Luciano Floridi

INFORMATION

A Very Short Introduction

OXFORD
UNIVERSITY PRESS

Chapter 8
The ethics of information

Our journey through the various concepts of information is almost complete. We started by looking at the information revolution, and we are now ready to see some of its ethical implications.

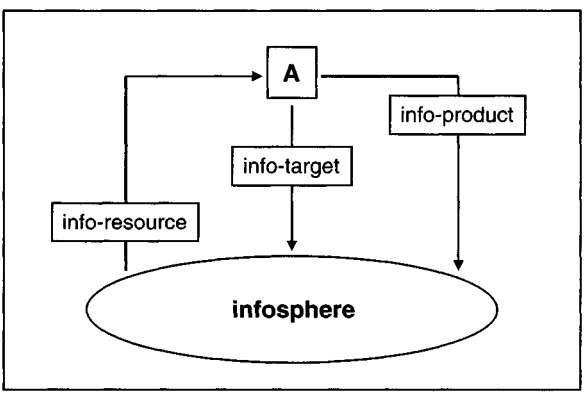
The previous chapters illustrated some crucial transformations brought about by ICTs in our lives. Moral life is a highly information-intensive game, so any technology that radically modifies the 'life of information' is bound to have profound moral implications for any moral player. Recall that we are talking about an ontological revolution, not just a change in communication technologies. ICTs, by radically transforming the context in which moral issues arise, not only add interesting new dimensions to old problems, but lead us to rethink, methodologically, the very grounds on which our ethical positions are based. Let us see how.

Information ethics as a new environmental ethics

ICTs affect an agent's moral life in many ways. For the sake of simplicity, they can be schematically organized along three lines, in the following way. Suppose our moral agent A is interested in pursuing whatever she considers her best course of action, given her predicament. We shall assume that A's evaluations and

interactions have *some* moral value, but no specific value needs to be introduced at this stage. Intuitively, *A* can avail herself of some information (information as a *resource*) to generate some other information (information as a *product*) and, in so doing, affect her informational environment (information as *target*). This simple model, summarized in Figure 19, will help us to get some initial orientation in the multiplicity of issues belonging to information ethics (henceforth IE). I shall refer to it as the RPT model.

The RPT model is useful to rectify an excessive emphasis occasionally placed on specific technologies (this happens most notably in *computer* ethics), by highlighting the more fundamental phenomenon of information in all its varieties and long tradition. This was also Wiener's position and the various difficulties encountered in the conceptual foundations of computer ethics are arguably connected to the fact that the latter has not yet been recognized as primarily an environmental ethics, whose main concern should be the ecological management and wellbeing of the infosphere.



19. The 'External' R(esource) P(roduct) T(arget) Model

Since the appearance of the first works in the 1980s, information ethics has been claimed to be the study of moral issues arising from one or another of the three distinct 'information arrows' in the RPT model.

Information-as-a-resource ethics

Consider first the crucial role played by information as a *resource* for *A*'s moral evaluations and actions. Moral evaluations and actions have an epistemic component, since *A* may be expected to proceed 'to the best of her information', that is, *A* may be expected to avail herself of whatever information she can muster, in order to reach (better) conclusions about what can and ought to be done in some given circumstances. Socrates already argued that a moral agent is naturally interested in gaining as much valuable information as the circumstances require, and that a well-informed agent is more likely to do the right thing. The ensuing 'ethical intellectualism' analyses evil and morally wrong behaviour as the outcome of deficient information. Conversely, *A*'s moral *responsibility* tends to be directly proportional to *A*'s degree of information: any decrease in the latter usually corresponds to a decrease in the former. This is the sense in which information occurs in the guise of judicial evidence. It is also the sense in which one speaks of *A*'s informed decision, informed consent, or well-informed participation. In Christian ethics, for example, even the worst sins can be forgiven in the light of the sinner's insufficient information, as a counterfactual evaluation is possible: had *A* been properly informed, *A* would have acted differently, and hence would not have sinned (Luke 23:34). In a secular context, Oedipus and Macbeth remind us how the mismanagement of informational resources may have tragic consequences.

From a 'resource' perspective, it seems that the moral machine needs information, and quite a lot of it, to function properly. However, even within the limited scope adopted by an analysis

based solely on information as a resource, and hence a merely semantic view of the infosphere, care should be exercised, lest all ethical discourse is reduced to the nuances of higher quantity, quality, and intelligibility of informational resources. The more the better is not the only, nor always the best rule of thumb. For the (sometimes explicit and conscious) withdrawal of information can often make a significant difference. *A* may need to lack (or preclude herself from accessing) some information in order to achieve morally desirable goals, such as protecting anonymity, enhancing fair treatment, or implementing unbiased evaluation. Rawls' 'veil of ignorance' famously exploits precisely this aspect of information-as-a-resource, in order to develop an impartial approach to justice in terms of fairness. Being informed is not always a blessing and might even be morally wrong or dangerous.

Whether the (quantitative and qualitative) presence or the (total) absence of information-as-a-resource is in question, there is a perfectly reasonable sense in which information ethics may be described as the study of the moral issues arising from 'the triple A': *availability*, *accessibility*, and *accuracy* of informational resources, independently of their format, kind, and physical support. Examples of issues in information ethics understood as an information-as-resource ethics are the so-called *digital divide*, the problem of *infoglut*, and the analysis of the *reliability* and *trustworthiness* of information sources.

Information-as-a-product ethics

A second, but closely related sense in which information plays an important moral role is as a *product* of *A*'s moral evaluations and actions. *A* is not only an information consumer but also an information producer, who may be subject to constraints while being able to take advantage of opportunities. Both constraints and opportunities call for an ethical analysis. Thus, information ethics, understood now as information-as-a-product ethics, may cover

moral issues arising, for example, in the context of *accountability*, *liability*, *libel legislation*, *testimony*, *plagiarism*, *advertising*, *propaganda*, *misinformation*, and more generally the *pragmatic rules of communication*. The analysis of the immorality of *lying* offered by Immanuel Kant (1724–1804) is one of the best known case studies in the philosophical literature concerning this kind of information ethics. Cassandra and Laocoon, hopelessly warning the Trojans against the Greeks' wooden horse, remind us how the ineffective management of informational products may have tragic consequences.

Information-as-a-target ethics

Independently of *A*'s information input (informational resources) and output (informational products), there is a third sense in which information may be subject to ethical analysis, namely when *A*'s moral evaluations and actions affect the informational environment. Examples include *A*'s respect for, or breach of, someone's information *privacy* or *confidentiality*. *Hacking*, understood as the unauthorized access to a (usually computerized) information system, is another good example. It is not uncommon to mistake it for a problem to be discussed within the conceptual frame of an ethics of informational resources. This misclassification allows the hacker to defend his position by arguing that no use (let alone misuse) of the accessed information has been made. Yet hacking, properly understood, is a form of breach of privacy. What is in question is not what *A* does with the information, which has been accessed without authorization, but what it means for an informational environment to be accessed by *A* without authorization. So the analysis of hacking belongs to an information-as-a-target ethics. Other issues here include *security*, *vandalism* (from the burning of libraries and books to the dissemination of viruses), *piracy*, *intellectual property*, *open source*, *freedom of expression*, *censorship*, *filtering*, and *contents control*. *Of the Liberty of*

Thought and Discussion by John Stuart Mill (1806–1873) is a classic of information ethics interpreted as information-as-target ethics. Juliet, simulating her death, and Hamlet, arranging for his father's homicide to be re-enacted, show how the unsafe management of one's informational environment may have tragic consequences.

The limits of any microethical approach to information ethics

At the end of this overview, it seems that the RPT model may help one to get some initial orientation in the multiplicity of issues belonging to different interpretations of information ethics. Despite its advantages, however, the model can be criticized for being inadequate, in two respects.

On the one hand, the model is too simplistic. Arguably, several important issues belong *mainly but not only* to the analysis of just one 'informational arrow'. The reader may have already thought of several examples that illustrate the problem: someone's testimony is someone else's trustworthy information; *A*'s responsibility may be determined by the information *A* holds, but it may also concern the information *A* issues; censorship affects *A* both as a user and as a producer of information; disinformation (i.e. the deliberate production and distribution of false and misleading contents) is an ethical problem that concerns all three 'informational arrows'; freedom of speech also affects the availability of offensive content (e.g. child pornography, violent content, and socially, politically, or religiously disrespectful statements) that might be morally questionable and should not circulate.

On the other hand, the model is insufficiently inclusive. There are many important issues that cannot easily be placed on the map at all, for they really emerge from, or supervene on, the interactions among the 'informational arrows'. Two significant examples may

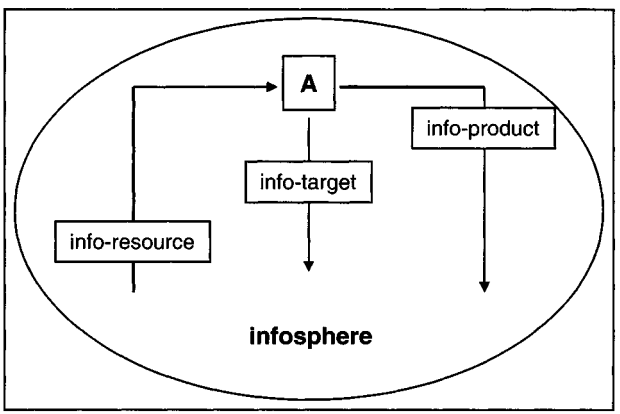
suffice: 'big brother', that is, the problem of *monitoring and controlling* any information that might concern *A*; and the debate about information *ownership* (including copyright and patents legislation) and *fair use*, which affects both users and producers, while shaping their informational environment.

So the criticism is reasonable. The RPT model is indeed inadequate. Yet why it is inadequate and how it can be improved are different matters. The tripartite analysis just provided is unsatisfactory, despite its partial usefulness, precisely because any interpretation of information ethics based on only one of the 'informational arrows' is bound to be too reductive. As the examples mentioned above emphasize, supporters of narrowly constructed interpretations of information ethics as a *microethics* (that is a practical, field-dependent, applied, and professional ethics) are faced by the problem of being unable to cope with a large variety of relevant issues, which remain either uncovered or inexplicable. The model shows that idiosyncratic versions of information ethics, which privilege only some limited aspects of the *information cycle*, are unsatisfactory. We should not use the model to attempt to pigeonhole problems neatly, which is impossible. We should rather exploit it as a useful scheme to be superseded, in view of a more encompassing approach to information ethics as a *macroethics*, that is, a theoretical, field-independent, applicable ethics.

A more encompassing approach to information ethics needs to take three steps. It must bring together the three 'informational arrows'. It has to consider the whole information cycle. And it needs to take seriously the ontological shift in the nature of the infosphere emphasized in the first chapter. This means analysing informationally all entities involved (including the moral agent *A*) and their changes, actions, and interactions, and treating them not apart from, but as part of the informational environment to which they belong as informational systems themselves. Whereas the first two steps do not pose particular

problems, and may be shared by other approaches to information ethics, the third step is crucial but involves an 'update' in the ontological conception of 'information' in question. Instead of limiting the analysis to (veridical) semantic contents – as any narrower interpretation of information ethics as a microethics inevitably does – an ecological approach to information ethics *also* treats information as an entity as well. In other words, we move from a broadly constructed epistemological or semantic conception of information ethics – in which information may be roughly equivalent to news or contents – to one which is typically ontological, and treats information as equivalent to patterns or entities in the world. Thus, in the revised RPT model, represented in Figure 20, the agent is embodied and embedded, as an informational agent or inforg, in an equally informational environment.

A simple analogy may help to introduce this new perspective. Imagine looking at the whole universe from a chemical perspective. Every entity and process will satisfy a certain chemical



20. The 'Internal' R(esource) P(roduct) T(arget) Model: the Agent A is correctly embedded within the infosphere

description. A human being, for example, will be mostly water. Now consider an informational perspective. The same entities will be described as clusters of data, that is, as informational objects. More precisely, our agent *A* (like any other entity) will be a discrete, self-contained, encapsulated package containing (i) the appropriate data structures, which constitute the nature of the entity in question, that is, the state of the object, its unique identity, and its attributes; and (ii) a collection of operations, functions, or procedures, which are activated by various interactions or stimuli (that is, messages received from other objects or changes within itself), and correspondingly define how the object behaves or reacts to them. At this level of analysis, informational systems as such, rather than just living systems in general, are raised to the role of agents and patients of any action, with environmental processes, changes, and interactions equally described informationally.

Understanding the *nature* of information ethics ontologically rather than epistemologically modifies the interpretation of its *scope*. Not only can an ecological information ethics gain a global view of the whole life cycle of information, thus overcoming the limits of other microethical approaches, but it can also claim a role as a macroethics, that is, as an ethics that concerns the whole realm of reality, as explained in the next section.

Information ethics as a macroethics

A straightforward way to introduce information ethics as a general approach to moral issues is by comparing it to environmental ethics. Environmental ethics grounds its analysis of the moral standing of biological entities and ecosystems on the intrinsic worthiness of *life*, and on the intrinsically negative value of *suffering*. It is biocentric. It seeks to develop a patient-oriented ethics in which the 'patient' may be not only a human being, but also any form of life. Indeed, land ethics extends the concept of patient to any component of the environment, thus coming close to

the approach defended by information ethics. Any form of life is deemed to enjoy some essential proprieties or moral interests that deserve and demand to be respected, at least minimally and relatively, that is, in a possibly overridable sense, when contrasted to other interests. So biocentric ethics argues that the nature and wellbeing of the patient (the receiver) of any action constitute (at least partly) its moral standing and that the latter makes important claims on the interacting agent. These claims in principle ought to contribute to the guidance of the agent's ethical decisions and the constraint of the agent's moral behaviour. The 'receiver' of the action, the patient, is placed at the core of the ethical discourse, as a centre of moral concern, while the 'transmitter' of any moral action, the agent, is moved to its periphery.

Substitute now 'life' with 'existence' and it should become clear what information ethics amounts to. It is an ecological ethics that is still patient-oriented but replaces *biocentrism* with *ontocentrism*. It suggests that there is something even more elemental than life, namely *being* – that is, the existence and flourishing of all entities and their global environment – and something even more fundamental than suffering, namely *entropy*. The latter is most emphatically *not* the concept of thermodynamic entropy discussed in Chapter 5, the level of mixedupness of a system. Entropy here refers to any kind of *destruction*, *corruption*, *pollution*, and *depletion* of informational objects (mind, not just of information as semantic content), that is, any form of impoverishment of reality. Information ethics then provides a common vocabulary to understand the whole realm of *being* informationally. It holds that *being*/information has an intrinsic worthiness. It substantiates this position by recognizing that any informational entity has a right to persist in its own status, and a right to flourish, i.e. to improve and enrich its existence and essence. As a consequence of such 'rights', information ethics evaluates the duty of any moral agent in terms of contribution to the growth of the *infosphere* and any process, action, or event

that negatively affects the whole infosphere – not just an informational entity – as an increase in its level of entropy and hence an instance of evil.

In information ethics, the ethical discourse concerns any entity, understood informationally, that is, not only all persons, their cultivation, wellbeing, and social interactions, not only animals, plants, and their proper natural life, but also anything that exists, from paintings and books to stars and stones; anything that may or will exist, like future generations; and anything that was but is no more, like our ancestors or old civilizations. Information ethics is impartial and universal because it brings to ultimate completion the process of enlargement of the concept of what may count as a centre of a (no matter how minimal) moral claim, which now includes every instance of *being* understood informationally, no matter whether physically implemented or not. In this respect, information ethics holds that every entity, as an expression of *being*, has a dignity, constituted by its mode of existence and essence (the collection of all the elementary proprieties that constitute it for what it is), which deserve to be respected (at least in a minimal and overridable sense), and hence place moral claims on the interacting agent and ought to contribute to the constraint and guidance of his ethical decisions and behaviour. This *ontological equality principle* means that any form of reality (any instance of information/*being*), simply for the fact of being what it is, enjoys a minimal, initial, overridable, equal right to exist and develop in a way which is appropriate to its nature. The conscious recognition of the ontological equality principle presupposes a disinterested judgement of the moral situation from an objective perspective, i.e. a perspective which is as non-anthropocentric as possible. Moral behaviour is less likely without this epistemic virtue. The application of the ontological equality principle is achieved, whenever actions are impartial, universal, and 'caring'. At the roots of this approach lies the *ontic trust* binding agents and patients. A straightforward way

of clarifying the concept of ontic trust is by drawing an analogy with the concept of 'social contract'.

Various forms of contractualism (in ethics) and contractarianism (in political philosophy) argue that moral obligation, the duty of political obedience, or the justice of social institutions gain their support from a so-called 'social contract'. This may be a hypothetical agreement between the parties constituting a society (e.g. the people and the sovereign, the members of a community, or the individual and the state). The parties agree to accept the terms of the contract and thus obtain some rights in exchange for some freedoms that, allegedly, they would enjoy in a hypothetical state of nature. The rights and responsibilities of the parties subscribing to the agreement are the terms of the social contract, whereas the society, state, group, etc. is the entity created for the purpose of enforcing the agreement. Both rights and freedoms are not fixed and may vary, depending on the interpretation of the social contract.

Interpretations of the theory of the social contract tend to be highly (and often unknowingly) anthropocentric (the focus is only on human rational agents) and stress the coercive nature of the agreement. These two aspects are not characteristic of the concept of *ontic trust*, but the basic idea of a fundamental agreement between parties as a foundation of moral interactions is sensible. In the case of the ontic trust, it is transformed into a primeval, entirely hypothetical *pact*, logically predating the social contract, which all agents cannot but sign when they come into existence, and that is constantly renewed in successive generations.

In the English legal system, a trust is an entity in which someone (the trustee) holds and manages the former assets of a person (the trustor, or donor) for the benefit of certain persons or entities (the beneficiaries). Strictly speaking, nobody owns the assets, since the trustor has donated them, the trustee has only legal ownership and the beneficiary has only equitable ownership.

Now, the logical form of this sort of agreement can be used to model the ontic trust, in the following way:

- the assets or 'corpus' is represented by the world, including all existing agents and patients;
- the donors are all past and current *generations* of agents;
- the trustees are all current *individual* agents;
- the beneficiaries are all current and future *individual* agents and patients.

By coming into being, an agent is made possible thanks to the existence of other entities. It is therefore bound to all that already exists, both *unwillingly* and *inescapably*. It *should be* so also *caringly*. *Unwillingly*, because no agent wills itself into existence, though every agent can, in theory, will itself out of it. *Inescapably*, because the ontic bond may be broken by an agent only at the cost of ceasing to exist as an agent. Moral life does not begin with an act of freedom but it may end with one. *Caringly* because participation in reality by any entity, including an agent – that is, the fact that any entity is an expression of what exists – provides a right to existence and an invitation (not a duty) to respect and take care of other entities. The pact then involves no coercion, but a mutual relation of appreciation, gratitude, and care, which is fostered by the recognition of the dependence of all entities on each other. Existence begins with a gift, even if possibly an unwanted one. A foetus will be initially only a beneficiary of the world. Once she is born and has become a full moral agent, she will be, as an individual, both a beneficiary and a trustee of the world. She will be in charge of taking care of the world, and, insofar as she is a member of the generation of living agents, she will also be a donor of the world. Once dead, she will leave the world to other agents after her and thus become a member of the generation of donors. In short, the life of a human agent becomes a journey from being only a beneficiary to being only a donor, passing through the stage of being a responsible trustee of the world. We

begin our career of moral agents as strangers to the world; we should end it as friends of the world.

The obligations and responsibilities imposed by the ontic trust will vary depending on circumstances but, fundamentally, the expectation is that actions will be taken or avoided in view of the welfare of the whole world.

The crucial importance of the radical change in ontological perspective cannot be overestimated. Bioethics and environmental ethics fail to achieve a level of complete impartiality, because they are still biased against what is inanimate, lifeless, intangible, or abstract (even land ethics is biased against technology and artefacts, for example). From their perspective, only what is intuitively alive deserves to be considered as a proper centre of moral claims, no matter how minimal, so a whole universe escapes their attention. Now, this is precisely the fundamental limit overcome by information ethics, which further lowers the minimal condition that needs to be satisfied, in order to qualify as a centre of moral concern, to the common factor shared by any entity, namely its informational state. And since any form of *being* is in any case also a coherent body of information, to say that information ethics is infocentric is tantamount to interpreting it, correctly, as an ontocentric theory.

The result is that all entities, *qua* informational objects, have an intrinsic moral value, although possibly quite minimal and overridable, and hence they can count as moral patients, subject to some equally minimal degree of moral respect understood as *a disinterested, appreciative, and careful attention.* As the philosopher Arne Naess (1912–2009) has maintained, 'all things in the biosphere have an equal right to live and blossom'. There seems to be no good reason not to adopt a higher and more inclusive, ontocentric perspective. Not only inanimate but also ideal, intangible, or intellectual objects can have a minimal

degree of moral value, no matter how humble, and so be entitled to some respect.

There is a famous passage, in one of Albert Einstein's letters, that summarizes well this ontic perspective advocated by information ethics. A few years before his death, Einstein received a letter from a 19-year-old girl grieving over the loss of her younger sister. She wished to know whether the famous scientist might have something to say to comfort her. On 4 March 1950, Einstein replied to her:

> A human being is part of the whole, called by us 'universe', a part limited in time and space. He experiences himself, his thoughts and feelings, as something separated from the rest, a kind of optical delusion of his consciousness. This delusion is a kind of prison for us, restricting us to our personal desires and to affection for a few persons close to us. Our task must be to free ourselves from our prison by widening our circle of compassion to embrace all humanity and the whole of nature in its beauty. Nobody is capable of achieving this completely, but the striving for such achievement is in itself a part of the liberation and a foundation for inner security.

Deep ecologists have already argued that inanimate things too can have some intrinsic value. In a famous article, the historian Lynn Townsend White, Jr (1907–1987) asked:

> Do people have ethical obligations toward rocks? [and answered that] To almost all Americans, still saturated with ideas historically dominant in Christianity... the question makes no sense at all. If the time comes when to any considerable group of us such a question is no longer ridiculous, we may be on the verge of a change of value structures that will make possible measures to cope with the growing ecologic crisis. One hopes that there is enough time left.

According to information ethics, this is the right ecological perspective *and* it makes perfect sense for any religious or spiritual

tradition (including the Judeo-Christian one) for which the whole universe is a divine creation, is inhabited by the divine, and is a gift to humanity, of which the latter needs to take care. Information ethics translates all this into informational terms. If something can be a moral patient, then its nature can be taken into consideration by a moral agent *A*, and contribute to shaping *A*'s action, no matter how minimally. In more metaphysical terms, information ethics argues that all aspects and instances of *being* are worth some initial, perhaps minimal and overridable, form of moral respect.

Enlarging the conception of what can count as a centre of moral respect has the advantage of enabling one to make sense of the innovative nature of ICTs, as providing a new and powerful conceptual frame. It also enables one to deal more satisfactorily with the original character of some of its moral issues, by approaching them from a theoretically strong perspective. Through time, ethics has steadily moved from a narrow to a more inclusive concept of what can count as a centre of moral worth, from the citizen to the biosphere. The emergence of the infosphere, as a new Athenian 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. Thus, information ethics represents the most recent development in this ecumenical trend, and an ecological approach without a biocentric bias. It translates environmental ethics in terms of infosphere and informational objects, for the space we inhabit is not just the earth.