argue in the future that OOA provides a powerful modelling tool (conceptual analysis and design) of the real world, and should represent an essential component of Information Ethics methodology in order to describe and understand particular ethical issues within a problem domain. In what follows, I have adopted the standard terminology and the conceptual apparatus provided by James Rumbaugh et al. in *Object-Oriented Modeling and Design* (Englewood Cliffs, NJ: Prentice Hall, 1991) whenever possible. Visual modeling provides a graphical representation of the structure and interrelationships of a system by constructing models of a design. UML (Unified Modeling Language), a *de facto* industry-standard language for specifying, visualising, constructing, and documenting software systems, is the language used in the graphic models in this paper. They have been developed by means of ROSE (Rational Object Oriented Software Engineering), a *de facto* standard visual modeling tool that implements UML. Of course, ROSE has been primarily thought as a tool to allow developers to define and communicate a software architecture. Thus, given the topic and the limits of space, I have relied on a more elementary and simplified formalism for expressing object models. For more information and a free, downloadable, educational copy of ROSE, based on the Rose/C++ 4.0 version, cf. <http://www.rational.com>.

³ In OOA, 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.

⁴ In OOA, 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. An operation may be implemented by more than one method.

⁵ More technically, an object has an identity, a state and a behaviour, three features which are also partly defined by the class it belongs to. A class in OOA is a intentional concept, i.e. it is used as a descriptor or constructor of objects.

⁶ The reader who finds it difficult to conceptualise the agent, the patient and the process as three information components may think of the agent as a mathematical model of a robot (whether it is implemented or not is irrelevant) indistinguishable from a human being, acting in a virtual environment. Even more intuitively, we may represent the system as a realistic fiction, in which characters behave more or less morally. Note that the modelling provided presupposes no commitment to any strong AI thesis.

⁷ The term comes from the operating system architecture vocabulary, not from OOA. It is the portion of the operating system that defines the interface between the operating system and its users.

⁸ Suppose that, through some cruel mechanism, switching on a computer causes enormous suffering in an innocent person imprisoned in the next room. The suffering is utterly unjustified and pointless. If the agent were informed about the causal connection, then by switching on the computer the agent would be committing a morally wrong action. If the agent is not informed, any ethical evaluation appears idle.

⁹ The reader may have seen *Seven* (1995), the psychological thriller directed by David Fincher, starring Morgan Freeman, Brad Pitt, Gwyneth Paltrow and Kevin Spacey. The plot of the film is simple: two homicide detectives hunt for a serial killer who justifies his horrible crimes as absolution for the world's ignorance of the Seven Deadly Sins. In the case of the sin of Lust, the serial killer forces a man to kill a prostitute through a horribly violent sexual act. When the man is in the police station, he recounts and justifies his action by saying that the killer had kept a gun in his mouth during the whole course of action. The police officer, who listens to the confession, does not challenge his justification, implicitly agreeing that, since his own life was seriously at risk, he could not be expected not to torture and kill the prostitute. Obviously, the philosopher's reaction is completely different: the man should have asked the killer to pull the trigger. We are told that he is a married man, we know his life is completely wretched, and one may even go as far as to say that he appears so shocked by his behaviour, which has left him with no integrity or dignity, that he may even commit suicide. We know, but he does not, that the killer would probably not have killed him, since his intention was to kill the prostitute by means of a sexual act and the man's death did not fit in the plan. The man turns out to be an instrument in the hands of the killer, but this is so only because he is unable to make a principled moral choice, and die rather than commit a horrible crime against an innocent. On the other hand, the killer probably expects him to be morally weak since he has already found him in a brothel, willing to use a human being simply as a means, not as an end in herself.

¹⁰ Cf. 'Information Ethics: On the Philosophical Foundation of Computer Ethics', an invited paper given at ETHICOMP98 The Fourth International Conference on Ethical Issues of Information Technology, Erasmus University, The Netherlands, 25 to 27 March 1998, available at <http://www.wolfson.ox.ac.uk/~floridi/ie.htm>; and *Philosophy and Computing* (London: Routledge, 1999).

¹¹ 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.

¹² Cf. 'Information Ethics: On the Philosophical Foundation of Computer Ethics'.

¹³ See I. Kant, *Practical Philosophy*, ed. and trans. by M. J. Gregor with introduction by A. W. Wood (Cambridge: Cambridge University Press, 1996), (all references are to this edition). In the *Groundwork of the Metaphysics of Morals* (henceforth *Groundwork*) Kant writes: '[p. 84] 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.'

¹⁴ See for example *Groundwork*, p. 85.

¹⁵ *Groundwork*, p. 79.

¹⁶ *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. 86-8. 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.

¹⁷ 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, pp. 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 4 anthropocentric examples.

¹⁸ This is Kant's 'fancy price', see Groundwork, p. 84.

¹⁹ Groundwork, p. 84: 'Hence morality, and humanity insofar as it is capable of morality, is that which alone has dignity'.

²⁰ In OOA, inheritance is the sharing of attributes and operations among classes based on a '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).

²¹ Adapting another OOA concept, persistence can here be defined as the property of any object that outlives the process that generates it.

²² Non-being is the absence or negation of any information, or entropy. In IE, entropy is a semantic, not a syntactic concept. As the opposite of information capacity, it indicates the decrease or decay of information leading to absence of form, pattern, differentiation or content in the infosphere. Broadly speaking, entropy is a quantity specifying the amount of disorder, degradation or randomness in a system bearing energy or information. More specifically, in thermodynamics, entropy is a parameter representing the state of randomness or disorder of a physical system at the atomic, ionic, or molecular level: the greater the disorder, the higher the entropy. In a closed system undergoing change, entropy is a measure of the amount of thermal energy unavailable for conversion into mechanical work: the greater the entropy the smaller the quantity of energy available. Thus, a glass of water with an ice cube in it has less entropy than the same glass of water after the ice cube has melted. According to the second law of thermodynamics, during any process the change in entropy of a system and its surroundings is either zero or positive, so the entropy of the universe as a whole inevitably tends towards a maximum. In information theory, entropy is a measure of the noise, or random errors, occurring in the transmission of signals or messages, whereas information is a measure of the probability of a message being selected from the set of all possible messages. Both concepts are therefore purely syntactic: neither information nor entropy refer to the actual meaning, content or interpretation of the message (a string of nonsense symbols and a meaningful sentence may be equivalent with respect to information content), but both quantitative parameters are based only on the presence of uninterpreted difference. The greater the information in a message, the lower its randomness, or noisiness, and hence the smaller its entropy. In IE, we still treat the two concepts of information and entropy as having the same inverse relation, but we are concerned with their semantic value: for example, as the infosphere becomes increasingly meaningful and rich in content, the amount of information increases and entropy decreases, or as entities wear out, entropy increases and the amount of information decreases.

Biography

Luciano Floridi (Laurea, University of Rome 'La Sapienza', M.Phil. and Ph.D. Warwicks., M.A. Oxon.) is Research Fellow, Wolfson College, and Lecturer in Philosophy, Jesus College, University of Oxford. His most recent publications include various articles on the epistemological aspects of information and communication technology. He is the author of *Scepticism and the Foundation of Epistemology - A Study in the Metalogical Fallacies* (Leiden: Brill, 1996) and *Philosophy and Computing* (London: Routledge, 1999). He is the consultant editor for the *Routledge Encyclopedia of Philosophy on CD-ROM* and the editor of P.O. Kristeller's *Iter Italicum on CD-ROM*.