JPE

Losing Your Way in the Fog

Reflections on *Evil Online*

PHILIP KITCHER
Columbia University

I

Many philosophers are inclined to dismissive judgments about 'applied ethics'. They think of this line of work as merely requiring straightforward use of ethical principles (the achievements of thinkers who tackle 'fundamental questions') in light of research into the technical details of the pertinent domain. So writing a book about the moral issues posed by the widespread use of the Internet, especially the impact of social media on contemporary lives, ought to be a simple matter. Learn the facts about what is going on, call up your favorite moral theory, and turn the crank.

One of the achievements of *Evil Online* is to provide a decisive refutation of this all-too-common view.[1] To be sure, Dean Cocking and Jeroen van den Hoven offer a parade of appalling stories about life on the World Wide Web. They also show, however, how hard it is to adapt standard ethical categories and principles to comprehend the online behavior that so concerns them. The habits they describe and condemn, often appearing as propensities to staggering cruelty,[2] require the introduction of a new concept: that of 'moral fog'.

What follows will be, in the main, an attempt to develop some of Cocking and van den Hoven's important themes. My principal focus will be on their

1. Dean Cocking and Jeroen van den Hoven, *Evil Online* (New York: Wiley, 2018).
2. As with the opening example, describing the hacking of the Epilepsy Foundation of America website, Cocking and van den Hoven, pp. 1–2.

Contact: Philip Kitcher <psk16@columbia.edu>

https://doi.org/10.3998/jpe.2378

conceptual innovation (explicitly motivated by the idea of 'the fog of war'[3]). First, however, I want to consider two independent points.

II

Virtually nobody maintains any more that the Internet is an unmitigated disaster. As Cocking and van den Hoven correctly remark, 'It is no longer an environment where, so far as getting correct answers to questions goes, it could be quite difficult to tell true from false'.[4] Yet their further judgment, hailing the Internet as having 'incredible epistemic power', by enabling people to answer 'just about any' of their questions,[5] needs qualification. Surely, there is a wide range of issues for which a few quick clicks can transform ignorance into true belief (maybe even knowledge?). Despite all the naysayers, Wikipedia has turned out to be an exceptionally valuable resource for those curious about all sorts of things. Or, more exactly, for all those things—and they are legion—about which the experts have achieved consensus *and* that do not threaten the ideas and values of a significant group of people. When either of these conditions lapses, activating a favorite search engine can do more epistemic harm than good. A banal fact of internet life is the invisible hand guiding searchers to sites fitting the profile already constructed for them.[6] Platforms live by the indulgence of advertisers, repaying the largesse by facilitating effective marketing. The pages that appear first on the screen are attuned to the prior history of touring the Web, and the accompanying advertisements harmonize with the interests and preferences attributed to searchers. As an epistemic consequence, a search can all too easily represent only one side of a disputed question, generating the false impression of consensus, where expert agreement has yet to be reached.

Even worse is the fostering and maintenance of controversy on issues that have already been settled, when those issues bear on questions of public policy. An outstanding example is the case of climate change. The reality of anthropogenic global warming has been recognized by climate scientists for well over three decades. Yet, even today, many Internet sites challenge the expert consensus, offering apparently 'scientific' graphs and figures to trace an alternative history of the Earth's mean temperature.[7] Slick videos contest the research, and

3. Cocking and van den Hoven, p. 86.

4. Cocking and van der Hoven, p. 40.

5. Cocking and van der Hoven, p. 40.

6. See Nicola Mössner and Philip Kitcher, 'Knowledge, Democracy, and the Internet', *Minerva*, 55 (2017), pp. 1–24.

7. See Philip Kitcher and Evelyn Fox Keller, *The Seasons Alter* (New York: Norton/Liveright 2017), especially chapter 1.

even the honesty, of leading climate scientists. The result is a mass of misinformation that has surely retarded action to combat a threat as severe as any our species has faced in its recorded history.

During the years since the advent of cable news, political divisions have fragmented the news media. In turn, the fragmentation has intensified the divisions. The vicious spiral has been accelerated by the Internet, with the consequence of making policy debates ever more contentious and difficult to resolve. Democracy relies on an informed citizenry, able to align voting preferences with genuine interests. When confusion is sown, and when an alleged source of knowledge partitions the electorate, democracy's claim to promote the freedom of the people is undermined. Citizens troop to the polls, and, 'instructed' in part by what they have read online, select candidates whose policies are at odds with their most central goals and aspirations. In the act supposed to express their freedom, they defeat what they most want to achieve. It is one of the great ironies of our times.

Cocking and van den Hoven's book is largely concerned with the dark side of the internet. Their wish to register some awareness of the epistemic benefits it has brought is easy to understand. Yet here, too, the large costs should be noted. What does it profit a species to gain the entire wisdom of Wikipedia and lose both the best (or least bad?) form of government and its planet as well?

III

Cocking and van den Hoven explore episodes and patterns of behavior whose pathologies are so striking as to silence any objections about the aptness of their title. 'Evil' is notoriously difficult to define, and, wisely, they don't venture a definition. As with obscenity, we know it when we see it, and we are shown it again and again throughout their discussions. Evil actions are often seen as extreme instances of moral badness. Nevertheless, for all the repugnance and even horror they arouse, they may not be the most significant wrongs committed in the online world.

Should those who design and write for sites disseminating disinformation relevant to major policy issues be included among the evildoers? They are certainly less flamboyant than the characters who star in Cocking and van den Hoven's horror stories. Yet the damage they do to human lives may be orders of magnitude greater than that perpetrated by the cyberbullies, not only in extent but also in intensity. Purveyors of 'alternative realities' today may open the path to the autocracies of tomorrow; the 'enemies of the people' rounded up by future dictators to be imprisoned and tortured may be connected by a long causal process to people who sowed Internet deceit in the interests of short-term political

gain. The 'merchants of doubt'[8] whose desires to maintain the profits derived from fossil fuels lead them to spread confusion on the Web are likely to retard action to address climate change; the costs of their contributions will be measured in future deaths, from starvation, drought, fire, flood, and (possibly agonizing) pandemics. What these malefactors do isn't obviously evil, at least not in any everyday sense. Given the high stakes, however, it should surely be a matter of moral concern.

An obvious reply: misconduct of the types just considered is importantly different from Cocking and van den Hoven's central focus. The turpitude of the actors is only contingently connected with the online world. The internet simply serves them as a useful means for implementing their designs. Fair enough. Nevertheless, there are aspects of Internet life Cocking and van den Hoven (rightly) consider, that share the potential for huge damage to human lives and to human society. In various passages, they are sensitive to the ways in which the increasing prominence—even dominance—of 'screens' in daily life affects social relations and the development of children.[9] Much of their discussion connects to the forms of online conduct on which they mostly focus. My own worries are more general. Even when there are no obvious pathologies, lives spent largely online threaten important values. They may bring about stunted and impoverished forms of human existence.

It's commonplace, fully appreciated by Cocking and van den Hoven, that Facebook and its ilk have modified our concept of friendship. As they rightly point out, when people measure themselves by the number of their 'friends', the quality and depth of individual friendships is likely to decrease.[10] Moreover, the pressure to fashion the most positive image of oneself, thus improving one's 'statistics' is a distortion of autonomous development.[11] Cocking and van den Hoven are somewhat reassured by the overlap between online and offline social connections.[12] All these are important points and deserve emphasis. I wonder, though, whether they go far enough.

Friendships vary in closeness and in the depth of mutual understanding.[13] Because our time is limited, multiplying the number of friends we have diminishes the amount of time—and attention—we can devote to each of them. Of course, the limited resource, time-spent-with-*X*, might be apportioned unequally.

8. See Naomi Oreskes and Erik Conway, *Merchants of Doubt* (New York: Bloomsbury, 2010).
9. Cocking and van der Hoven, *Evil Online*, pp. 41, 50, 75–77, and 78.
10. Cocking and van der Hoven, p. 76.
11. Cocking and van der Hoven, pp. 75 and 78.
12. Cocking and van der Hoven, p. 78.
13. For a sensitive philosophical discussion of friendship, from which I have learned much, see Alexander Nehamas, *On Friendship* (New York: Basic Books, 2016).

A large number of superficial friendships could coexist with a small number of relationships achieving the depth and intimacy realized in the bonds we most admire. Whether that occurs for any large proportion of those for whom life is infused with hours spent on social media is a matter for sociological research. Assuming, however, that many people now enter into friendships more diluted—thinner—than those central to the lives of the past, the consequence is likely to be registered in new standards and norms of human relationships. At a time when many reflective people are worried about the ways in which conceptions of living well are often dominated by consumerism—'He who dies with the most toys wins'—the abandonment of intimacy with its rewards and its often stringent demands would be an enormous loss.

The metaphorphosis I fear might occur without any of the pathologies Cocking and van den Hoven focus on. Imagine a world in which the proportion of the day spent online continues to increase. Suppose the internet becomes as safe and well-mannered as you like. Cyberbullying, sexual predation, revenge porn, cruel pranks all become things of the past. Yet the pressure to advertise oneself positively on social media remains, and even intensifies. When pretenders are unmasked, they are not humiliated or publicly scorned. The offenses are noted. Reprimands are administered firmly but with restraint. Liars lose friends by the score and, because worth is measured by the sheer number of 'friends', the punishment is felt.

In this envisaged world, people connected on the internet still sometimes interact offline. Because of the hours they devote to social media, the interactions are typically less frequent and shorter than those through which friendships of the kinds most admired are developed and sustained. Most importantly, some *dimensions* of friendship become rarer. The spate of online chat doesn't offer much opportunity for serious exploration of goals, for thinking through uncertain prospects together, for providing and receiving aid or consolation, for sharing the deepest joys. When an intimate friendship has already been formed, contact online can provide resources for maintaining it (although it may still require occasions on which friends can talk face to face or act in a joint project). What strikes me as less clear is how multidimensional intimacy is achieved without shared experiences, without episodes of standing together against some common threat, without the moments when troubles are confessed and advice is sought. The world I have imagined has banished online evil. Despite that, it is a world in which one of the most valuable aspects of human life has been reduced and cheapened.

New technologies have frequently invited jeremiads. In an older generation, concerned parents often worried about the effects of television on their children's development. So, it is reasonable to reply to my concerns by charging me with repeating a familiar Luddite complaint. There's no evidential basis

for supposing the future I've imagined to have any large probability. Nonetheless, it is surely a possibility, one that should arouse concern. Discussions of the damage wrought by the internet should not be restricted to the manifest evils. Vigilance about more subtle and insidious ways of making the world worse is, I suggest, a good idea.

IV

At the heart of Cocking and van den Hoven's project is a cluster of related questions. How does it happen that, placed in an online environment, some people—not obviously 'bad people'—do appalling things? What has caused them to deviate from their normal, morally acceptable, patterns of conduct? How does our understanding of morality and of the moral training of the young need to be expanded to reduce the frequency with which the pathologies occur?

Cocking and van den Hoven approach these issues by proposing that the evildoers are in a moral fog. Their moral education has instilled into them habits of action and reflection able to guide them correctly under many of the conditions they experience. In front of their screens, however, something goes wrong. The fog descends. They lose their compass and their way. And they do appalling things.

These strike me as good questions, and I am sympathetic to Cocking and van den Hoven's approach to them. Much of what they say is insightful. Yet, as they are surely aware, the notion of a moral fog is itself—well—foggy. In what follows, I shall try to suggest a different way of articulating their problem and of developing techniques for addressing their questions. My reformulation of their central themes already hints at the character of my alternative. Where they emphasize the fog, I consider the compass.

Compasses can fail to guide for all sorts of reasons. They may be mislaid, or unreadable in the available light, or damaged, or useless because one doesn't know which of the points marks their destination or because they are affected by the presence of a sufficiently strong magnet—or one may simply not see how to proceed in the indicated direction. Switching metaphors from fog to compass opens up the possibility of viewing what initially appears as a unified phenomenon as covering a range of different cases. I'll exploit that possibility.

A good place to start is with a philosophical debate about which Cocking and van den Hoven write with admirable cogency and clarity. The rich history of discussion of how apparently good (normal, law-abiding) people can commit evils has raised questions about whether it is apt to ascribe a standing character to a person and thus whether talking of people as 'good people' or 'bad people' makes sense. Cocking and van den Hoven's review of work in social psychology

is well-informed and accurate. One important source doesn't figure in it: the joint work of Walter Mischel, Yuichi Shoda, and Rodolfo Mendoza-Denton.[14] The study I have in mind was a follow-up to Mischel's earlier research, in which he had amassed considerable evidence for situationism (the thesis that conduct varies in important ways with the surrounding circumstances). That earlier conclusion was further confirmed by Mischel and his coauthors. Observing behavior at a summer camp, they recorded how traits like *being outgoing* and *being willing to take risks* varied across contexts. But they were also able to show how individuals had distinctive profiles. When placed in situations of one kind, a child would be deferential; in the rest of the observed circumstances, deference vanished. Their research questioned the oversimple way in which we use vocabulary to discuss character. We assume people to have stable characters in being disposed to exhibit a particular type of behavior across the board: the brave person will act bravely, come what may. Our actual character traits are more complex, consisting in a spectrum of behavioral responses keyed to different classes of situations.

Following Mischel, Shoda, and Mendoza-Denton, let's think of moral training (considered as something that occurs not only in youth but throughout a lifetime) as setting up, at any given point in a person's existence, a complex of dispositions. Some of these dispositions are *habits of action*. Much of our conduct occurs without reflection. People often don't ask themselves whether to wave to an acquaintance, or if they should step back to let someone pass by, or whether they should head off to work in the morning. With respect to these and myriad other parts of our conduct, we frequently don't wonder if what we are doing is permissible. We simply do it.

Of course, there are occasions on which we do stop to think, when some feature of the context triggers another disposition—a morally crucial disposition—that suspends the habit of action and calls on us to take stock. As Bernard Williams famously argued, it's a mistake to suppose reflection always to be justified.[15] The husband who engages in moral inquiry before plunging in to save his drowning wife has had 'one thought too many'. Indeed, as Shakespeare taught us in *Hamlet*, people can have *many* thoughts too many. An ideally trained moral agent would be disposed to moral meditation in all and only the situations for which the proper course is unclear, leaving habit to operate unchecked in each of the rest. Much of everyday moral education consists in equipping the young with the habits society has approved, freeing the growing moral agent to work things through on the occasions on which matters are confused.

14. 'Situation-Behavior Profiles as a Locus of Consistency in Personality', *Current Directions in Psychological Science*, 11, no. 2 (2002), pp. 50–54.

15. Bernard Williams, 'Persons, Character, and Morality', in *Moral Luck* (Cambridge: Cambridge University Press, 1981).

Among the habits also required are those put to work in moral reflection. It isn't enough to be aptly disposed to pause or to let habit generate action. Besides knowing *when* to reflect, one also has to know *how*. According to many influential views of moral life, this know-how is largely (even completely) reducible to a system of moral principles. When a problematic situation has caused one to stop and think, they consult their corpus of precepts, find the one that fits the circumstances, determine what it counsels them to do on this occasion, and then they follow that advice. Views like this lay at the root of the 'turn-the-crank' conception of applied ethics with which I began. (Perhaps I caricature—but many discussions of biomedical ethics, environmental ethics, and business ethics seem to adopt this kind of simplistic picture.)

In any number of taxing moral predicaments, recommending this type of procedure is hopeless. We can't find any precept fitting the case at hand, or we don't understand the moral dimensions of the circumstances well enough to decide which of several potential principles best suits it, or there are several principles judged to be applicable and they pull in contrary directions.[16] What we have explicitly been taught doesn't seem to help. Nor, when we try to find analogies with other situations, whether they are real instances we have experienced or historical cases that have been thoroughly analyzed, or fictitious examples we or others have constructed, do we take them to resolve our quandary. The moral propositions we believe don't apply themselves.

Moral philosophy is often beguiled by a faulty analogy with natural science. The complete moral theory is supposed to be akin to the picture Newton (and many scientists—and even more fans of science—after him) envisaged for the complete system of the world. There will be a small number of fundamental principles (of extraordinary generality) from which any correct moral statement can be derived by inserting the pertinent boundary and initial conditions. The subsequent history of the natural sciences has belied the Newtonian dream.[17] Yet, even if it had been sustained, the daily *practice* of the sciences is remarkable for the common difficulty of bringing high theory to bear on the local situation. As Thomas Kuhn pointed out, anyone who has ever learned any significant amount of any science encounters a recurrent phenomenon: the chapter presents a small number of new principles; the student understands those principles, committing them firmly to memory; but, faced with the exercises at the end, the mind goes completely blank.[18]

16. Dewey offers a cogent statement of the difficulty. See *Human Nature and Conduct*, in *The Middle Works*, vol. 14 (Carbondale: Southern Illinois University Press, 1988), p. 74.

17. See John Dupré, *The Disorder of Things* (Cambridge, MA: Harvard University Press, 1993); and Nancy Cartwright, *The Dappled World* (Cambridge: Cambridge University Press, 1999).

18. T. S. Kuhn, *The Structure of Scientific Revolutions* (Chicago: University of Chicago Press, 1962).

Contemporary accounts of scientific practice emphasize the skills, acquired in training and subsequently put to work. Much of the professional's knowledge is tacit. Why should moral practice be different? Instead of thinking of moral training as simply presenting a list of moral principles and as disciplining the will to follow the propositions engraved on consciousness by proper education, conceive it as delivering a set of dispositions. These dispositions—*sensitivities* to give them a name—guide moral agents as they encounter new situations.

The basic sensitivity distinguishes those contexts in which action-guiding habits are allowed to proceed from those triggering a need for reflection. Probably none of us has a perfect version of this sensitivity. We go wrong in two main ways, either by plunging ahead when we should stop to reflect or by dithering when we ought to act decisively. (The contrast is encapsulated in two familiar proverbs, whose opposition brings out the failure of the morality-as-a-system-of-principles view: 'Look before you leap!' and 'He who hesitates is lost'.) Within these two large classes of potential errors, we have our own individual propensities to go astray. For each of us, there is a set of contexts in which we mistakenly let habit prevail when we should pause to reflect, as well as a set of circumstances in which we wrongly inhibit our customary behavior.[19] The two sets define a personal profile, a complex example of the kinds of profiles studied by Mischel, Shoda, and Mendoza-Denton. These profiles connect directly to the questions raised by Cocking and van den Hoven: Is it common for internet users either to follow habits where they ought to stop and think or to suspend habits of everyday consideration for others after engaging in morally distorted musings, reassuring themselves of the permissibility of behavior they would otherwise firmly reject?

The latter question raises issues about the derivative sensitivities our moral training instills in us. Mapping those sensitivities, and recognizing the general form of the associated personal profiles, requires an understanding of the methodology of moral inquiry. Elsewhere, I have offered some proposals about proper methods of moral inquiry.[20] On the view I suggest, moral inquiry is primarily social. Societies make moral progress through identifying problems with the currently accepted moral framework and solving (more exactly: partially solving) those problems. A satisfactory moral methodology must answer two

19. I'm inclined to speculate that, for the vast majority of people, the principal deficiency is thoughtlessness, the tendency to charge ahead when reflection is required. It also seems likely that this type of mistake causes greater moral harm. But these conjectures call for serious empirical research. So far as I know, nobody has yet done it.

20. For an early version, see Philip Kitcher, *The Ethical Project* (Cambridge, MA: Harvard University Press, 2011), chapter 9. My thoughts about moral methodology are more systematically developed in *Moral Progress* (New York: Oxford University Press, 2021).

main questions: How are problems discerned? And, How are they then properly addressed? For the sake of simplicity, I'll focus here on one particular way of recognizing problems and of tackling them.[21]

Historically, moral advances have been slow, messy, and vulnerable to many contingent factors. (The movement to abolish slavery, the fight for greater opportunities for women, and the overcoming of homophobia testify to these features of progressive transitions.) They have often been initiated when the sufferings of a conventionally marginalized group become a matter for wider discussion. William James was far too optimistic in declaring that if the philosopher 'makes a bad mistake the cries of the wounded will soon inform him of the fact'.[22] Sometimes, however, the cries of the wounded do arouse sufficient social concern to spark a debate and, usually after long and bloody struggles, most members of the relevant society are prepared to admit that a mistake has been made. When moral methodology focuses on such episodes, the task is to recognize the ways in which successful resolution was eventually achieved and to propose a procedure for moral inquiry capable of streamlining the process.

As I have argued, human moral life has a long history, measured in tens of thousands of years.[23] It results from a central problem in the human condition: our need to live in societies mixed by age and sex while still lacking a full psychological capacity for responding to our fellows. We are able on occasion to recognize the goals and aspirations of others and to modify our own actions so that they harmonize. Yet this ability frequently breaks down, and we thwart the intentions of people with whom we causally interact. The moral project amplifies our responsiveness. The shortcomings of our evolved psychology are partially remedied by the social working out of accepted patterns of conduct.

The major historical examples of moral progress reveal how this can occur. What is needed is the clear representation of the ways in which hitherto accepted practices bear on the lives of different classes of people, followed by a sympathetic response to the various perspectives and predicaments. In principle, that can be done without fighting a civil war, or campaigning against monstrous 'men-women', or stoning gays. Identify moral problems by taking the protests that arise seriously enough to investigate them; conduct the investigation by assembling a body inclusive enough to represent all those who are causally affected; let that body deliberate, using only information well-supported by the

21. *Moral Progress* distinguishes two generic cases: problems of exclusion and problems of false consciousness. Here, I shall only consider the former.

22. William James, 'The Moral Philosopher and the Moral Life', in *The Will to Believe and Other Essays in Popular Philosophy* (Cambridge MA: Harvard University Press, 1979), p. 158.

23. See Kitcher, *The Ethical Project*.

available evidence and committed to finding a conclusion all can live with. If the verdict takes the complaint to be justified, expand the body by including representatives of all groups potentially affected by the various options for amending the status quo. Let this new panel of deliberators discuss further, until they can arrive at a modification all can tolerate.[24]

How does that social method help the individual? How should each of us proceed when the moral resources acquired in our development seem inadequate to the current situation? None of us can snap our fingers and assemble appropriately constituted advisory boards to counsel us. What we can manage, however, is a simulation of how we imagine a properly conducted social inquiry would go. The adequacy of our imaginative attempt will depend on a number of sensitivities. First, our ability to discern the range of options available to us. Second, an ability to identify the kinds of people who would be affected by each of the alternatives. Third, an understanding of their various perspectives. Fourth, some sense of how they might respond to one another's needs and aspirations, if they were firmly committed to leaving nobody unsatisfied. The first two of these sensitivities are primarily cognitive, depending on talents for recognizing causal structure. The third and fourth introduce affective dimensions, requiring skills in empathizing with the viewpoints and circumstances of others.

An adequate moral psychology would analyze the four sensitivities into more basic elements, thus enabling empirical investigations of the kinds of profiles individuals (and groups) come to have. Even in advance of a more sophisticated account, it's possible to extend Cocking and van den Hoven's picture of moral fog by differentiating a number of ways in which an internet user's moral compass may break down. As already noted, one class of instances of online evil may come about through thoughtlessness: seated at the computer, the wrongdoer just follows impulse or habit, without pausing to examine their conduct. Others check their behavior. They ask if what they were going to do is OK—and, after some thought, they decide to go ahead. Although they have engaged in moral inquiry, it has gone badly.

The trouble may occur in any number of ways. Here are some principal varieties.

(1) Choosing the wrong class of "deliberators": instead of focusing on those affected, people consider the opinions of their friends, gaining reassurance from the thought that 'everyone does it'.

24. This paragraph compresses an approach I defend at length in *Moral Progress*. Many qualifications and nuances have been omitted. But I hope a rough précis provides enough to explain my version of the methodological project.

(2) Misrepresenting the impact on others: adverse effects on particular people are minimized; 'it's just a prank (they should be able to take a joke)'.
(3) Failure to represent the situation of others or to empathize with it: this may occur as a general insensitivity to others' feelings or as hostility toward members of a particular group ('they deserve it—they're losers'[25]).
(4) Overestimation of the value expressed in the action: while a negative impact on some others is acknowledged, it is seen as outweighed ('the whole point of the web is to let people feel free to be themselves').

There are surely other scenarios, and those I've briefly noted can combine and interact. A moral compass—like its everyday namesake—can let the user down for all sorts of reasons. Even if the prankster, the cyberbully, the poster of revenge porn, and the cruel hacker are all lost in moral fog, it would be useful to understand the various processes through which they have lost their way.

V

Useful because understanding the different etiologies might help decrease the incidence of online evil. If my general approach to moral life is right, moral training is always likely to be incomplete. Societies discover methods to educate their members so that most of them have sensitivities adapted to most of the circumstances they encounter. Moral mistakes occur most frequently when the situation poses a novel challenge. Unprepared for this type of context, the agent's acquired resources prove inadequate. Online evil results from a spectacular change in the frequency with which the previously cultivated sensitivities are not up to the job—at least for a significant number of people. What we would like to do is to reform moral education so that fewer Internet users engage in the practices Cocking and van den Hoven so powerfully describe. That requires (in my view) recognizing the sensitivities needed for success in moral navigation of the internet world and devising educational programs for instilling them. (Although we shouldn't expect any such program to anticipate future technological changes and the new predicaments they generate.)

25. This is the best rationalization I can come up with for the case of the people who hacked into the Epilepsy Help site. But perhaps this judgment reflects the limits of my own moral imagination.

Why do good people do appalling things? The question arose in response to the behavior of ordinary citizens under Nazism and it was intensified by Milgram's notorious experiments. Cocking and van den Hoven reasonably devote space to considering the ensuing discussion. In the end, however, it seems to me right to abandon the question. People are variously sensitive. Internet usage exposes differences we might previously not have suspected. Individuals whose moral status is indistinguishable outside their time at the keyboard are sharply differentiated by what they do online. In contexts beyond those considered in their early training, some are better able to appreciate causal impact, better at imagining the reactions of the affected groups, more empathetic, and so forth. So far as we fix the class of challenging situations by including just those currently experienced, they turn out to be better people. If we concentrated on the predicaments encountered in the world in which they grew up, that judgment would not hold. And—who knows?—in some range of environments arising in their future, they might turn out to be worse. If blame has a role in responding to the online evildoers, it should be by teaching them—and others potentially like them—how to do better.

I hope it is clear how dismissive judgments of applied ethics are completely wrongheaded. As the world changes, our species constantly needs adjustments of its moral practices to keep pace. *Online Evil* begins an important enterprise, that of understanding how the sensitivities of ordinary people are often morally inadequate in a world transformed by internet technology. We should all be grateful to Cocking and van den Hoven for having initiated a crucial project in moral and educational reform.