



There's something about Mary: The moral value of things *qua* information objects

Kenneth Einar Himma

Department of Philosophy, Seattle Pacific University, 3307 Third Avenue, West, Seattle, WA 98119, USA
E-mail: himma@spu.edu

Abstract. Luciano Floridi argues that every existing entity is deserving of at least minimal moral respect in virtue of having intrinsic value *qua* “information object.” In this essay, I attempt a comprehensive assessment of this important view as well as the arguments Floridi offers in support of it. I conclude both that the arguments are insufficient and that the thesis itself is substantively implausible from the standpoint of ordinary intuitions.

Key words: ethics, foundations of information ethics, information objects, intrinsic value, moral standing

Introduction

Luciano Floridi grounds his conception of Information Ethics (IE) as a general ethical theory in a very expansive claim about moral standing.¹ On Floridi's view, every entity in the universe can be understood as being an “information object” that is intrinsically valuable and hence deserving of at least minimal moral respect. Every existing entity, whether sentient or non-sentient, living or non-living, natural or artificial, has some minimal moral worth, as he puts it, in virtue of its existence “*qua* information object.” Information Ethics, on his view, represents a departure from traditional general ethical theories (or “macroethical theories”), like utilitarianism and Kantian moral theory, insofar as it entails expanding the moral community to include all information objects as morally valuable entities.

In this essay, I attempt a comprehensive assessment of the thesis that information objects have minimal moral worth and supporting arguments. First, I evaluate the principal lines of argument that Floridi offers in support of this thesis. Second, I assess the plausibility of the thesis against certain ordinary ethical and meta-ethical intuitions. I conclude both that Floridi's arguments provide inadequate support for the thesis and that the thesis itself is substantively problematic judged from the standpoint of ordinary moral intuitions and practices.

Intrinsic value and moral standing

Theoretically prior to the issue of what moral agents are obligated to do is the more general issue of to whom could an obligation be owed. To say that an entity *X* has *moral standing* (i.e., is a moral patient) is, at bottom, simply to say that it is possible for a moral agent to commit a wrong against *X*. Thus, *X* has

moral standing if and only if (1) some moral agent has at least one duty regarding the treatment of *X* and (2) that duty is owed to *X*. The importance of (2) should not be overlooked. Even if my dog lacks moral standing, you still have a duty not to kick my dog; but that duty is owed to *me*, and not to my dog. Such a duty is typically called an “indirect” duty because it immediately concerns the treatment of something other than the subject to whom the duty is owed (though, of course, a duty owed to me not to kick my dog ultimately protects me from the suffering *I* would experience if you kicked my dog). In contrast, a duty that immediately concerns the being to whom the duty is owed is typically called a “direct” duty. To have moral standing and be a moral patient, then, is to be the beneficiary of at least one direct duty.²

² Theorists typically distinguish two general categories of moral standing. A being may, of course, enjoy full-strength moral standing; such a being has a full array of moral rights that are defined by constraints on the outward behavior of agents. My right to life, for example, is constituted in part by certain obligatory constraints on the behavior of other moral agents; in particular, others are constrained from intentionally killing me unless I am culpably posing a threat of death or grievous bodily harm to some other rights-holder. But a being may also have a much weaker form of moral standing in the following sense: a being that is merely *morally considerable* has only a right to have its well-being taken into consideration in the deliberations of moral agents.

While this weaker form of moral standing is principally a constraint on an agent's private deliberations, it probably also entails some oblique behavioral constraints. It is generally thought, for example, that, other things being equal, we can do anything we want to a stone for any reason whatsoever precisely because it lacks all moral standing. It is reasonable to think that moral considerability requires, at the very least, some sort of morally respectable reason for destroying an object and hence precludes at least the act of destroying an object for no reason at all.

¹ Floridi (2002).

While the notion of having moral standing is traditionally associated with having “intrinsic value,” there are two distinct senses of the notion that figure into moral theorizing. The first is concerned with the sort of ends that are characteristically pursued by practically rational agents. In this sense of the phrase, a thing, state, or entity has intrinsic value if and only if practically rational agents typically value it as an end-in-itself; a thing has merely instrumental value, in contrast, if and only if it conduces as a means to some agent’s end. Mill, for example, famously argues that the only thing that people characteristically pursue for its own sake and not for the sake of something else is happiness. Accordingly, all other things, like money or vacations, are instrumentally valuable as a means to securing the intrinsically valuable end of pleasure.³

There are different views about how intrinsic value in this sense figures into moral theorizing. Mill deduced his utilitarianism from the view that the only thing people characteristically value intrinsically is happiness; if happiness is the only thing that people characteristically value for its own sake, then it is the sole ground of moral value. Much more modestly, it is reasonable to think that people have some sort of morally protected interest in what they characteristically value intrinsically – though this tells us nothing about the strength or nature of such an interest (e.g., it tells us nothing about whether it rises to the level of a right).

The second sense of the phrase “intrinsic value” is concerned to identify a class of objects that are entitled to some measure of moral respect. Entities that have intrinsic value in this sense are moral patients deserving of at least minimal respect from practically rational moral agents. Unlike an entity that has only instrumental value, an entity with intrinsic value may not be treated by practically rational agents as just an object to be “used” for purely instrumental purposes. Whereas the appropriate manner for deciding how to treat things with only instrumental value is cost-benefit analysis, things with intrinsic value are entitled to some level of moral consideration in deliberations by moral agents about what to do. Things with intrinsic value in this second

sense count as moral patients with entitlements that must be satisfied.

It is quite natural to think that, as a logical matter, a thing has intrinsic value in the first sense if and only if has it in the second sense, but the relationship between the two notions is more complicated than this. First, the claim that a thing has intrinsic value in the second sense does not imply that it has it in the first sense. While it is true that a thing that has intrinsic value in the second sense is entitled to respect as a moral patient and hence morally *ought not* to be treated as just a means, the claim that a thing has intrinsic value in the first sense makes a weaker claim about what practically rational agents happen to pursue. Mill’s claim is, after all, an empirical one: on his view, when all is said and done, the ultimate motivation for any particular piece of behavior is, on his view, to achieve pleasure. To say that something *S* is morally entitled to respect is to make a *normative* claim about what agents should do; as such, it does not imply that agents actually value *S* as an end-in-itself.

Second, the claim that a thing has intrinsic value in the first sense does not imply that it has it in the second sense. While it is probably true that practically rational agents pursue pleasure for its own sake, it doesn’t follow that pleasure is an object that has moral standing in the sense of being entitled to some measure of respect. Indeed, the idea that one has a moral obligation to respect pleasure that is owed *to pleasure*, if ordinary intuition is correct, involves a category mistake. Neither a general abstract category of mental states nor a particular mental state seem to be the sort of thing that could, as a conceptual matter, be owed an obligation. One can, of course, owe an obligation to an agent to respect *her* pleasure, but that is a fundamentally different claim than the claim that one owes an obligation of respect to *pleasure per se*.

Indeed, it is sometimes thought that the theoretical foundation of utilitarianism is flawed precisely because it falsely equates the two senses. On this line of analysis, utilitarianism explains the moral standing of a person entirely in terms of being an instantiator of the intrinsically valuable states of happiness and pleasure. Since only happiness and pleasure are intrinsically valuable, particular persons are not the direct objects of moral respect; rather, the states of happiness and pleasure are the appropriate subjects of moral respect. Thus, while persons benefit from moral respect, they are not the direct subjects of that attitude. Utilitarianism, then, values persons in only an unacceptable derivative sense: utilitarianism entails that human beings are worthy as respect insofar as they are, so to speak, receptacles of an intrinsically

Footnote 2 (Continued)

While the class of morally respectable reasons might be quite large, there must be at least one reason that is not morally respectable. Otherwise, moral considerability confers no protection whatsoever and is no different, in effect if not in principle, from lacking moral standing entirely.

³ See J.S. Mill, *Utilitarianism*; available from <http://www.utilitarianism.com/mill1.htm>.

valuable state of utility, however this notion is defined.⁴ Though ordinary intuitions might very well be mistaken about all of this, it should nonetheless be clear that the first sense of intrinsic value does not logically entail the second sense.

While it is difficult to fully identify the relationship between the two senses, this much, I think, can safely be hazarded: a state of affairs in which an intrinsically valuable being in the second sense (i.e., morally deserving of respect) achieves *its* intrinsically valuable ends is, other things being equal, morally better than a state of affairs in which an intrinsically valuable being does not achieve its intrinsically valuable ends. This, of course, does not entail either that moral agents are obligated to assist intrinsically valuable beings in achieving their intrinsically valuable ends or that such beings have a right to what they intrinsically value. But it does entail – quite plausibly I think – that respect for an intrinsically valuable thing (in the second sense of having moral standing) requires some concern for whether it realizes its ultimate ends (which presumably would constitute, as Paul Taylor puts the matter, “a good of its own”). Indeed, the idea that one could simultaneously instantiate a moral attitude of respect for a thing with moral standing while being utterly indifferent with respect to whether it achieves what things of its type characteristically regard as intrinsically valuable ends seems conceptually incoherent.

The concept of an entity *qua* information object

To evaluate the IE claim that “all entities... have minimal moral worth *qua* information objects” (1), we need to get clear on the concept of an information object. Floridi explains the concept by way of a helpful analogy:

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 (p. 288).⁵

What constitutes the being of a pawn, then, is defined by the rules of chess that govern what can be done with it. Thus, what Floridi terms the “real pawn” is a “cluster of information.”

As a cluster of information, the most reasonable conception of an information object is as constituting a set, which is also an abstract object. The members of the set, as Floridi explains it, includes a description of the physical properties of the relevant subject (e.g., a pawn) along with a description of the rules and functions that define the subject’s behaviors and operations; as Floridi puts the matter using the formal concepts of object-oriented programming, “data structures (cf. the pawn’s property of being white) and their behavior (programming code, cf. the pawn’s power to capture pieces only by moving diagonally forward) are packaged together as information objects” (p. 288). Thus, the information object defining a pawn is a set that contains all the propositions that describe its characteristics, including physical and behavioral properties.

The claim that “the real pawn is an information object” is probably something of an overstatement. It is true that there is nothing in the nature of chess as a game that requires that it be played on a particular physical surface by pushing pawns around; after all, two people with imagination and memory enough could play a game of chess simply by visualizing a board and pieces. Even so, the fact remains that most people play chess games at particular locations in space–time by pushing particular pieces around the board. The pawn that sits on my chessboard at this moment at K8 is ontologically distinct from every other pawn in the world. It thus has a physical reality that is distinct from the purely abstract existence of

⁴ Other theories of moral standing do not make this mistake (though, of course, they might be flawed for other reasons). A theory that makes the capacity for rationality a sufficient condition for having intrinsic value is not the same as a theory that makes the state of rational deliberation intrinsically valuable. On the first view, a human being is the sort of thing that has intrinsic value in virtue of being able to think rationally. On the second, a human being might not have intrinsic value at all; she might simply be the sort of thing that from time to time realizes states that have intrinsic value.

⁵ Since rules are propositions and hence abstract objects, the term “abstract” more accurately describes the status of information objects. Propositions can, of course, be understood – and the corresponding idea defines a mental state and hence a “mental entity”; but, strictly speaking, propositions are abstract objects in the same way that numbers and sets are abstract objects. While the ontology of such objects presents a great deal of difficulties, they are not plausibly thought of as identical with mental states or as any other kind of mental entity.

the information object that Floridi identifies as “the real pawn.” Though there exist real pawns that have a purely mental existence (such as those in a purely mental game), most real pawns exist in time and space. Since information objects are purely abstract and most pawns are partly physical, most pawns cannot be identical with the information objects that describe them.⁶

Though Floridi often speaks of “persons having intrinsic value *qua* information objects” (which suggests that persons *are* information objects), he concedes, as he must, that biological entities like human beings and moral-biological entities like persons can’t be identical with the information objects they define or pick out. In describing the “information modeling” of a moral agent *a* (i.e., an entity with moral obligations) and a moral patient *p* (i.e., an entity entitled to some moral consideration), Floridi puts the point as follows:

When *a* and *p* are analysed as information objects... they are considered and treated as discrete self-contained, encapsulated 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 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 (p. 288).

To say that a particular person, Mary, can be “modeled” or “analysed” as an information object is, then, to say that there exists a set of propositions that contain a description of the various states, properties, and attributes of Mary over time and a collection of functions defining Mary’s reactions, behaviors, etc. *Qua* human being, Mary is a collection of molecules arranged in a particular way that function in various

⁶ Indeed, it is hard to see how a pawn could be identical with the information object that describes *its* properties and operations. If we conceive of the pawns as nothing more than information objects’ then all of the propositions in the set constituting the relevant information object are propositions that describe *that set*. Such an assumption would, of course, immediately render some of the sentences obviously false (information objects lack spatio-temporal location and hence can’t be moved around). More importantly, it would endow, as a matter of necessity, information objects with a disturbing self-referential quality: the information object defining a pawn is the set of sentences that describe the information object defining a pawn. If this is a coherent interpretation, it certainly cannot be what we have in mind by what might be called a pawn’s information properties.

self-sustaining ways; *qua* information object, the human being Mary is described by the set of propositions and functions that constitute an information object. Strictly speaking, then, the entity we refer to as “Mary” *defines* an abstract information object, but is *not identical* with that object.

In anticipation of the discussion below, it is worth noting that this suggests that what counts as respect for Mary *qua* person (i.e., the human agent herself) might not *necessarily* count as respect for Mary *qua* information object (i.e., the set of propositions describing her). After all, Mary *qua* person is as different from Mary *qua* information object as the entity referred to as “Kenneth Einar Himma” is from the proposition expressed by the sentence “Kenneth Einar Himma enjoys playing and watching basketball.” It should not be surprising that if, as Floridi maintains, information objects deserve respect *qua* information objects, the appropriate attitudes and gestures should be very different from what is owed to human beings, animals, or plants. As we will see, this creates a variety of problems for Floridi’s analysis.⁷

⁷ At this point, however, one can see the problems this creates for a related IE thesis. While Floridi claims that “information objects *qua* information objects can be moral agents” (p. 290), this cannot be literally correct. To say that something is a moral agent is, as a conceptual matter, to say that its behavior is subject to governance by moral standards and hence that it is rightly held accountable for its behavior. But an “information object *qua* information object” is, as an abstract object, incapable, even in principle, of either *behaving* in the relevant sense or being guided by moral standards.

As it turns out, Floridi does not seem to mean this claim literally. He clarifies this thesis immediately after stating it as follows: “This means not just *analyzing* an interpreted *a* as an information object (e.g. *a* = Mary) – this is elementary, as it requires only the adoption of the right *LoA* [i.e., level of abstraction] – but rather showing that *a* can be correctly *interpreted* as an information object” (p. 290). This, however, is not at all what is meant by the sentence, if interpreted literally, “some entity *a* is a moral agent.”

Nevertheless, Floridi takes the preceding sentence to be logically equivalent with the claim “that an artificial agent, like a piece of software, can play the role of a moral agent” (p. 290). While there is nothing in the idea of a moral agent that rules out the possibility of an artificial agent, the claim that artificial agents are possible is logically independent of the claim that a moral agent “can be correctly interpreted as an information object.” There are a wide variety of standards of “correctness” under which Mary can correctly be interpreted as an information object; it seems trivially true, for example, that every entity, possible and actual, defines an information object. But this has nothing to do with whether any particular class of entities satisfies the necessary conditions for being a moral agent.

Floridi's arguments for the intrinsic value of information objects

In evaluating the overall viability of the IE claim that all entities have minimal intrinsic value *qua* information objects, it would be helpful to consider Floridi's reasons for the view. Floridi offers a number of arguments that he believes support his view that every entity has moral standing in virtue of being (or, more accurately, corresponding to) an information object. In this section of the paper, I argue that none of these arguments provides adequate support for his view.

The historical argument

Historically, moral anthropocentrism has been the dominant theory of moral standing. On this view, all and only human beings have intrinsic value and hence have moral standing. While there are a number of ways to ground a moral anthropocentric position, the most common way has been to argue that it is the capacity to reason abstractly (which is, of course, behind the capacity for using language) that makes human beings intrinsically valuable beings (in the second sense) entitled to respect from moral agents. This position, as Floridi notes, was championed by Kant; but it has also been accepted by such diverse theorists as Aristotle and Descartes.

The twentieth century brought a number of increasingly expansive conceptions of moral standing. Moral animocentrists, like Peter Singer and Tom Regan, take the position it is the capacity for conscious experience (particularly of pain) that gives rise to the sort of value that endows a thing with a moral standing that demands respect;⁸ thus, on this view, all and only sentient beings have moral standing. Moral biocentrists, like Paul Taylor and Kenneth Goodpastier, take the position that the property of being alive is sufficient to entitle something to minimal moral respect; thus, on this view, all and only biologically living entities have moral standing. Deep ecologists, like Aldo Leopold, find intrinsic value (in

the second sense) in the natural community as a whole and hence expand the moral community beyond individuals to include ecosystems. Indeed, some theorists go so far as to take the position that artworks have intrinsic value that is entitled to moral respect.

Floridi argues that such developments provide support for further expansion of the moral community to include information objects; according to "the historical argument":

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 (p. 291).

If the historical trend described above is intended as an argument for thinking informational objects have intrinsic value, it falls short in a couple of conspicuous ways. First, one cannot simply assume, as Floridi seems to, that all of these positions are correct. Even the most modest position of the bunch, moral animocentrism, remains a minority position among professional philosophers and laypersons who continue to believe there is nothing wrong with eating factory-farmed animals. The more extreme positions that afford moral standing to ecosystems and art objects are fairly characterized as fringe positions. But to the extent that these positions are false, or even reasonably thought false, it is hard to see how they could provide much support for an even more expansive conception that extends moral standing to information objects.

More importantly, however, this is simply the wrong kind of argument – even if we assume that these theories are correct. If we construe the argument as citing the history of such theorizing as a reason for us to expand the moral community to in-

Footnote 7 (*Continued*)

While it might be true that an extraordinarily sophisticated computer with extensively parallel processing capabilities running the right software might be conscious, as some artificial-intelligence theorists believe, and hence capable of counting as a moral agent, this has nothing to do with its corresponding to some information object. (It should be noted, *contra* the passage above, that a piece of software, by itself, cannot be or "play the role of" a moral agent.) Rocks correspond to information objects but could not play the role of moral agent because they lack other properties.

⁸ See, e.g., Singer (1974). References are to the reprint.

clude information objects, it is problematic because the mere fact that theorists have, as an empirical matter of historical fact, expanded the moral community to include animals, plants, and land doesn't give us any reason, by itself, to think that theorists *ought* to further expand the moral community to include information objects. Alternatively, if we construe the argument as simply deducing the moral standing of information objects from the moral standing of animals, plants, and land, it is problematic because it is not obvious that information objects are morally analogous to animals, plants, and land. While such considerations might justify an open mind towards the possibility of further expansion of the moral community, they fall well short of providing significant support for such a move.

There's something about Mary

Floridi's principal argument for the IE claim is grounded in an analysis of what various macroethical theories would say about what might be called a marginal case of humanity. Floridi asks us to consider the moral standing of a human being, Mary, who is born brain-dead. Despite lacking the various capacities that we usually think give rise to moral standing, we have the strong intuition that Mary deserves some measure of respect – though, as Floridi rightly observes, Mary is entitled to less respect than she would have been entitled to had she been born alive. Given the strength and the stubbornness of this judgment about Mary, Floridi concludes, it is a necessary condition for the adequacy of a general moral theory that it harmonizes with the intuition that Mary has moral standing.

Floridi argues that none of the traditional general theories can be reconciled with this intuition. First, he points out that it cannot be reconciled with the Kantian view that the only property that gives rise to intrinsic value is the capacity for free and rational agency because Mary is dead and hence obviously lacks this capacity.⁹ Second, he points out that it cannot be reconciled with either moral anthropocentrism or moral biocentrism because, by hypothesis, Mary lacks the properties of sentience and life. (It

is clear, moreover, that there is nothing in the land ethic that could explain why Mary deserves respect.)

On Floridi's view, what is needed to explain our intuitions about Mary is an IE approach “that looks for the minimal... conditions of moral worth (p. 296)” and holds that all entities are entitled to respect *qua* information objects. Since the IE approach explains why Mary deserves respect and no other macroethical approach can do so, the case of Mary provides strong support for the claim that all entities have moral standing *qua* information objects. As he puts the matter, “[Mary's] 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” (p. 296).

Accordingly, we can accurately summarize the key steps in Floridi's argument as follows:

1. Mary has moral standing and is entitled to respect.
2. If moral standing is exhausted by Kantian anthropocentrism, moral animocentrism, moral biocentrism, and the land ethic, then it is not the case that Mary has moral standing and is entitled to respect.
3. Therefore, moral standing is not exhausted by Kantian anthropocentrism, moral animocentrism, moral biocentrism, and the land ethic.
4. The only other possible explanation of Mary's moral standing is the IE claim that Mary deserves respect *qua* information object.
5. Therefore, Mary deserves respect *qua* information object.

In the following subsections, I will offer a number of objections to this interesting line of reasoning.

Is it true that Mary's standing is incompatible with anthropocentrism?

Premise 2 is problematic insofar as it is not entirely clear that anthropocentric theories are inconsistent with the claim that Mary is entitled to respect. Many anthropocentric theorists believe, plausibly enough, that a human being who has achieved moral standing at any point in her development is entitled to a certain degree of respect after her demise. For example, if Mary was born with a brain that had the proper hardware in place needed for rational agency (e.g., a fully developed cerebral cortex), then anthropocentrists would claim, presumably including Kant, that Mary became a living person deserving of respect at that point in the pregnancy where the requisite cerebral structures emerged. But if this is correct, then the claim that Mary's corpse deserves respect is no more problematic for an anthropocentric theory than the claim that *any* person's corpse deserves respect:

⁹ Floridi rejects as unsatisfactory four alternatives he believes are open to Kant. The four alternatives are (1) Mary is not deserving of respect because she is not a free and rational agent; (2) Mary deserves respect because “in principle, though not in practice, she is still a member of the class ‘free agents’” (p. 294); (3) Mary deserves respect because it is logically possible that she is a free agent; and (4) Mary deserves respect because she *potentially* has the property of free will.

anthropocentric theories typically claim that respect for persons requires that we respect their corpses, as well as those pre-mortem wishes that are expressed in a will.

Indeed, an anthropocentrist can coherently (and plausibly) claim with minor modification to the theory that even potential persons have moral standing. On this line of analysis, there are two sorts of entities that have moral standing: entities that have the characteristics requisite to personhood (e.g., the capacity for free rational agency) and entities that are likely, given nomological laws, to potentially develop those characteristics. On this view, for example, most living human fetuses qualify, while chickens do not, as having moral standing *at the moment of conception* because they would develop the requisite characteristics under normal circumstances.¹⁰ If having had moral standing at any point during its life (or pre-life in the womb) entitles a thing to some measure of respect after its demise, then Mary would be entitled to respect even if she died during a miscarriage early in the pregnancy. Of course, it is true, on this line of analysis, that she would have a lesser moral standing in such circumstances than she would have had if she developed to term, but that harmonizes nicely with common intuitions: while people typically bury still-born infants, they rarely bury miscarried fetuses.

What sorts of considerations are capable of explaining Mary's standing?

Premise 4 is problematic for two different reasons – one more serious than the other. First, it is simply not clear that there are no alternative explanations for Mary's moral standing beyond the four theories discussed in the argument. There are a number of different religious and quasi-religious views that support the claim that Mary's corpse deserves respect. According to one fairly common view, for example,

¹⁰ This is compatible with the claim that, for example, a genetically defective fetus that is constitutionally incapable of developing the requisite characteristics because it lacks the appropriate genetic sequences needed to drive the development of the brain is not a potential person and hence lacks moral standing. However, such a result is not necessarily intuitively problematic: many people do not have clear intuitions regarding the moral standing of anencephalic infants – much less fetuses that are genetically fated to become anencephalic infants – and many people are prepared to argue that anencephalic infants are not persons (though they may be deserving of some respect). This is why, for example, the issue of whether it would be morally permissible to clone anencephalic individuals as a source of donor organs is being hotly debated by medical ethicists.

Mary's body deserves respect because Mary continues to exist in the form of a disembodied soul that can be harmed by maltreatment of her body. According to another, Mary's body deserves respect, like everything else in the biotic community, because it is the beloved creation of an all-perfect God; on this line of analysis, a being or entity may have intrinsic value because of its relationship to God or to some specific divine purpose.¹¹ While such conceptions are, of course, controversial, they cannot simply be assumed to be false.

More intriguingly, one might take the position that the property of material existence, by itself, suffices to endow a thing with some measure of moral standing. Here it is worth noting that the class of things with minimal moral standing on this view is no broader than the class of things with minimal moral standing on Floridi's view; since everything that has material existence also exists, as it were, *qua* information object, it follows from the IE claim, as Floridi realizes, that all material entities have moral standing.¹² Though neither of these views strike me as especially intuitive, the claim that some extant object *x* has moral standing in virtue of actually being in the physical universe strikes me as considerably more plausible from the standpoint of ordinary intuitions than the claim that *x* has moral standing in virtue of defining an abstract object containing a set of propositions describing it.

The more worrisome problem, however, with Premise 4 is that IE simply cannot adequately explain the kind of things that we do out of respect for Mary's corpse. While, as we saw above, a stillborn child is treated with far more reverence than a miscarried fetus, both are utterly on the same level *qua* information objects. *Qua* information objects, they consist of a set of propositions that describe their various properties, operations, and functions. Mary's nature *qua* information object, then, cannot explain

¹¹ Indeed, one might go so far as to argue, as Spinoza might, that Mary's body deserves respect because it forms part of the totality that constitutes God and is hence divine.

¹² The converse, however, does not hold: merely having a nature as an information object and hence having moral standing on Floridi's view, does not entail the material existence of any particular thing. For example, it is possible that Jane Roe, who died childless, might have had a son named John Roe who lived a particular life with a particular story. Thus, there exists an information object that corresponds to Jane Roe's possible son. It follows, quite implausibly, that John Roe's possible son has some minimal level of moral standing *qua* information object despite the fact that he never had an actual existence in the physical universe. This counterintuitive implication will be discussed in more detail below.

why we show more respect for her if she were stillborn than if she were miscarried. But if this is correct, then the IE claim is no more successful in explaining our moral practices towards Mary than any other problematic theory.

The propositions contained in the information object describing a stillborn child (IO_{SC}) are, of course, quite different from the propositions contained in the information object defining a miscarried fetus (IO_{MF}), but this does not support premise 4. While it is true both that IO_{SC} contains a proposition that attributes to the decedent a fully developed nervous system capable of supporting sentience and rational agency and that IO_{MF} contains the negation of this proposition, this cannot plausibly ground a justification for treating the two decedents differently in an IE view. To begin with, this difference does not in any way bear on the nature of the two decedents *qua* information objects. Moreover, to the extent that such a distinction operates at all to justify treating the two decedents differently, it is far more intuitive to think it is because one of the decedents *instantiates* the property of having the right sort of nervous system while the other one does not than it is to think that it is because the sets defining the two information objects that describe the decedents differ with respect to such a proposition.

Indeed, the problem with Premise 4 runs much deeper. If IE lacks sufficient resources to explain the different levels of respect we show to a stillborn child and a miscarried fetus, it also lacks sufficient resources to explain the different levels of respect we show to a stillborn child and a rock. From the standpoint of the claim that a thing has moral standing in virtue of its nature *qua* information object, there is no difference between Mary *qua* information object and a rock *qua* information object; to put it in roughly Quinean terms, to be *qua* information object is to be a set of propositions that has a certain logical structure. For this reason, the IE approach cannot explain even the intuition that Mary's corpse deserves more respect than a rock.

It is important to note that the problem for IE cannot be resolved by noting, as Floridi does, that other factors (like rationality, sentience, and life) also contribute to moral standing.¹³ While this claim affords Floridi with *other* resources that may help him to explain why we treat Mary with a certain level of

respect, those resources are logically independent of the claim that objects have moral standing in virtue of having an informational nature. As such, they are available to any theorist who rejects the IE approach in favor of an aggregate approach that includes all the other views. Thus, if Floridi has left himself sufficient resources to explain why we treat Mary the way we do, the relevant resources do not include Mary's nature *qua* informational object.

Is accepting the IE claim more intuitively plausible than rejecting the claim that Mary has moral standing?

Perhaps the most telling objection to Floridi's argument, however, has to do with the underlying strategy. To see the problem, let us suppose for the sake of argument that Floridi is correct in thinking that the *only* possible explanation for Mary's having moral standing is that she (or, more accurately, her corpse) has intrinsic value as an information object.¹⁴ One possible response, of course – and the one that is presupposed by the argument – is to accept the IE claim as a means of preserving our commitment to Premise 1's claim that Mary's corpse is entitled to respect. But this is not the only possible response: one could instead reject the claim both the IE claim and Premise 1.

As odd as this may seem at first glance, a deeper look reveals that it is not an entirely unreasonable response. It has been my experience in teaching applied ethics, for example, that students are quite willing to challenge our ethical practices regarding the dead. Students, for example, who subscribe to a physicalist conception of mind and who hence believe that a person's existence is utterly extinguished by physical death will often, upon reflection, give up the view that we owe duties *to the dead* to respect their bodies and pre-mortem wishes; such practices are either explained by indirect duties owed to *other* people (e.g., the decedent's friends and family) or by consequentialist considerations (e.g., respecting a person's wishes regarding post-mortem disposition of their property encourages people to work harder to save for their loved ones). Not everyone, of course, will adopt this view. But most people, quite reasonably, find it intuitively plausible enough at least to take it seriously – suggesting that our gut-level commitment to these practices may not be as durable as Floridi believes.

More importantly, the IE claim has a number of implications that are inconsistent with fairly deeply held moral intuitions about the world. The claim that

¹³ Floridi does not claim that intrinsic value is defined only by having an informational nature; indeed, he endorses only the weaker position that the list of characteristics thought to give rise to moral standing should be expanded to include having a nature *qua* information object.

¹⁴ This would, of course, entail the falsity of a number of ethical and non-ethical views (including theism).

an entity is entitled to minimal respect *qua* information object entails not only that Mary's corpse deserves respect, but also, as Floridi concedes, that every existing thing in the universe deserves respect. But this claim implies that such diverse objects as rocks, boxes, nuclear weapons, and heaps of garbage are all deserving of minimal respect because they define information objects – an implication that is difficult to reconcile with ordinary intuitions and practices.

However, if this is not sufficiently counterintuitive, the IE claim also implies that all possible, but unactualized, objects and states of affairs also deserve minimal moral respect. To say that some unactualized state of affairs *S* is logically possible is to say that there exists a consistent set of propositions that includes the proposition expressed by the sentence "it is true that *S*"; to say that it is causally possible is to say that there is a consistent set of propositions that include the proposition expressed by "it is true that *S*" and the propositions describing the nomological laws that obtain in the actual world. If such abstract objects do not themselves have the right logical structure to constitute information objects of the relevant kind (and it is not entirely clear that they do), then they certainly give rise to or define information objects that do have the correct logical structure. This much is clear about Floridi's view: there is a one-to-one correspondence between things can be described and information objects. Logically possible and causally possible states of affairs, at the very least, correspond to information objects that would, on Floridi's view, be deserving of respect.

From the standpoint of ordinary intuition, practices, and judgment, this is strikingly implausible. While I am not prepared to argue that most people would give up Premise 1 if the only way to consistently hold it was to accept the IE claim, it should be clear that it would not be unreasonable to give up both claims.¹⁵ If I were truly convinced that the logical cost of preserving the belief that Mary (or her corpse) is owed a duty of respect is that I also owe a duty to minimally respect not only rocks and nuclear weapons, but also abstract objects like unactualized possibilities, I would give up Premise 1. My intuition that I owe no such duty to things like this is some-

what stronger than my intuition that a human corpse is owed a duty of respect (though I share this view). Of course, my judgments might well be wrong about this issue, but the point is that Floridi's argument assumes a great deal of highly contestable ground here.

Floridi's reasoning, then, is problematic because what it establishes, at the very most, is that a reader must either accept Premise 1 and the IE claim or reject them both. To succeed in establishing the IE claim, Floridi needs to give some sort of reason to think that the proper response to the dilemma he creates is to accept both claims. Of course, what is needed is not so much an argument for Premise 1; many people, including myself, already accept this – and those that don't remain pretty sympathetic. What is really needed here is an argument that shows that the plausibility of the IE claim on independent intuitive grounds.

What type of argument would make the IE claim plausible?

To get a sense for what is needed to make the IE claim intuitively plausible, it would be helpful to consider the sorts of arguments that make it reasonable to think that human beings, non-human animals, plants, and biotic communities have moral standing. Consider, for example, the Kantian view that rationality and free agency give rise to moral standing. Kant takes the position that what really matters with respect to determining the moral worth of a particular action is the agent's will. The basic idea is that the consequences of any particular act are subject to luck and hence beyond the direct control of the agent. Since moral culpability cannot turn on matters of luck¹⁶ and since, on Kant's view, people have direct control only over their will, it follows that the quality of the will wholly determines the moral quality of the act. Moreover, since a good will cannot be perverted to do evil, only a good will is unqualifiedly good.

Accordingly, having a will that is capable of being good, on Kant's view, is both necessary and sufficient to have moral standing. To have a will capable of being good, one must, of course, have a will and hence the capacity to freely determine what one does.

¹⁵ And part of the reason for this is that this claim would entail the falsity of other claims that ground judgments about respect for corpses (such as classical theism). Classical theists typically ground many of their specific judgments, either directly or indirectly, on the claim that an all-perfect God exists. It is reasonable to hypothesize that a theist would rethink at least some of her specific judgments if convinced that God does not exist.

¹⁶ Consider, for example, two equally drunk drivers who leave a bar at the same time. One barely misses a pedestrian who is able to leap to safety at the very last instant while the other kills a pedestrian who just happens to be there. It strikes many people as counterintuitive to think that there is any difference in moral culpability. After all, the two drivers *did* exactly the same thing; the difference between them is wholly a matter of the behavior of some third person over whom the drivers exercise no direct control.

But one must also have a capacity for rational thought; for the good will consists, on Kant's view, in a commitment to do one's duty according to the moral law and someone who lacks the capacity to understand the moral concepts of right and wrong cannot make a commitment of this kind. Thus, Kant's view ultimately explains moral standing in generally intuitive terms that make reference to the capacities for rationality and free agency.

The argument for the claim that animals have moral standing typically begins from the intuition, almost universal in this culture, that the infliction of gratuitous suffering on any sentient being is wrong; as Bonnie Steinbock puts it, "we all agree that cruelty is wrong, whether perpetrated on a moral or non-moral, rational or non-rational agent."¹⁷ What ultimately explains this intuition, according to Peter Singer, is that animals, unlike rocks, have interests deserving of respect because they are capable of suffering:

The capacity for suffering and enjoyment is a *pre-requisite for having interests at all*, a condition that must be satisfied before we can speak of interests in a meaningful way. It would be nonsense to say that it was not in the interests of a stone to be kicked along the road by a schoolboy. A stone does not have interests because it cannot suffer. Nothing that we can do to it could possibly make any difference to its welfare. A mouse, on the other hand, does have an interest in not being kicked along the road because it will suffer if it is... So the limit of sentience... is the only defensible boundary of concern for the interests of others. To mark this boundary by some other characteristic like intelligence or rationality would be to mark it in an arbitrary manner.¹⁸

What explains the wrongness of cruelty to animals, then, is that animals have intrinsic value in virtue of having a capacity – i.e., sentience – that is both necessary and sufficient for having interests and hence for having moral standing.

In justifying moral biocentrism, Paul Taylor accepts the relevance of having interests to determinations of intrinsic value, but denies that the capacity to suffer is a necessary prerequisite for having interests. On Taylor's altogether intuitive view, it makes sense to think of living things as having interests in virtue of having "goods of their own":

Every organism... has a good of its own which moral agents can further or damage by their ac-

tions. To say that an entity has a good of its own is to say that, without reference to any *other* entity, it can be benefited or harmed. One can act in its overall interest or contrary to its overall interest, and environmental conditions can be good for it (advantageous to it) or bad for it (disadvantageous to it). What is good for an entity is what "does it good" in the sense of enhancing or preserving its life and well-being. What is bad for an entity is something that is detrimental to its life and well-being.¹⁹

Insofar as living beings necessarily have a good of its own, it follows, on Taylor's view, that "rational, autonomous agents [ought to] regard such entities as possessing inherent worth."

The land ethic is grounded in the common intuition that a human community is something that deserves moral respect for its own sake. Since local and global ecosystems exhibit the same interdependence that, as a conceptual matter, *constitutes* a particular group of human beings as a "community," it follows that these ecosystems are also communities and hence possess the same characteristics that give rise to moral standing in human communities. Just as a human community deserves respect as such, so too the biotic community deserves respect as such; as Leopold puts it, "a land ethic changes the role of *Homo sapiens* from conqueror of the land-community to plain member and citizen of it... [and] implies respect for his fellow-members, and also respect for the community as such."²⁰

It is important to note that the arguments for these various positions, whether correctly or incorrectly, attempt to justify them by grounding them in common moral intuitions that are related to rationality, suffering, life, and community. Each one attempts to argue that we are already committed in virtue of our common intuitions and moral practices to regarding rationality, the capacity to suffer, being alive, and being a community as having intrinsic value deserving of moral respect. Insofar as we believe cruelty to animals is wrong, for example, it must be because we regard animals as having intrinsic value in virtue of having interests that can be harmed. Insofar as we believe being vulnerable to harm gives rise to intrinsic value, it follows that living things have intrinsic value. Insofar as we believe communities are intrinsically valuable, the biotic community as a whole is intrinsically valuable.

Thus, while Floridi disparages these positions as "biased" (p. 299), this characterization is misleading

¹⁷ Steinbock (1978).

¹⁸ Singer (1977).

¹⁹ Taylor (1981, p. 199, emphasis added).

²⁰ See, e.g. Leopold (1977).

at best. The anthropocentric, animocentric, biocentric, and land ethics positions are explicitly grounded in the exactly the right kind of considerations needed to persuade moral agents: they appeal to deeper commitments implicit in our ordinary intuitions and moral practices. That they are intended to appeal to *our* commitments is not an indicator of bias; it is the hallmark of a good argument. No argument – not even a mathematical argument – can pick itself up by its own bootstraps; since reasoning designed to persuade finite agents must itself be finite, any persuasive argument must begin with premises that are regarded by the intended audience as plausible or correct.²¹

There is a sense, of course, in which Floridi attempts to ground the IE claim in a common moral intuition; after all, he attempts to ground it in our commitment to the claim that Mary's corpse deserves respect. But notice that the logical distance between the intuitive claim that Mary's corpse deserves respect and the intended conclusion that objects have intrinsic value in virtue of being information objects is vastly greater than, for example, the distance between the intuitive claim that cruelty to animals is wrong and the intended conclusion that beings have interests that deserve respect in virtue of being susceptible to cruelty. The connection between the two latter claims, unlike the connection between the two former claims, approximates that of premise to corollary.

This is important because the point of the other arguments is to identify some common moral commitment that can be *directly* applied to show that the relevant entity deserves respect. The intuitive claims that (1) cruelty to animals is wrong, (2) having a good of its own gives rise to intrinsic value; (3) community is intrinsically valuable result immediately in the respective conclusions (1) animals have intrinsic value; (2) living things have intrinsic value; and (3) biotic communities have intrinsic value. In contrast, the intuition that Mary deserves respect does not in the same way warrant an inference that information objects deserve respect.

²¹ Mathematical theories typically assume a set of axioms and a more or less formal set of inference rules. It sometimes turns out that these assumptions are problematic. For example, Euclid's Parallel Postulate (i.e., given a line and a point off the line, there is one and only one line passing through that point parallel to the given line) is inconsistent with Einstein's theory of relativity. Most people, however, have a very difficult time imagining how Euclid's Parallel Postulate could falsely model our physical reality.

Is there an intuitive argument for the IE claim?

As it turns out, there is one argument for the IE claim that is grounded in the right kind of intuitive claim. We pursue knowledge, the argument begins, for two qualitatively different reasons. There are some things we want to know – some information that we seek – because it has instrumental value as a means for solving certain problems; most medical and technological knowledge, for example, falls in this category. But there are many things we want to know, i.e., pieces of information we seek, without any regard for whether they have such benefits. For example, mathematicians have long regarded the question of whether Fermat's Last Theorem is true as one of the most important problems in mathematics – despite the fact that no one thinks it is even remotely likely that the answer to this question bears on any practical problem we face. Indeed, even when information has instrumental value, we tend to regard it as important for its own sake: knowing how the body works, for example, is valuable not only as a means to enable us to treat diseases, but also as an end-in-itself. Insofar as we seek knowledge for its own sake, it follows, according to this argument, that information is intrinsically valuable and hence deserving of respect.

As plausible as such an argument may seem,²² however, it will not support Floridi's views about information objects. At the very most, it will support an attribution of moral standing to the individual propositions that express information – i.e., to particular units of information – and not to entities *qua* information objects as Floridi defines that notion. The claim that we value knowledge for its own sake shows, at most, that we regard information units (i.e., true propositions) as having intrinsic value. But it does not show that we value rocks, nuclear weapons, or even people *qua* information objects; that is, the claim that we regard knowledge as having intrinsic value does not show that we value rocks, nuclear weapons, or people as having intrinsic value insofar as they define certain data structures that describe them. The intuitive claim that information is intrinsically valuable is very different from the unintuitive claim that every entity has intrinsic value *qua* information object.

But the problem with this argument runs even deeper: it is not even clear that the claim that we value knowledge for its own sake implies even the claim that information has moral standing. As will be recalled, there are two different senses of "intrinsic value." To say that we value information for its own

²² I have elsewhere argued that it is false. See Himma (2004) (a) and (b).

sake is to say that it is intrinsically good for beings like us that have moral standing. But while this claim may imply that it is morally good, other things being equal, that beings like us have information, it doesn't imply that information has moral standing in the sense of being owed obligations. Happiness is valued by us for its own sake, but this simply means that it is morally good for moral patients who intrinsically

²³ Floridi offers a second argument for his claim that information objects have intrinsic value, but it is vulnerable to the same sort of objection. In Section 4.2 of the paper, Floridi challenges Paul Taylor's reasoning for the claim that only biological entities can have intrinsic value. First, he points out that if God exists, then God is a non-biological entity that has intrinsic value; thus, the conclusion of the argument is problematic. Second, he gives examples of non-biological entities that he believes have goods of their own; on Floridi's view, companies, parties, and governments can have goods of their own. (The problem with this claim, however, is that the good of such entities is most plausibly explained in terms of the goods of various parties: shareholders, members, and citizens.)

Floridi, then, reconfigures Taylor's reasoning in the following way: "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... The proper LoA [level of abstraction] is not represented by the analysis of what x dynamically strives to be, but by the properties that x has 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
 - (a) x 'has a good of its own', that is x can be enhanced or damaged; and
 - (b) x 's existence as x is a good thing;
- (iii) all existing entities, including information objects, satisfy (a) and (b);
- (iv) it follows that all existing entities have some intrinsic value and are subject to moral respect" (p. 299).

This argument is vulnerable to exactly the same objection as the argument from the example of Mary. The strategy here is, in essence, to argue, first, that Taylor's reasoning is not strong enough to capture the value of God and certain artificial entities made up of persons; second, that allowing intrinsic value to information objects can do so; and, finally, to conclude that all existing entities have some intrinsic value. But notice that the key premise in the argument – premise (iii) – receives no intuitive defense whatsoever.

value happiness to have happiness; it does not imply the very counterintuitive claim that moral agents owe an obligation of respect *to* happiness.

In any event, it seems that, to meet his burden of persuasion, Floridi needs to supplement his analysis with an argument that is very different in character from the arguments he supplies. In particular, he needs to try to ground the idea that information objects have moral standing in ordinary intuitions about either the treatment of information or the treatment of people, animals, plants, and other entities *as* entities that are representable by the sorts of data structures that constitute information objects. This, as should be evident, presents a very different challenge from those that Floridi's OOP-influenced analysis attempts to meet.²³

Information objects as appropriate objects of respect

To say that information objects have intrinsic value (in the first sense described above) and are hence deserving of respect is only part of the job that must be done in establishing IE as a separate "macroethical" theory. As a general normative approach, macroethical theories define systems of ethical rules, principles, or considerations that purport to guide the behavior of moral agents. If IE is to stand on its own as a general macroethical theory, it must also ground a system of normative standards in the foundational IE claim that information objects deserve moral respect.

Floridi attempts to meet this burden, plausibly enough, by identifying a number of ethical principles that purport to govern the behavior of agents in the "infosphere"; as he describes the matter:

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 and 2001b this frame has been built in terms of ethical stewardship of the information environment, the infosphere (p. 298).

Footnote 23 (*Continued*)

Floridi simply asserts that "the harm/benefit pair is only a biocentric and teleological kind of the more general pair damage/enhancement," insinuating – rather than arguing – that there is no morally significant difference between the various notions; thus, the argument assumes, rather than shows, that information objects have a good of their own. What is needed here is an intuitively grounded reason to think that the existence of information objects is "a good thing" (i.e., that information objects satisfy (b)).

Thus, Floridi posits four “universal laws against information entropy”: (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; and (4) the infosphere ought to be protected, extended, improved, enriched, and enhanced. Taken together, these laws express the basic normative claim that, other things being equal, preservation of information is morally good and destruction of information is morally bad.

In one respect, these laws are intuitively plausible. As discussed above, we regard information as being the sort of thing that has tremendous value to us. Given the importance of information to us, it is intuitively reasonable to think that there would be ethical standards, however general or modest, that obligate us to protect and preserve information. If it is morally good that beings with moral standing have what they value intrinsically and we value information intrinsically, there must be moral standards that express that, other things being equal, we ought to have information. If protecting against information entropy amounts to promoting human acquisition of information, then the four general standards described above harmonize well with our moral intuitions about the importance of information and the associated practices.

The problem for IE, however, is that the universal laws are more firmly grounded in other claims than they are in the IE claim. To begin with, it is important to note that the intuitive considerations adduced in the preceding paragraph ground the universal laws in the weaker claim that units of information – as opposed to information objects – have intrinsic value in the same sense that happiness has intrinsic value. As we have seen, the claim that happiness has intrinsic value implies that, other things being equal, a state of affairs with more happiness is morally preferable to a state of affairs with less happiness, but this does not entail or presuppose the claim that happiness itself has moral standing deserving of respect. Accordingly, the universal laws of the infosphere can be grounded in an intuitive way in the claim that information has intrinsic value (in the first sense of being characteristically valued as an end-in-itself by beings like us) and hence that it is a morally good thing that beings like us with moral standing have information – without assuming either that information or information objects have moral standing.

As it turns out, the universal laws of the infosphere can, as a logical matter, be firmly grounded in the more modest claim that information has instrumental value. After all, property is commonly regarded as having only instrumental value as a means

for satisfying other more important desires. But the close connection between property and the satisfaction of these desires is sufficient to warrant strong moral protection of property. It is not only true, for example, that we have a duty to respect the property of other persons; it is also true that we sometimes have duties not to squander or waste our own property – as it is commonly thought that it is morally bad to be wasteful. But this means that, assuming the universal laws of the infosphere can be grounded in the IE claim, it is false that they the IE claim is logically necessary to ground these laws.

As it turns out, it is not clear how these laws can be plausibly grounded in the IE claim that information objects have moral standing. Part of the problem here is that the notion of an information object, as Floridi defines it, is an abstract object having a very specific structure. Information objects, as will be recalled, are “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)... packaged together” (p.288). To claim that information objects have moral standing and are entitled to respect does not imply that information ought to be protected and preserved; at most, it implies that *information objects* ought to be protected and preserved.

Of course, one might reasonably think that standards protecting information are easily derived from the claim that information objects ought to be protected. On this line of analysis, information objects are sets that contain propositions – the basic units of information. Since the existence and identity of any set is defined, according to the set theoretic axiom of extensionality, by the existence and identity of its members, it follows that protecting the existence of information objects requires protecting the existence of the information units (i.e., the propositions) that are its members.

This reasoning is problematic, however, because it supports the wrong kind of laws – or, put more precisely, the wrong construction of the universal laws above. Construed as protecting *the existence of information objects*, the universal laws are impossible to violate. Information objects are abstract objects (i.e., sets having a particular logical structure) that are composed entirely of other abstract objects (i.e., other sets and propositions). The “existence” of these objects does not depend on the existence of any physical object. If it makes sense, as many philosophers of mathematics and logic believe, to talk of the existence of abstract objects, they exist in an abstract “logical space” that is ontologically independent from the spatio-temporal universe we inhabit.

Strictly speaking, there is nothing we can do to causally affect the existence of abstract objects in logical space. Propositions regarding, for example, what is needed to cure cancer exist in logical space whether anyone ever utters a sentence that correctly expresses any of those propositions. Every possible world, understood as maximally consistent sets of propositions, exists in logical space. We have not the power to create or destroy propositions or possible worlds; while our behavior can, of course, determine which propositions are true and which possible world is actual, this is an entirely different matter. We can apprehend what is in logical space, but we cannot change it – and this means that we cannot create or destroy the sorts of entities that Floridi defines as information objects. Strictly construed, the universal laws are not really *moral* laws at all: insofar as they stipulate standards that are impossible for beings like us to violate, they more accurately express something like laws of nature that define the nature of logical space.

This, of course, is not the construction that Floridi has in mind. Normally, when we think of moral norms protecting information, what we have in mind are norms that protect *sources* of information. Someone who claims, for instance, that information about what might cure cancer should be preserved means that certain kinds of physical objects should be protected; these items include documents that record the results of research on cancer and physical objects (like corpses) that might provide oncologists with the insights they need to develop a cure. Thus construed, the norms protecting information really protect physical objects – like databases.

This construction of the universal laws can plausibly be grounded in the weaker claim that information (*qua* abstract objects that can be “apprehended” by human beings) is intrinsically valuable to human beings. Since information can be discovered and hence apprehended only by encountering the sorts of objects considered above, those objects have tremendous instrumental value because they are the *only* practical way for us to gain information. Thus, if, other things being equal, a state of affairs in which human beings apprehend a particular piece of information is morally better than a state of affairs in which human beings do not apprehend that particular piece, then it follows that the state of affairs in which the objects that convey such information is preserved is morally better, other things being equal, than the state of affairs in which these objects are not preserved. This construction of the universal laws is a corollary of the intuitive claim that information is intrinsically valuable *to us*.

But it is not at all clear how Floridi’s claim that information objects have intrinsic value could ground

this view. While Floridi wants to argue, for example, that a world *W* in which a particular database is preserved – presumably in the physical form of a disk containing it – is morally preferable to a world that is exactly like *W* except that the database is missing (p. 302), the claim that the database *qua* information object has moral standing will not do this work. The disk containing the database is not the information object; the relevant information object is an abstract object that contains a description of the disk (and its contents). Regardless of what happens to the disk, the relevant information object continues to exist in logical space, the only “place” in which it exists, because it is an abstract object – and not a physical one.

Of course, if the disk is destroyed, then human access to that information object – or, rather, the propositions contained in the database – is thereby eliminated, but respect for the existence of information objects *qua* information objects isn’t the sort of moral consideration that would require that easy *human access* to such objects be preserved. If, as Floridi believes, it makes sense to think of information objects as having goods of their own, that good extends only so far as the continued existence of those objects; it does not include human consumption or apprehension of those objects. Assuming it makes sense to think of information objects as having interests, it is clear, at most, that such objects have an interest in continued existence. The idea that it somehow conduces to the interest of an abstract object that this or that particular human being grasp, consume, or apprehend that object is very hard to make sense of. For example, while I suppose it makes sense to think that the number 2 (an abstract object consisting of a set theoretic object called an equivalence class) has an interest in continued existence, it is hard to see how my using or understanding the number 2 could do it any good.

At this point, one might object that the preceding arguments rest on a misinterpretation of Floridi’s view. While it is true that Floridi says things like “Respect information objects for their own sake if you can” (p. 300), the point of Floridi’s IE claim, on this line of analysis, is not so much to secure the point that abstract information objects *themselves* are deserving of respect; rather it is that every existing *physical* entity is deserving of respect in virtue, so to speak, of having an informational nature – a very different claim. For example, he points out that “[the IE claim] 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” (p. 299). Similarly, he argues that “all entities share a lowest common denominator, their nature as information objects, and... this... can contribute to our

ethical understanding" (p. 302). His point, the argument goes, is to show that all existing physical entities deserve respect "qua information objects."

While there is a sense in which Mary and a disk containing a database might plausibly be characterized as information objects, this is not the relevant sense of the notion as Floridi defines it. Human beings and the disk are *sources* of information in the following sense: each contains information in a form that can be made useful to human beings. Living human beings can convey that information by speaking or by having their bodies examined; a disk can convey its contents by being read and translated into natural language by a computer running the right program. But this is not what Floridi means by information object. As he defines the notion, an information object IO_E is a set of propositions describing the properties, attributes, and behaviors of an entity E and is thus ontologically distinct from E . The relevant information object pertaining to Mary (IO_{MARY}) is not Mary herself *qua* person with memories, beliefs, etc. Rather, it is a set of propositions that describe various aspects of Mary. As such, they are ontologically distinct objects.

This creates a profound problem for the IE claim. If information objects were identical with the entities they describe, then there would be no problem: the duty to respect IO_{MARY} and the duty to respect IO_{DISK} would simply amount to a duty to respect Mary and a duty to respect the disk. But Mary and the disk are radically different objects from IO_{MARY} and IO_{DISK} – indeed, so different that what constitutes moral respect for one has little bearing on how one should treat the other. Respecting Mary's wishes for how to dispose of her corpse cannot, for example, be construed as respect for the information object describing Mary. Dumping Mary's body in the river does not promote information entropy – the hallmark of informational disrespect on Floridi's view – regardless of whether Mary consents to having her corpse dumped in the river: if respect for IO_{MARY} doesn't require preserving her body at all costs, as seems eminently reasonable, then whether or not Mary consents to what is done to her body seems irrelevant from the standpoint of respecting IO_{MARY} .

Accordingly, while Floridi wants to derive minimal moral standing for Mary *qua* information object

from the minimal moral standing of the information object that describes Mary, he has not provided an adequate analytic framework that would link the standing of such very different objects as Mary the person and IO_{MARY} the information object. A concern to preserve abstract informational objects in logical space simply doesn't translate in any obvious way into moral respect for any physical entity – including disks that store symbol sequences that express databases.²⁴ At the very least, some kind of argument – one that is grounded in some sort of ordinary intuition or moral practice – is needed to bridge the gap between the two classes of objects. In the absence of a plausible argument supporting such a link, the thesis that information objects have moral standing will not successfully ground Floridi's view of IE as a distinct macroethical theory.

References

- L. Floridi. On the Intrinsic Value of Information Objects and the Infosphere. *Ethics and Information Technology*, 4(4): 287–304, 2002.
- K.E. Himma. The Question at the Foundation of Information Ethics: Does Information have Intrinsic Value? In T.W. Bynum, N. Pouloudi, S. Rogerson and T. Spyrou, editors, *Challenges for the Citizen of the Information Society: Proceedings of the Seventh International Conference on the Social and Ethical Impacts of Information and Communications Technologies*, ETHICOMP, 2004(a).
- K.E. Himma. The Moral Significance of the Interest in Information: Reflections on a Fundamental Right to Information. *Journal of Information, Communication, and Ethics in Society* 2(4), 191–202, 2004(b).
- A. Leopold. *A Sand County Almanac*. Oxford University Press, New York, 1977.
- P. Singer. All Animals are Equal. *Philosophic Exchange*, 1(1), 1974. Reprinted in P. Singer. *Animal Liberation*. Random House, Inc., New York, 1975.
- P. Singer. *Animal Liberation*. Avon Books, New York, 1977.
- B. Steinbock. Speciesism and the Idea of Equality. *Philosophy*, 53(204), 1978.
- P. Taylor. The ethics of Respect for Nature. *Environmental Ethics*, 3(3): 197–218, 1981.

²⁴ Indeed, this is why, at bottom, the universal laws that protect the infosphere are irrelevant to the issue of what should be done with Mary's corpse.