This paper has been accepted for publication in

Ethics and Information Technology (Springer)

Information Technologies and the Tragedy of the Good Will

Luciano Floridi

Dipartimento di Scienze Filosofiche, Università degli Studi di Bari; Faculty of Philosophy and IEG, Computing Laboratory, Oxford University.

Address for correspondence: Wolfson College, OX2 6UD, Oxford, UK; luciano.floridi@philosophy.oxford.ac.uk

Abstract

Information plays a major role in any moral action. ICT (Information and Communication Technologies) have revolutionized the life of information, from its production and management to its consumption, thus deeply affecting our moral lives. Amid the many issues they have raised, a very serious one, discussed in this paper, is labelled the *tragedy of the Good Will*. This is represented by the increasing pressure that ICT and their deluge of information are putting on any agent who would like to act morally, when informed about actual or potential evils, but who also lacks the resources to do much about them. In the paper, it is argued that the tragedy may be at least mitigated, if not solved, by seeking to re-establish some equilibrium, through ICT themselves, between what agents know about the world and what they can do to improve it.

Keywords

Agents, Evil, Globalization, Good Will, Information Ethics, Tragedy of the Good Will.

Introduction: The Moral Values of Information

ICT affect an agent's moral life in many ways. Recently,[1] I suggested that these may be schematically organized along three lines (see Figure 1).

Suppose our agent *A* is interested in pursuing whatever she considers her best course of action, given her predicament. She is embodied and embedded, as an informational agent, in an equally informational environment, to which we may refer as the *infosphere*.[2] 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*).

The scheme in Figure 1 may help one to get some initial orientation in the multiplicity of issues belonging to Information Ethics.[3] It is also useful to explain why any technology that radically modifies the "life of information" is bound to have profound moral implications. Moral life is a highly information-intensive activity and ICT, by radically transforming the informational context in which moral issues arise,

[1] Luciano Floridi. Information Ethics, Its Nature and Scope. In Jeroen van den Hoven and John Weckert, editors, *Moral Philosophy and Information Technology*. Cambridge University Press, Cambridge, forthcoming.

[2] *Infosphere* is a neologism I coined years ago (see for example Luciano Floridi. *Philosophy and Computing: An Introduction*. Routledge London, New York, 1999.) on the basis of "biosphere", a term referring to that limited region on our planet that supports life. It denotes the whole informational environment constituted by all informational entities (thus including informational agents as well), their properties, interactions, processes and mutual relations. It is an environment comparable to, but different from cyberspace (which is only one of its sub-regions, as it were), since it also includes off-line and analogue spaces of information.

[3] The interested reader may find a detailed analysis of the model in Floridi. Information Ethics, Its Nature and Scope.

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.[4]

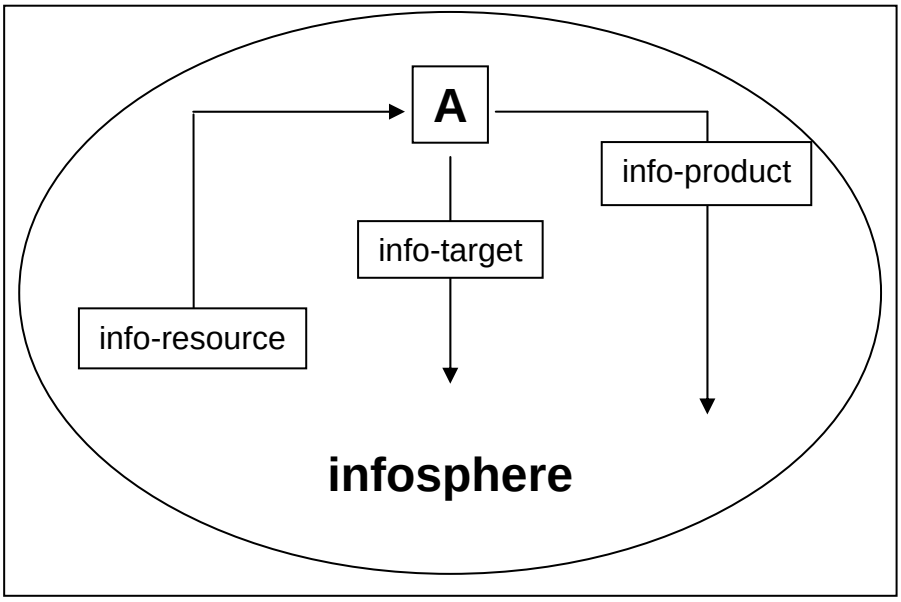

Figure 1. The Agent Embedded in the Infosphere

At the same time, the previous analysis rectifies an excessive emphasis occasionally placed on specific technologies (this happens most notably in *computer* ethics), by calling our attention to the more fundamental phenomenon of information in all its varieties and long tradition. This was also Wiener's position[5] and the various difficulties

[4] For a similar position in computer ethics see Walter Maner. Unique Ethical Problems in Information Technology. *Science and Engineering Ethics*, 2(2): 137-154, 1996. On the so-called "uniqueness debate" see Luciano Floridi, and Jeff Sanders. Mapping the Foundationalist Debate in Computer Ethics. *Ethics and Information Technology*, 4(1): 1-9, 2002. Herman T. Tavani. The Uniqueness Debate in Computer Ethics: What Exactly Is at Issue, and Why Does It Matter? *Ethics and Information Technology*, 4(1): 37-54, 2002.

[5] The classic reference here is to Norbert Wiener. *The Human Use of Human Beings: Cybernetics and Society, Rev. Ed.* Houghton Mifflin Boston, 1954. Terrell Bynum. Computer Ethics: Basic Concepts and Historical Overview. In Edward N. Zalta, editors, *The Stanford Encyclopedia of Philosophy*, 2001. has convincingly argued that Wiener may be considered one of the founding fathers of information ethics.

encountered in the conceptual foundations of computer ethics are arguably[6] connected to the fact that the latter has not yet been recognized as primarily an environmental ethics, whose main concern is (or should be) the ecological management and well-being of the infosphere.

Returning to the ways in which ICT affect an agent's moral life, in this paper I shall not be concerned with ontological issues (the third arrow in Figure 1, to simplify) but rather with the identification and solution of a key problem that arises in the context of the first two "arrows" in our model, namely when information is taken in a semantic sense, as a resource and a product. I shall refer to the problem as *the tragedy of the Good Will*.

The problem is simple, but making it explicit and precise, as well as suggesting some fruitful strategies for tackling it, will require careful analysis. This will be the goal of the next two sections. I hope the reader will bear with me while I make explicit a few simple assumptions and then outline an interpretation of the *tragic* and the *scandalous*.

Six Assumptions

Our first assumption has already surfaced, so let me make it fully visible:

1) "information" will be used here in its strongly semantic sense, in order to refer to syntactically well-formed, semantically meaningful and veridical data, like "Paris is the capital of France" or "the train to London leaves at 11 am". So we shall not be concerned with information in the *probabilistic* sense (Shannon's theory), in the *structural* sense (consider the local and global information that help one to solve

[6] Luciano Floridi. Information Ethics: On the Philosophical Foundations of Computer Ethics. *Ethics and Information Technology*, 1(1): 37-56, 1999.

crosswords or sudoku problems), in the *ontological* sense, especially as natural patterns (e.g. the DNA) or in the *instructional* sense (e.g. an algorithm).

Moral evaluations and actions have an informational component in the semantic sense just introduced. This component is crucial, since *A* may be expected to proceed "to the best of her information", availing 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. This is our second, Aristotelian assumption:

2) our moral agent *A* is interested in gaining as much relevant information as required by the circumstances, or, as Aristotle puts it at the beginning of his *Metaphysics*, we shall assume that "all men by nature desire to know". This may be for evolutionary reasons (one naturalistic way of reading Aristotle's "by nature") or because well-informed agents are more likely to do the right thing (a Socratic way of reading Aristotle's "by nature"). One can accept the assumption without necessarily embracing the ensuing ethical naturalism or intellectualism (which analyses evil and morally wrong behaviour as the outcome of deficient information).

Our third assumption concerns *A*'s limited powers:

3) *A* does not have boundless resources but is realistically constrained, especially by time, memory (i.e., amount of information storable and available), energy expendable to increase her information and capacities to handle it. This is not as bad as it looks. As is well known, moral action cannot presuppose any form of omnipotence. So one of the axioms of Standard Deontic Logic requires that, if it ought to be that *a*, then it is

permissible that *a* ($Oa \rightarrow Pa$), which in our context means that, if *A* must do *a* then *A* can do *a*.[7]

The previous condition goes some way towards mitigating the impact of our next assumption:

4) *A*'s moral *responsibility* tends to be directly proportional to *A*'s *degree of information*, any decrease in the latter usually corresponding to a decrease in the former. This is the important sense in which information may occur in the guise of judicial evidence, for example. It is also the sense in which one speaks of *A*'s informed decision, informed consent, or well-informed participation. The assumption also allows counterfactual evaluations: had *A* been properly informed *A* would have acted differently and hence would not have made the moral mistake that she made.

5) no *akrasia*. We shall assume that our agent *A* is capable of carrying out the course of action that she judges to be morally best. Although not very realistic (the practicing vs. preaching dichotomy is common to the point of being proverbial), this assumption is still plausible and it merely satisfies a simplicity requirement. *A*'s lack of *akrasia* means that *A* does not act against her judgment, but here it is not taken to mean that *A* has an intrinsic desire to act morally. For this anti-Hobbesian motivation, we need a last assumption:

6) *eudokia*. This Greek word means "good will", an expression made famous by the *Vulgata* version of Luke 2.14 ("pax hominibus *bonae voluntatis*" "peace to all men of *good will*"). It is in this original sense of *benevolent attitude*, or a willingness/desire to

[7] On the connection between epistemic and deontic logic see now Eric Pacuit, Rohit Parikh, and Eva Cogan. The Logic of Knowledge Based Obligation. *Synthese*, 149(2): 57-87, 2006.

do the right thing, that it is going to be used in this article.[8] This use of "Good Will" is eccentric with respect to Kant's well-known interpretation. According to Kant, a good will is the only thing that "can be taken as good without qualification". Its decisions are entirely dictated by moral demands, that is, by the moral law. In this paper, the Good Will overlaps with Kant's description deontologically, insofar as she (I use "it" to refer to Kant's conception) is identified as a privileged centre of morally good action. On the other hand, the God Will differs partly from Kant's description in a way that may be defined as "care-ethically", that is, insofar as she includes not only a purely rational but also a caring attitude. Our Good Will is expected to exhibit a *willingness* to engage with the world for its own sake and an *attentiveness* to (that is, interest in, concern with, and compassion for) its well-being. Both attitudes are extraneous to Kant's conception, as each requires an emotional and emphatic involvement. In our case, the rational and caring attitudes are supposed to be complementary and to add value to each other.

To summarize, we shall assume that *A* is a Good Will, endowed with some but limited resources, who bases her decisions and actions on the proper management of her information about the world, who is reasonably capable of implementing whatever she thinks ought to be done morally, whose responsibilities increase with the amount of information she enjoys (and who knows that this is so), and who is motivated by a genuine desire to know and by a sincere *eudokia,* while not suffering from *akrasia*. For the sake of simplicity, I shall refer to this type of agent as *the Good Will*.

[8] The reader should be warned that the discussion about the proper reading of the passage is a scholarly battlefield. Depending on whether one adds an "s" at the end of *eudokia* and make it a genitive, the reading changes from "Glory to God in the highest, and on earth peace, good will toward men" – which is the classic reading (but note that the good will in question is God's, not men's) – to "Glory to God in the highest, and peace among men in whom he is well pleased", which has strong Calvinist implications in favour of the predestined. Either way, the *Vulgata* seems a misleading translation, if suggestive.

The Good Will is an ideal but not an idealized agent. As in any scientific experiment in which one tries to abstract from irrelevant details and obtain ideal conditions (e.g. by referring to frictionless models in dynamic experiments), we can use the Good Will to bring to light and properly formulate an important problem caused by ICT. But first, one last round of clarifications, as promised.

The Tragic and the Scandalous

To understand the tragedy of the Good Will we need to appreciate what the tragic means. The suggestion developed in this section is that the tragic arises from a lack of balance between information and power in the presence of *eudokia*, i.e. of a Good Will's (the agent's) inclination to act morally. "Power" refers to the bounded skills, resources, means, etc. needed to implement a morally good action (see point 3 above). "Information" refers to how much (or little) the Good Will knows about the world, including past events, current circumstances and future implications or effects (see point 1 above). Without *eudokia* there is no sense of the tragic, but the presence of *eudokia* is insufficient to give rise to the tragic, since the Good Will might actually succeed in her endeavours. For the tragic to arise, there also needs to be a fundamental lack of balance. A few classic examples will help to clarify the point.

1) Lucretius: no Good Will, no tragedy.

Lucretius in his *De Rerum Natura* (Book II, Proem) provides a beautiful illustration of information without either Good Will or power: "Tis sweet, when, down the mighty main, the winds / Roll up its waste of waters, from the land / To watch another's labouring anguish far, / Not that we joyously delight that man / Should thus be smitten,

but because 'tis sweet / To mark what evils we ourselves be spared [...]". Lucretius is presenting here the detached and content ataraxia to be developed by the philosophical mind. If there is a lack of involvement (*apathia*) and Good Will – in this case a desire to help and intervene – then it is not tragic but sweet to witness someone else's anguish, for the struggle is only in the object observed, and not in the observer. Compare this to the following, equally famous scene of shipwrecking.

2) Miranda: the tragic as a result of Good Will, Information and Power.

When in *The Tempest* (Act I, Scene II) Shakespeare describes Miranda watching from afar the apparent sinking of "a brave vessel", he makes her utter the following words (emphasis added): "If by your art, my dearest father, you have / Put the wild waters in this roar, allay them. / The sky, it seems, would pour down stinking pitch, / But that the sea, mounting to the welkin's cheek, / Dashes the fire out. *O, I have suffered/ With those that I saw suffer*: a brave vessel, / Who had, no doubt, some noble creature in her, / Dash'd all to pieces. O, the cry did knock / Against my very heart. Poor souls, they perish'd. / *Had I been any god of power*, I would / Have sunk the sea within the earth or ere / It should the good ship so have swallow'd and / The fraughting souls within her." Two points deserve our attention.

First, both Lucretius and Miranda may be assumed to be witnessing the same disaster. But Miranda is a Good Will ("I have suffered with those that I saw suffer"). Her *eudokia* makes her wish she were able to match her alleged information (in fact, it will turn out that no "noble creature" is "dashed to pieces") with some equal power, which, in this case, would require a god-like (*demiurgic*, more on this later) degree of control over the elements ("had I been any god of power I would have sunk the sea

within the earth"). She knows that the tragic would disappear if only her (the Good Will's) power were equal to her information.

The second point is that the tragic will indeed later vanish when Miranda/the Good Will realizes that she was misinformed. So we, readers and audience, are confronted by a lack of balance between power and information that can be restored either by making the former match the latter (what Miranda would like to do), or by making the latter match the former (what in fact will happen). Such a lack of balance, as the essence of the tragic, is openly evident in Oedipus and Cassandra.

3) Oedipus: the tragic as a result of Good Will, Power and Lack of Information.

On the one hand, Oedipus has only some limited information about his horrific future (he is told that he will kill his father and marry his mother) but lacks the relevant information (he was adopted; the man he kills on his journey is his real father; the woman he later marries is his real mother). On the other hand, Oedipus has quite a lot of power to implement his *eudokia* and try to avoid his destiny (he leaves his home town and those whom he believes to be his parents, thus hoping to escape his destiny; he later becomes king). It is because Oedipus is a Good Will that his fate is tragic. But his tragedy is entirely informational: his desire to do the right thing is combined with the (royal) power to carry over his decisions but also with the wrong sort of information. So it is not accidental that Oedipus becomes king of Thebes (marrying his mother Jocasta) through an informational rite of passage, by answering the riddle of the Sphinx; that it is a *blind* source who sees better than him (the seer Teiresias) who reveals to Oedipus his real fate; and that Oedipus, in the end, punishes himself by forcing his mother's brooch pins into his eyes. Greek epistemology is very visual, being informed is seeing.

Our last example is equally classic, but shows a lack of balance in terms of lack of power, not of information.

4) Cassandra: the tragic as a result of Good Will, Information but Lack of Power.

Cassandra can predict ("hear") the future, a gift from Apollo, but is also cursed by the same god, so that her predictions will never be believed. This is a source of endless frustration and pain, as nobody acts on her accurate warnings. She is the Good Will that has all the necessary information (about the Trojan Horse and Troy's destruction; or about Agamemnon's and her murder) but who is powerless when it comes to avoiding the foreseen events.

To summarize: the tragic occurs in the presence of a Good Will (Miranda), when she is sufficiently powerful but insufficiently informed (Oedipus), or sufficiently informed but insufficiently powerful (Cassandra). Since the tragic is due to a lack of balance, and any balance is a matter of fine tuning, the risk of the tragic in either form is constant. When the tragic occurs, it is a scandal.

The *scandalous* is how the tragic may be perceived by its observers. Oedipus' and Cassandra's tragic predicaments are scandalous not because they set bad examples (for nobody would follow them), but because they show to the observers the ultimate, titanic failure of the Good Will. In a context in which the essence of agenthood is largely constituted by its *eudokia*, the agent who "gives scandal" has, by the same token, annihilated her essence, and thus ceased to be an agent altogether. For the Good Will, giving scandal is tantamount to committing suicide or being terminated. This is how one may interpret the famous quote from the Matthew's Gospel: "He that shall scandalize one of these little ones, that believe in Me, it were better for him that a mill-

stone should be hanged about his neck, and that he should be drowned *in the depth of the sea*." (Mt. 18:6, emphasis added). In the desperate sea of Miranda, that is, not of Lucretius.

We are finally ready to analyse the relation between ICT, the tragic and the scandalous.

The Twofold Relation between ICT and the Tragic

Given the forms in which the tragic (and hence the scandalous) may occur, it is not surprising that the relation between the informational revolution, brought about by ICT, and the tragic, might be twofold.

On the one hand, we have what I have labelled elsewhere the *IT-heodicean problem*.[9] New or digital ICTs provide the Good Will with increasing opportunities – directly or indirectly, from nanotechnology to risk assessment modelling, from bioinformatics to neuroscience, from genetic engineering to telemedicine and so forth – to prevent, defuse, control or eradicate evil. Information is power, as we all know. It follows that, the more powerful the Good Will becomes – in terms of science and technology and ICT in particular – the wider becomes the scope of her responsibilities for what is within her power to influence is. Thus, ICT greatly contribute to the increasing moral pressure put on the Good Will and her insufficient information about what ought to be done. It is as if the Good Will had the means to do something for the welfare of the world, but did not see how. Like Oedipus, when evil finally occurs, the

[9] The theodicean debate concerns whether it is possible to reconcile the existence of God and the presence of evil; the IT-heodicean debate concerns the problem of artificial evil. I have discussed it at length in Luciano Floridi, and Jeff Sanders. Artificial Evil and the Foundation of Computer Ethics. *Ethics and Information Technology*, 3(1): 55-66, 2001 .

Good Will can only blame herself, for had she been better informed, evil might have been avoidable. The tragedy of her inability is also the scandal of her annihilation as a moral agent.

ICT erode the scope of natural evil, re-cataloguing it as moral, or, as André Gide once put it, “man’s responsibility increases as that of the gods decreases” and ICT play a major role in this shifting process. Not that the process itself is either new or limited to ICT. Already Homer could write “Look you now, how ready mortals are to blame the gods. It is from us, they say, that evils come, but they even of themselves, through their own blind folly, have sorrows beyond that which is ordained.” (*Odyssey*, I.30-35). But ICT have made the process snowball.

On the other hand, if ICT have increased by orders of magnitude a Good Will’s capacity to cope with the world, they have also submerged her with information about the endless evils that she should be worried about. This is Cassandra’s predicament, which I suggest may be labelled, to differentiate it from the *IT-heodicean problem*, the *tragedy of the Good Will*.

The Tragedy of the Good Will

Good Wills are regularly submerged and often overwhelmed by information about evils in the world that they are unable to prevent, defuse, control or eradicate. In the past, less information meant less responsibility. Nowadays, ICT keep inundating the Good Wills with distressing news about famine, diseases, wars, violence, corruption, injustices, environmental disasters, poverty, lack of education, racism and so forth, on a daily basis. The list is endless, the disasters hearth-breaking, the responsibilities mounting, the sense of scandalous powerlessness nauseating. Confronted by so much information about so many moral failures, the Good Will cannot help feeling frustrated, aggrieved and guilty. A concrete example will render the analysis less academic and ivory-tower. It concerns the sea again.

On 14 August 2003, *The Economist* published an article ("The next big wave"), in which one could read that "[in the western Pacific] […]. Since 1990, ten big tsunamis have claimed more than 4,000 lives. So it would be nice to be able to detect such tsunamis far enough in advance for people to be evacuated. […] What is needed are specific detectors that take advantage of the fact that tsunamis are felt throughout the ocean's depths, unlike wind-generated waves, which affect only its surface." The article carried on to discuss several technologies and techniques for detecting, analysing, classifying and predicting tsunamis. It concluded: "Technology, though, can do only so much. […] Coastal dwellers must be able to recognize the signs of a possible tsunami – such as strong, prolonged ground shaking – and seek higher ground at once. As with any hazard, the more informed the public are, the better their chances of survival."

Despite all this information, on 26 December 2004, the Sumatra-Andaman earthquake caused a series of devastating tsunamis that spread throughout the Indian Ocean, killing approximately a quarter of a million people, with thousands of others missing. No ICT (tsunami warning systems) were in place to mitigate the impact of the catastrophe. It was one of the deadliest disasters in modern history. On the other hand, thanks to ICT, Good Wills everywhere in the world "suffered with those whom they saw suffer", almost in real time. Morally speaking, it was an instance of the tragedy of the Good Will.

It would be easy to speculate about future disasters that will be equally tragic and scandalous in the technical sense of the words specified above. Think of global warming, nuclear proliferation, the Palestinian problem, or AIDS in Africa, for example. But the point should be sufficiently clear to require no further illustration. Instead, one aspect that is worth emphasising here is how the Good Will might be inclined to develop skilful forms of ignorance or blind spots. As Plato remarks in the *Republic* (478c), the soul might decide not to pursue *nous* (knowledge and understanding) but *agnoia* (ignorance and irrationality), and dwell in "that which *is not* (*at all*)", *entropy*, in the vocabulary I introduced some time ago. Let me explain.

If the analysis offered so far is even roughly correct, a Good Will will feel pain and frustration when informed about evil events, and the more so the more she is informed about dramatic events with respect to which she is powerless. But it is also reasonable to assume that no Good Will will be inclined to leave open such a perennially bleeding wound. If one suffers too much with those whom one sees suffer, one may soon wish to avert one's eyes. So the risk that the Good Will constantly runs is

that of unwittingly (when not consciously) and innocently trying to avoid her Cassandra-like predicament by shutting herself off in her own informational niche. The dialectic is simple, and well captured by two well-worn phrases: since "what the eye does not see, the heart cannot grieve", the Good Will is constantly tempted to "bury her head in the sand". ICT have made the need for such hiding more strongly and widely felt, insofar as they have increased the potential exposure of the Good Will to evil.

The result, once again, is well epitomized in our digital age by the phenomenon of the so-called *The Daily Me*. The term, coined by Negroponte[10] some time ago, refers now to any news system (including news feeds) tailored to, customized by, or personalized for the reader's interests and tastes. The problem with *The Daily Me* is that it can easily become a mere mirror of one's own idiosyncratic biases, thus contributing to what David Weinberger has called the "echo chamber", information spaces where like-minded people unwittingly (and this is the risk) communicate only with people who already agree with them, reinforcing and never really challenging their belief systems.[11]

This filtering phenomenon, however, is not new. On the contrary, it might help to explain, for example, why the Germans managed to organize the concentration camps (recall: no Good Will no tragedy) while largely failing to grasp the horror of the Holocaust in all its magnitude (the agnostic Good Will).[12] What I am suggesting here is not that the Germans did not know at all, or that there was insufficient information

[10] Nicholas Negroponte. *Being Digital*. Knopf: Distributed by Random House New York, 1995.
[11] For a critical discussion of The Daily Me effects see Cass R. Sunstein. *Republic.Com*. Princeton University Press Princeton, N.J., 2001.
[12] The issue of how much the German population knew about the Holocaust is still debated. Robert Gellately. *Backing Hitler: Consent and Coercion in Nazi Germany*. Oxford University Press Oxford, 2001. has provided mass media evidence in favour of the hypothesi that Germans knew quite a lot about the Holocaust, but it seems that what the research shows, rather, is that they could have known quite a lot, had they wished to know it.

available to anyone who cared to check it, but that many Germans, confronted by such horrors and by the costly consequences of any disagreement with the Nazi regime, preferred not to see what was happening. As Dahrendorf wrote "It is certainly true that most Germans 'did not know' about National Socialist crimes of violence; nothing precise, that is, because they did not ask any questions."[13] Not asking questions, not seeing, not believing what one hears, filtering and rationalising evil: this is the common trap into which *weak* (see the comment above about *akrasia*) Good Wills tend to fall. No one is less informed than the person who does not want to be informed.

Paradoxically, Good Wills may therefore be the worst witnesses, the more so the more they are morally good and hence sensitive to evil.[14] As a consequence, Good Wills may have to be forced to keep their eyes open in front of the horrors that are being committed in their backyards. This might seem almost a torture. It reminds one of the "Ludovico technique" in *A Clockwork Orange* (1971), the cult film directed and produced by Stanley Kubrick. There, the protagonist, Alex, is forced to keep his eyes mechanically and painfully wide-open, while being shown scenes of intense violence, cruelty, and social aberration, including *The Triumph of the Will* by Leni Riefenstahl, the infamous propaganda documentary about the 1934 Nazi Party Congress in Nuremberg. Alex is not a Good Will but a psychopath, who enjoys violence. His conditioning is supposed to rehabilitate him. In the case of the Good Will, the metaphorical "Ludovico technique" that should be applied by ICT has a different effect,

[13] Ralf Dahrendorf. *Society and Democracy in Germany*. 1st ed. Doubleday Garden City, N.Y., 1967.
[14] Compare this to the conclusions reached by Pacuit, Parikh, and Cogan. The Logic of Knowledge Based Obligation. about the Kitty Genovese case: "In 1964, a young woman was stabbed to death while 38 neighbours watched from their windows but did nothing. The reason was not indifference, but none of the neighbours had even a default obligation to act, even though, as a group, they did have an obligation to take some action to protect Kitty".

for it is supposed to prevent her from burying her head in the sand of ignorance. It is one of the ethical tasks that a free press and uncensored ICT should have in any decent democracy.

Escaping the Tragic Condition

There may be plenty of reasons for being pessimistic about the tragedy of the Good Will, not least historical records. Perhaps information about preventable or solvable evils will keep pouring in, and we will forever be unable to do anything about it. One good thing about such pessimism, however, is that, if correct, it would require no action and Lucretius' attitude might be the only serious alternative. In contrast, if some optimism is even partially justified, the bad news is that this is cause for further toil, and not just pragmatically, but also theoretically. For more discussion of the possible strategies available to escape the tragic becomes indispensable. In this section, I hope to make a first step in such a direction.

There seem to be four main ways in which the tragedy of the Good Will might be escaped. Luckily, they are mutually compatible and hence possibly synergetic. Before discussing them, let me briefly outline them here:

1. the information/power gap may decrease, as information has already reached its peak, whereas power is catching up;
2. from quantity to quality of information: better informed Good Wills can act and exercise their augmented power better;

3. from the powerless observation of the single Good Will to the empowered interactions of multiagent systems of Good Wills: global problems and distributed morality require global agents;
4. the ontological side of information: the need for an augmented ethics.

Each strategy requires some comments.

1) More power. To begin with, although ICT and the corresponding amount of available information have seen an extraordinary development in the last half century, Good Wills have also witnessed a steady increase in their powers. For a rough estimate, one may adopt a brute translation into dollars per person. According to the World Bank, the proportion of people living in extreme poverty (less than $1 a day) in developing countries dropped by almost half between 1981 and 2001, from 40 to 21 percent of the global population (23 April, 2004). Again in 2004, the Bank's annual statistical report showed a drop in the absolute number of people living on less than $1 a day in all developing countries from 1.5 billion in 1981, to 1.1 billion in 2001, with much of the progress occurring in the 1980s. Of course, these are merely quantitative measures, but they do provide some ground for cautious optimism. Good Wills might be able to put ever more dollars where the bad news events conveyed by their ICT occur, thus helping to restore some balance between information and power.

2) Better information. The second way of tackling the tragedy of the Good Will is by using the same ICT, which can bring so much information about the evils in the world, to empower the *individual* Good Will. This is not a simple matter of more or less

information. Depending on contexts and usage, more information might be a benefit (more control, more competition, more choice and less censorship) or a curse, since sometimes less information might be preferable (more fairness and less bias, more privacy, more security). Too often these issues are left unqualified (what information?), and uncircumstantiated (information for whom? under which conditions? for what purpose?). Rather, empowering the single Good Will seems to be a matter of more "quality information", in the sense that future ICT should provide her with more feedback (whether and how the single agent's efforts and resources are affecting reality), more transparency (information constrains other agents' misbehaviour, as speed camera show), more forecasting (information is prevention) and more engineering (information as building capability).

3) Global agents. The careful reader might have noticed a tension between, on the one hand, the IT-heodicean problem and Oedipus' predicament (sufficient power, insufficient information) and, on the other hand, the tragedy of the Good Will and Cassandra's predicament (sufficient information, insufficient power). How is it possible that ICT can generate both predicaments? If they are empowering both pragmatically and informationally, surely these are two sides of the same coin, so their effects should overlap and cancel each other out, at least to a large degree. Make Oedipus and Cassandra work together, as it were, and it won't be necessary to escape the tragic condition because none will arise in the first place.

The tension is indeed there, but the inference drawn from it is mistaken, for it is based on a confusion of levels of agenthood. The IT-heodicean problem affects the

Good Will in so far as the latter refers to supra-individual agents. In this sense, it is ultimately *humanity* that is empowered by ICT. For example, none of us individually could have done anything to prevent the Sumatra-Andaman devastating tsunamis, but humanity as a whole could and should. The tragedy of the Good Will, on the other hand, affects single individual agents: it is you and I, John and Mary, Peter and Jenny who are subject by ICT to the dialectic of being informed about evils against which we are largely unable to do anything of comparable magnitude. It is us individually who give scandal.

It follows that the third strategy consists in identifying this mismatch and re-aligning individual and global agents, in order to make sure that the latter inherit the *eudokia* of the former and act on it. It might be easier to overcome both the IT-heodicean problem and the tragedy of the Good Will if we could work on developing global artificial agents – i.e., non-human (engineered) and/or social (e.g. groups, organizations, institutions), global agents – capable of channelling and guiding the energies of the single Good Wills who constitute them. National states, NGOs, international organizations or multinational companies are just some examples of this sort of supra-individual, global, artificial agents that are hybrids of other artificial agents (imagine the member states of the EU, or the software and hardware systems that contribute to the existence of a company) and individual people. This general strategy calls for more conceptual analysis, in order that we might understand artificial agents better, and clear outlines about how moral artificial agents may be built, morally educated or trained, and controlled.

4) Augmented ethics not super-ethics. It might be felt that the impact of ICT on our lives could be entirely reduced to a matter of DUMB effects: Doing & Understanding More & Better. If this were the case, then DUMB effects would transform man (the supra-individual Good Will) into superman. Superman has super-responsibilities and so ICT would require a super-ethics. The problem would then be that any super-ethics would be, for each of us single human agents, supererogatory, as it would require super-heroes. The mistake, in this case, is to confuse not only the level of agenthood, but also the scope of the impact of ICT. ICT are not just a matter of DUMB effects. New or digital ICT *re-ontologize*[15] the very nature of the infosphere, that is, of the environment itself, of the agents embedded in it and of their interactions. Since they also have an essentially ontic impact, they radically transform old realities and create entirely new ones. And because of their ontic impact, ICT require an augmented ethics for the whole of humanity as the ultimate Good Will, not for individual super-heroes. It follows that nowadays the IT-heodicean problem and the tragedy of the Good Will call for an ethics of creators (demiurges, as in Plato's *Timaeus*) and not of mere end-users of reality.[16] Or, to put it slightly differently, since the Good Will is increasingly morally responsible for designing and implementing the world the way it is, the moral question concerning her responsibilities is as much ethical as ontological, namely how she (both as an individual and as a supra-individual or global agent) could act as a morally good demiurge. It seems that her augmented responsibilities require an ecological approach to the whole reality.

[15] The neologism is constructed following the word "re-engineering" ("to design and construct anew").
[16] Luciano Floridi, and Jeff Sanders. Internet Ethics: The Constructionist Values of Homo Poieticus. In Robert Cavalier, editors, *The Impact of the Internet on Our Moral Lives*. SUNY, New York, 2005.

Conclusion: Towards a Global Consensus?

By way of conclusion, I would like to point to a specific case in which some of the previous suggestions seem to have found a first application. This is the *Copenhagen Consensus*, a project conceived and organized by Bjørn Lomborg (http://www.copenhagenconsensus.com/).

What would be the best ways to spend additional resources on helping the developing countries? Resources are scarce, and their allocation is therefore a specific case of triage, which demands difficult choices among good projects. In 2004, the Copenhagen Consensus project attempted to set priorities among a range of suggestions on how to improve standards of life in developing countries on the basis of a cost-benefit analysis. Eight economists, including four Nobel laureates, met on 24-28 May, 2004 at a roundtable in Copenhagen, and produced a ranking, based on applied welfare economics, of the 30-50 identified opportunities on which $50 billion of new money for development initiatives might be best spent. Ten global challenges were chosen: civil conflicts, climate change, communicable diseases, education, financial stability, governance, hunger and malnutrition, migration, trade reform, water and sanitation. With something close to unanimity, the panel put measures to restrict the spread of HIV/AIDS at the top of the ranking. It also rated all four top proposals “very good”, as measured by the ratio of social benefit to cost. The bottom of the list, however, aroused more controversy. All three of the schemes proposed to the panel for mitigating climate change (including the Kyoto protocol on greenhouse-gas emissions) were rated “bad”, meaning that their costs were estimated to exceed their benefits. The panel may meet again in 2008 for a second round.

Regardless of whether one shares the conclusions of the panellists, several aspects of the Copenhagen consensus resonate positively with the analysis developed in this paper.

First, the Copenhagen Consensus itself should be interpreted as a supra-individual Good Will, that is, as a multiagent system constituted by individuals, institutions and communication systems satisfying those conditions laid down in the second section of this article ("six assumptions").

Second, the Good Will gave priority to information above any other consideration, including politics and religion. Of all the problems tackled, it was clear that the most pressing was to have some reliable information on which problems to tackle first. An ethics of information was the stage against which the decisional procedure took place.

Third, the Copenhagen Consensus clearly meant to offer a series of strategies to other global Good Wills (again, understood as supra-individual agents) while at the same time informing individual Good Wills (the public) about what it considered to be the most economically fruitful and morally justifiable approach to global challenges. So there was no confusion in levels of agenthood, while the needs of both individual and global agents were addressed.

Fourth, despite appearances, the Consensus adopted a strongly ecological approach: it was clear that it wished to provide a balanced assessment of how limited resources could be best employed to improve the world. That some solutions to solve environmental problems were deemed to be unsatisfactory said nothing about the importance of the issues they were addressing.

Fifth, in a way that complements the previous remark about ecologism, the Consensus was an explicit attempt to develop a demiurgic approach to global issues. One of the assumptions behind the Copenhagen exercise is that the world will change according to human initiatives and that sorting them out and prioritising them is of vital importance.

Last, but equally importantly, since its beginning the Copenhagen Consensus project has itself been subjected to open discussion and made the subject of that flow of information that ICT have taught us to take for granted.

In a phrase, it was information about information, or an open and rational process of discussion about what needs to be done first.[17] There seem to be few better ways of dealing with the world's most serious problems.

Acknowledgements

I discussed several drafts of this paper at many meetings. It was first the topic of a talk, "From Augmented Intelligence to Augmented Responsibility", that I gave many years ago at the Computing and Philosophy Conference at Oregon State University (Corvallis, 24-26 January, 2002). I am grateful to Jon Dorbolo, for his kind invitation and the great opportunity to exchange some thoughts with Douglas Engelbart, and to OSU for its financial support. I then discussed a revised version at one of the seminars, "Some Ethical Consequences of Global Information", which I gave at the Norwegian University of Science and Technology (Trondheim, 2 February 2005). For that

[17] A reference to Habermas or Rorty might be "de rigueur" here, but I would rather point in the direction of Charles Sanders Peirce's ethics of research, and his famous invitation "do not block the way of inquiry" (Charles S. Peirce. *Collected Papers*. Edited by Charles Hartshorne and Paul Weiss. 8 vols. Harvard University Press Cambridge Mass., 1931-1958. Vol. I, par. 135.

opportunity and for the feedback received, I wish to thank Eric Mointeiro, Knut Rolland, Johnny Søraker and May Thorseth, and for the financial support, the Programme for Applied Ethics, the Department of Computer and Information Science, the Department of Philosophy and the Forum for Theory of Science at NTNU. I presented a further revised version at the Dipartimento di Scienze della Comunicazione, Università degli Studi di Salerno (Fisciano, 30 November 2005). Thanks to Roberto Cordeschi, Alessandro Laudanna and Teresa Numerico for their kind invitation, and to the Dipartimento di Scienze della Comunicazione for the financial support. I gave a virtually final version at the 2005 Uehiro/Carnegie Joint Conference "Information Ethics: Agents, Artifacts and New Cultural Perspectives", which Julian Savulescu and I organised in Oxford (8-9 December, 2005). I am grateful in this case to Julian himself, also for his invitation to organize and coordinate the conference with him, to Terry Bynum, Roger Crisp, Charles Ess and Kei Hiruta, Deborah Johnson and Jim Moor for our discussions, and to the Carnegie Council on Ethics and International Affairs, the Oxford Uehiro Centre for Practical Ethics, and the Uehiro Foundation on Ethics and Education for their financial support. I had the opportunity to test the final version during an invited Departmental Seminar at OUCL (Oxford, UK, 2 March 2006), for which I am grateful to Alexandru Baltag. Finally, I wish to thank the participants in these meetings for their helpful discussions, Ken Herold for calling my attention to the paper by Pacuit, Parikh and Cogan, and Paul Oldfield for his copyediting suggestions. Terry Bynum and Charles Ess kindly agreed to act as referees and provided a wealth of crucial suggestions. As usual, any remaining mistakes, even after so much feedback and revision, are all mine.

References

Terrell Bynum. Computer Ethics: Basic Concepts and Historical Overview. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*, 2001.

Ralf Dahrendorf. *Society and Democracy in Germany*. 1st ed. Doubleday Garden City, N.Y., 1967.

Luciano Floridi. Information Ethics: On the Philosophical Foundations of Computer Ethics. *Ethics and Information Technology*, 1(1): 37-56, 1999.

Luciano Floridi. *Philosophy and Computing: An Introduction*. Routledge London, New York, 1999.

Luciano Floridi. Information Ethics: An Environmental Approach to the Digital Divide. *Philosophy in the Contemporary World*, 9(1): 39-45, 2002

Luciano Floridi. Information Ethics, Its Nature and Scope. In Jeroen van den Hoven and John Weckert, editors, *Moral Philosophy and Information Technology*. Cambridge University Press, Cambridge, forthcoming.

Luciano Floridi, and Jeff Sanders. Artificial Evil and the Foundation of Computer Ethics. *Ethics and Information Technology*, 3(1): 55-66, 2001

Luciano Floridi, and Jeff Sanders. Mapping the Foundationalist Debate in Computer Ethics. *Ethics and Information Technology*, 4(1): 1-9, 2002.

Luciano Floridi, and Jeff Sanders. Internet Ethics: The Constructionist Values of Homo Poieticus. In Robert Cavalier, editors, *The Impact of the Internet on Our Moral Lives*. SUNY, New York, 2005.

Robert Gellately. *Backing Hitler*: *Consent and Coercion in Nazi Germany*. Oxford University Press Oxford, 2001.

Walter Maner. Unique Ethical Problems in Information Technology. *Science and Engineering Ethics*, 2(2): 137-154, 1996.

Nicholas Negroponte. *Being Digital*. Knopf: Distributed by Random House New York, 1995.

Eric Pacuit, Rohit Parikh, and Eva Cogan. The Logic of Knowledge Based Obligation. *Synthese*, 149(2): 57-87, 2006.

Charles S. Peirce. *Collected Papers*, edited by Charles Hartshorne and Paul Weiss. 8 vols. Harvard University Press Cambridge Mass., 1931-1958.

Cass R. Sunstein. *Republic.Com*. Princeton University Press Princeton, N.J., 2001.

Herman T. Tavani. The Uniqueness Debate in Computer Ethics: What Exactly Is at Issue, and Why Does It Matter? *Ethics and Information Technology*, 4(1): 37-54, 2002.

Norbert Wiener. *The Human Use of Human Beings: Cybernetics and Society, Rev. Ed.* Houghton Mifflin Boston, 1954.