

Too Much Information: Questioning Information Ethics

John Barker

University of Illinois at Springfield

jbark3@uis.edu

1. Introduction

In a number of recent publications¹, Luciano Floridi has argued for an ethical framework, called Information Ethics, that assigns special moral value to information. Simply put, Information Ethics holds that all beings, even inanimate beings, have intrinsic moral worth, and that existence is a more fundamental moral value than more traditional values such as happiness and life. Correspondingly, the most fundamental moral *evil* in the world, on this account, is entropy – this is not the entropy of thermodynamics, but entropy understood as “any kind of *destruction, corruption, pollution, and depletion* of informational objects.” (Floridi 2007, p. 9) Floridi regards this moral outlook as a natural extension of environmental ethics, in which non-human entities are treated as possessors of intrinsic moral worth, and more specifically, of *land ethics*, where the sphere of moral patients is further extended to include inanimate but naturally occurring objects. On Floridi’s view, artifacts can also be moral patients, including even such “virtual” artifacts as computer programs and web pages.

In general, then, Floridi holds that *all objects* have some moral claim on us, even if some have a weaker claim than others; moreover, they have this moral worth intrinsically, and not because of any special interest we take in them. In this paper, I want to consider the

motivation and viability of Information Ethics as a moral framework. While I will not reach any firm conclusions, I will note some potential obstacles to any such moral theory.

2. Background and Motivation

Before continuing, we need to clarify the notions of *object* and *information*, as Floridi uses those terms. Briefly, an informational entity is any sort of instantiated structure, any pattern that is realized concretely. In particular, information is not to be understood semantically. An informational object need not have any semantic value; it need not represent the world as being this way or that. Instead, information should simply be thought of as structure.

Now any object whatsoever may be regarded as a realization of some structure or other. Floridi realizes this, and indeed expands on it in a view he terms “Informational Structural Realism” (ISR)². ISR is a metaphysical account of the world that basically dispenses with substrates in favor of structures. On this view, the world should be regarded as a system of realized structures, but it is a mistake to ask what substrate the structures are ultimately realized in: it is structures “all the way down.”

ISR is a fascinating thesis, but it will not be my purpose here to offer any further examination or critique of it. I mention it simply to show that when Floridi speaks of informational entities, he is really speaking of *arbitrary* entities. Information ethics is in fact a theory of arbitrary objects as moral patients. By casting it in terms of information, Floridi is stressing that the class of entities we should be concerned about, as moral

patients or otherwise, is broader than the familiar concrete objects of our everyday experience. It should include any sort of instantiated information whatsoever, be it a person, a piece of furniture, or a “virtual” web-based object.

Now Floridi’s central claim, that all entities have some (possibly very minimal) moral claim on us, while fascinating, certainly runs counter to most moral theories that have been proposed. It therefore seems reasonable to ask for some argument for it, or at least some motivation. The main rationale Floridi provides seems to be an argument from precedent. Before people started thinking systematically about ethics, they withheld the status of moral patient from all but the members of their own tribe or nation. Later, this status was extended to the whole of humanity. Many if not most people would now treat at least some non-human animals as moral patients, and some would ascribe moral worth to entire ecosystems and even to inanimate parts of nature. Thus, the history of ethical thinking is one of successively widening the sphere of our moral concern, and the logical end result of this process is to extend our moral concern to all of existence – or so Floridi argues.

However, as it stands this argument seems weak. True, there has been some historical tendency for moral theories to broaden the sphere of appropriate targets of moral concern. This tendency may continue indefinitely, until all of existence is encompassed. And then again, it may not. Here it is worth considering *why* at least some non-human animals are now generally considered moral patients. The main rationale, both historically and for most contemporary moral theorists, is that animals have a capacity for pleasure and

suffering. It does not matter for my purposes whether this is the only or best rationale for extending moral consideration to animals. The point is that *some* rationale was needed; the mere precedent of extending moral consideration from smaller to larger groups of humans was not itself a sufficient reason to further extend it to animals. Likewise, if we are to extend the sphere of moral patients still further, we will need a specific reason to do so, not just precedent.

The most ardent supporters of animal rights have always been Utilitarians, and Utilitarianism justifies the inclusion of animals with a specific account of what constitutes a benefit or harm. Namely, benefit and harm are identified with pleasure and suffering, respectively. Once this identification is made, all it takes to show that a given being is a moral patient is to show that it can experience pleasure and pain. If Floridi were to give a specific account of what constitutes benefit or harm to an arbitrary entity, that would go some way toward providing a rationale for Information Ethics. It is also desirable for another reason. Floridi's ethical account is "patient-oriented" (Floridi 2007, p. 8). This may or may not mean that it is a consequentialist theory; however, it seems fair to assume that in a patient-centered moral theory, an action's benefit or harm to moral patients plays a preeminent role in determining its rightness or wrongness. Thus, it seems desirable for such a theory to include a general account of benefit and harm. Moreover, such a theory can presumably be action-guiding only if it provides at least some such account.

Does Floridi offer such an account of benefit and harm? He does identify good and evil with “existence” and “entropy,” respectively; but as I will argue below, it is not clear that this amounts to a general account of benefit and harm, either to individual entities or to the universe or “infosphere” at large. Now to some extent this omission is understandable, given the pioneering nature of Floridi’s work. However, I will argue below that there are substantial obstacles in principle to providing any such account. The notion of an arbitrary object, I suspect, is simply too broad to support any substantive account of harm and benefit.

3. Information and Entropy

Let us first see why there is even a question about fundamental good and evil in Information Ethics. Floridi identifies existence as the fundamental positive moral value, and inexistence as the fundamental negative value. Thus, it might seem natural to suppose that an action is beneficial if it creates (informational) objects, and harmful if it destroys them, with the net benefit or harm identified with the net number of objects created or destroyed.

The trouble with this proposal is that given the broad conception of objects that we are working with, *every* act both creates and destroys objects. Since objects are simply instantiated patterns, there are indefinitely many objects present in any given physical substrate. Any physical change whatsoever involves a change in the set of instantiated patterns, thus creating and destroying informational objects simultaneously. Moreover, even if it were possible to count the number of informational objects in a given medium,

such a count would ignore the fact that some beings have more inherent moral worth than others: this is fairly obvious in its own right, and Floridi himself insists on it, asserting that some moral patients have a strong claim on us while for others, the claim is “minimal” and “overridable” (Floridi 2007, p. 9).

Thus, if we are to take seriously the idea that “being” is the most fundamental good and “inexistence” or “entropy” is the most fundamental evil, we cannot calculate good or evil by simply counting objects. A natural idea, and one which is somewhat suggested by Floridi’s term “entropy,” is that fundamental moral value should be identified with some overall measure of the informational richness or complexity of a system. This would preserve the idea of being and nonbeing as fundamental moral values while avoiding the difficulties involved in the simple counting approach.

One of the best-developed accounts of non-semantic information is *statistical information theory*. This theory, developed by Claude Shannon in the 1940s³, has been used very successfully to describe the amount of information in a signal without describing the signal’s semantic content (if any). Thus, it seems like a natural starting point for describing the overall complexity or richness of a system of informational objects.

Statistical information theory essentially identifies high information content with low probability. Specifically, the Shannon information content of an individual message M is defined to be $\log_2(1/p(M))$, where $p(M)$ is the probability that M occurs.⁴ As a special case, consider a set of 2^n messages, each equally likely to occur; then each message will

have a probability of 2^{-n} , and an information content of $\log_2(2^n) = n$ bits, exactly as one would expect. The interesting case occurs when the probability distribution is non-uniform; low probability events occur relatively rarely, and thus convey more information when they do occur.

As is well known, the definition of Shannon information content is formally almost identical to that of statistical entropy in physics. The entropy S of a given physical system is defined to be $S = k_B \ln \Omega$, where k_B is a constant (Boltzmann's constant) and Ω is the number of *microstates* corresponding to the system's *macrostate*. (A system's macrostate is simply its macroscopic configuration, abstracting away from microscopic details; the corresponding microstates are those microscopic configurations that would produce that macrostate.) Now for a given microstate q and corresponding macrostate Q , Ω is simply the probability that the system is in microstate q given that it is in macrostate Q . In other words, the entropy of a system is simply $k_B \ln (1/p_Q(q))$, where p_Q is a uniform probability distribution over the microstates in Q . Alternatively, if we posit a uniform probability distribution p over all possible microstates q , then we have $p_Q(q) = p(q) / p(Q)$, and thus $S = (k_B/p(q)) \ln p(Q) = -(k_B/p(q)) \ln (1/p(Q))$; the quantity $k_B/p(q)$ is a constant because the measure p is uniform. In any case, we have $S = K \log (1/p)$, where K is a constant and p is the probability of the state in question under some probability measure (the base may be omitted on the log because it only affects the result up to a constant, and may thus be subsumed in K). Thus, up to a proportionality constant, statistical entropy is a special case of Shannon information content.

However, it is the *wrong* special case, since as Floridi states very clearly, the fundamental evil which he refers to as “entropy” is not thermodynamic entropy. And indeed, in light of the second law of thermodynamics, thermodynamic entropy is not a reasonable quantity for moral agents to try to minimize. Thus, if we are to use Shannon information theory to capture the morally relevant notion of complexity, we will have to use a probability measure other than that described above. However, information theory does not offer us any guidance here, because it does not specify a probability measure: it simply *assumes* some measure as given. Typically, when applying information theory, we are working with a family of messages with well-defined statistics; thus, a suitable p is supplied by the context of the problem at hand.

Thus, Shannon information theory provides a measure of a system’s information content, but this measure is relative to a probability measure p . This presents an obstacle to explaining complexity in terms of Shannon information and simultaneously claiming that complexity is a fundamental, intrinsic moral value. If we allow complexity to be relative to a probability measure, then intrinsic moral worth will also be relative to a probability measure. Conceivably, different probability measures could yield wildly different measures of complexity and, thus, of intrinsic moral worth. Thus, it would appear to be necessary to pin down a single probability measure, or at least a family of similar probability measures, in a non-arbitrary manner.

And here is where things get tricky. What probability measure is the right one for measuring the complexity of *arbitrary* systems? Whatever it is, it must be a probability

measure that is in some sense picked out by nature, rather than by our own human interests and concerns. Otherwise complexity, and thus inherent moral worth, is not really objective, but is tied to a specifically human viewpoint. This goes against the whole thrust of Information Ethics, which seeks to liberate ethics from an anthropocentric viewpoint. Thus, we need to find a *natural* probability measure for our task. What might such a probability measure look like?

The best-known conception of objective probability is the frequentist conception. According to that conception, the probability of an outcome O of an experiment E is the proportion of times that O occurs in an ideal run of trials of E . To apply this notion, we need a well-defined outcome-type O , a well-defined experiment-type E , and a well-defined set of ideal trials of E – and if the latter set is continuous, a well-defined measure on that set. This is all notoriously difficult to apply to non-repeatable event tokens and to particulars in general. To assign a frequentist probability to a particular x , it is necessary to subsume x under some general type T , and different choices of T may yield different probabilities. In other words, the frequentist probability of a particular depends among other things on how that particular is described. Different ways of describing a particular will correspond to different conceptions of what it is to repeat that particular, and thus, to different measures of how frequently it occurs in a run of cases.

What this means for us is that the information content of a concrete particular depends, potentially, on how we choose to carve up the world. Again, this is not a problem in practice for information theory, since in any given application, a particular (frequentist)

probability measure is likely to be singled out by the problem's context. But in describing the information context of completely arbitrary objects, there is no context to guide us. In particular, if we subsume a concrete particular x under a commonly occurring type T , it receives a high frequentist probability, and correspondingly low Shannon information content. If we subsume that same particular under a rarely occurring type T^* , it receives a low probability and correspondingly high information content.

Thus, it is by no means obvious that there is a choice of probability measure that (a) is natural independently of our own anthropocentric interests and concerns, and (b) gives us a measure of complexity that is a plausible candidate for inherent moral worth, even assuming that the latter has any special tie to complexity in the first place. To be fair, it is also not obvious that there is not such a probability measure. As the measure p from thermodynamics shows, there is at least one natural way of assigning probabilities to physical states, one which does indeed yield a measure of complexity, albeit not the measure of complexity we are looking for. It also raises a further worry. The reason thermodynamic entropy is a bad candidate for basic moral disvalue is simply that it is always increasing, regardless of our actions. That is simply the second law of thermodynamics. What guarantee do we have that complexity, measured in any other way, is not also decreasing inexorably? Thermodynamic entropy can decrease *locally*, in the region of the universe we care about, at the expense of increased entropy somewhere else, and the same may be true for other measures of complexity. But this fact is surely irrelevant to a patient-centered, non-anthropocentric moral theory.

4. Information Everywhere

Statistical information theory is of course not the only way to capture the idea of complexity and structure. However, I would argue that the whole notion of complexity or information content becomes trivial unless it is tied to our interests (or someone's interests) as producers and consumers of information.

How much information is there in a glass of water? The obvious, intuitive answer is: very little. A glass of water is fairly homogeneous and uninteresting. Yet the exact state of a glass of water would represent an enormous amount of information if it were described in its entirety. There are approximately 7.5×10^{24} molecules in an eight ounce glass of water.⁵ If each molecule has a distinguishable pair of states, call them A and B , then a glass of water may be regarded as storing over seven trillion terabits of data. Further, let f be any function from the water molecules into the set $\{A, B\}$. Relative to f , we may regard a given molecule M as representing the binary digit 0 if M is in state $f(M)$, and 1 otherwise. Clearly, there is nothing to prevent us from regarding a glass of water in this way if we so choose, and with any encoding function f we like. And clearly, by a suitable choice of f , we may regard the water as encoding any data we like, up to about seven trillion terabits. For example, by choosing the right encoding function, we may regard the water as storing the entire holdings of the Library of Congress, with plenty of room to spare. Alternatively, a more "natural" coding function, say $f(M) = A$ for all M , might be used, resulting in a relatively uninteresting but still vast body of information.

Now if ordinary objects like glasses of water really do contain this much information, then there is too much information in the world for information content to be a useful measure of moral worth. The information we take a special interest in – the structures that are realized in ways that we pay attention to, the information that is stored in ways that we can readily access – is simply swamped by all the information there is. The moral patients we normally take an interest in are vastly outnumbered by the moral patients we routinely ignore. Floridi's estimate of the world's information, a relatively small number of exabytes, is several orders of magnitude lower than the *yottabyte* of information that can be found in a glass of water. Thus, if information content is to serve as a measure of moral worth, the information described in the previous paragraph must be excluded.

But on what basis could it be excluded? We might try to exclude some of the more unconventional encoding functions, such as the encoding function that represents the water as storing the entire Library of Congress. Such encoding functions, it may be argued, are rather unnatural and do not represent the information that is objectively present in the water. Even if this is so, there is no getting around the fact that a glass of water represents a vast amount of information, in that it would take much information to accurately describe its complete state. That information might be rather uninteresting – uninteresting *to us*, that is – but so what? If moral worth is tied to information content per se, then it does not matter whether that information is interesting. If moral worth is tied to *interesting* information, then it appears that moral worth is directly tied to human concerns after all.

But there is a more fundamental problem with dismissing some encoding functions f as unnatural. Whenever information is stored in a physical medium, there needs to be an encoding function to relate the medium's physical properties to its informational properties. Often, this function is "natural" in that it relates a natural feature of information (e.g. the value of a binary variable) to a natural feature of the physical medium (e.g. high or low voltage in a circuit, the size and shape of a pit on an optical disk, magnetic field orientation on a magnetic disk, etc.). However, there is absolutely no requirement to use natural encoding functions. There need be no simple relation whatsoever between, say, a file's contents and the physical properties of the media that store the file. The file could be encrypted, fragmented, stored on multiple disks in a RAID, broken up into network packets, etc.

In practice, we always disregard the information that is present, or may be regarded as present via encoding functions, in a glass of water. But the reason does not seem to be a lack of a natural relation between the information and the state of the water. The reason is that even though the information is in some sense there, we cannot easily use or access it. We can regard a glass of water as storing a Library of Congress, but in practice there is no good reason to do so. By contrast, a file stored in a possibly very complicated way is nonetheless accessible and potentially useful to us.

If this is right, then there is a problem with viewing information's intrinsic value as something independent of our own interests as producers and consumers of information. The problem is that information does not *exist* independently of our (or someone's)

interests as producers and consumers of information. Or alternatively, information exists in an essentially unlimited number of different ways: what we count as information is only a minute subset of all the information there is. Which of these two cases obtains is largely a matter of viewpoint. On the former view, even if inanimate information has moral value, it has value in a way that is more tied to a human perspective than Floridi lets on. On the latter, there is simply too much information in the world for our actions to have any net effect on it.

5. Conclusion

The immediate lesson of the last two sections is that overall complexity, or quantity of information, is a poor measure of intrinsic moral worth. Now this conclusion, even if true, may not appear to be terribly damaging to Information Ethics, as the latter embodies no specific theory of how to measure moral worth. It may simply be that some other measure is called for. However, I would argue that the above considerations pose a challenge to any version of Information Ethics, for the following reason.

As we have seen, the number of (informational) objects with which we interact routinely is essentially unlimited, or at least unimaginably vast. If each object has its own inherent moral worth, what prevents the huge number of informational objects that we do not care about from outweighing the relatively small number that we do care about, in any given moral decision? For example, I might radically alter the information content of a glass of water by drinking it, affecting ever so many informational objects; why does that fact carry less moral weight than the fact that drinking the water will quench my thirst and

hydrate me? The answer must be that virtually all informational objects have negligible moral value, and indeed Floridi seems to acknowledge this by saying that many informational objects have “minimal” and “overridable” value. But that claim is rather empty unless some basis is provided for distinguishing the few objects with much value from the many with little value.

Of course, one answer is simply to assign moral worth to objects based on *how much we care about them*. That would just about solve the problem. Moreover, that is more or less what it would *take* to solve the problem, insofar as the objects that must be assigned minimal value (lest ethics become trivial) are in fact objects that we do not care about. However, this is not an answer Floridi can give. Moral worth is supposed to be something objects possess intrinsically, as parts of nature. It is not supposed to be dependent on our interests and concerns. Thus, what is needed is an independent standard of moral worth for arbitrary objects which, while not based directly on human concern, is at least roughly in line with human concern. And so far that has not been done.

References

- Floridi, Luciano. 2008a. “Information Ethics, its Nature and Scope.” In *Moral Philosophy and Information Technology*, edited by Jeroen van den Hoven and John Weckert. Cambridge: Cambridge University Press.
- . 2008b. “A Defence of Informational Structural Realism.” *Synthese* 161.2. 219-253.
- . 2007. “Understanding Information Ethics.” *APA Newsletter On Philosophy and Computers* 7.1. 3-12.

MacKay, David J. C. 2003. *Information Theory, Inference, and Learning Algorithms*.
Cambridge: Cambridge University Press.

Shannon, Claude. 1948. "A Mathematical Theory of Computation." *Bell System
Technical Journal* 27. 379-423 & 623-656.

¹ See, for example, Floridi (2007), Floridi (2008a), etc.

² See Floridi (2008b).

³ See Shannon (1948). For a good modern introduction see MacKay (2003).

⁴ A base-2 logarithm is used because information is measured in bits, or base-2 digits. If information is to be measured in base-10 (decimal) digits, then a base-10 logarithm should be used. In general, the Shannon information content is defined to be $\log_b(1/p(M))$, with b determined by the units in which information is measured (bits, decimal digits, etc.).

⁵ This figure is obtained from the number of molecules in a mole (viz. Avogadro's number, approximately 6×10^{23}), the number of grams in one mole of water (equal to water's atomic weight, approximately 18), and the number of grams in 8 ounces (about 227).