

The Possibility of an Ongoing Moral Catastrophe

Evan G. Williams

Accepted: 28 January 2015 / Published online: 6 February 2015
© Springer Science+Business Media Dordrecht 2015

Abstract This article gives two arguments for believing that our society is unknowingly guilty of serious, large-scale wrongdoing. First is an inductive argument: most other societies, in history and in the world today, have been unknowingly guilty of serious wrongdoing, so ours probably is too. Second is a disjunctive argument: there are a large number of distinct ways in which our practices could turn out to be horribly wrong, so even if no particular hypothesized moral mistake strikes us as very likely, the disjunction of all such mistakes should receive significant credence. The article then discusses what our society should do in light of the likelihood that we are doing something seriously wrong: we should regard intellectual progress, of the sort that will allow us to find and correct our moral mistakes as soon as possible, as an urgent moral priority rather than as a mere luxury; and we should also consider it important to save resources and cultivate flexibility, so that when the time comes to change our policies we will be able to do so quickly and smoothly.

Keywords Hedging · Moral mistakes · Moral uncertainty · Progress

1 Introduction

Consider the most terrible manmade disasters of history, such as institutionalized slavery or the Holocaust. Certainly they were catastrophic for their victims, who lost their freedom or their lives, or had to make horrifying compromises in order to survive. However, the disasters were also catastrophic, in a different way, for their perpetrators. While the victims lost their freedom, their lives, or even their innocence, the perpetrators lost something arguably even more precious: I am tempted to say “their souls,” but let us settle for “their moral admirability.” They stained their hands with the blood of millions, left a legacy of grief and remorse, and are viewed by their descendants as a source of great shame.¹

¹I have inherited a family heirloom, a handbell with an attached note identifying it as: “Bell used to call slaves to meals before Civil War.” I do not know exactly how it came to be in my possession, but I fear the worst—and wonder what else I may have inherited from the same tainted source.

E. G. Williams (✉)
Department of Philosophy, University of Wisconsin Oshkosh, 800 Algoma Blvd, Oshkosh, WI 54901, USA
e-mail: william@uwosh.edu

E. G. Williams
e-mail: evangreggwilliams@gmail.com

I shall say that the societies which perpetrated those great manmade disasters suffered *moral* catastrophes, in contrast to the material catastrophes suffered by their victims. For my purposes, moral wrongdoing counts as *catastrophic* when three elements are present. First, it must be *serious* wrongdoing: for example, in the case of actions which are wrong by virtue of harming people—not necessarily the only type of seriously-wrong action, but certainly *a* type—the harm must be something closer to death or slavery than to mere insult or inconvenience. Second, the wrongdoing must be *large-scale*; a single wrongful execution, although certainly tragic, is not the same league as the slaughter of millions. Third, responsibility for the wrongdoing must also be widespread, touching many members of society. For example, the blame for slavery and the Holocaust does not rest merely on slave-owners and Nazi bigwigs; it also stains non-slave-owning Americans who consumed the products of slave labor, non-genocidal Germans who supported the Nazis for reasons of economics or patriotism, and indeed anyone who failed to oppose the evils as actively as he should have.² We would never wish such a fate for ourselves or our loved ones—in fact, I hope I speak for most of my readers when I say that we would give almost anything to avert it.

In Part One of this article, I will argue that it is quite possible that our society is presently suffering a moral catastrophe. My argument is non-constructive: I argue that one or another of our society's policies or institutions will likely turn out to be catastrophically immoral, but I do not identify *which* policy or institution it is. Part Two will then be an exercise in *applied moral uncertainty*: I argue that the likelihood that we are suffering a moral catastrophe has tangible policy implications even *before* the catastrophe has been identified, since what we do now will affect how quickly the catastrophe is brought to a halt. The more likely it is that we are making a major moral mistake, the more we should be prepared to sacrifice to permit change to come about sooner rather than later.

2 Two Reasons to Think We are Suffering a Moral Catastrophe

The first question is: how likely is it that we—by which I mean liberal Westerners, and in particular Americans—are currently suffering a moral catastrophe? I think there are two major reasons for concern: call them the Inductive Worry and the Disjunctive Worry.

2.1 The Inductive Worry

Is it possible to be unknowingly complicit in great evil—i.e., to commit grave moral transgressions in spite of acting in accordance with the moral views of oneself and one's society? Some philosophers would say “no.” For example, a moral relativist or non-realist would deny that there is any objective morality to transgress against: for them, individual and social moral views are the only kind of morality that can meaningfully be discussed. More subtly, one could also hold that there are objective moral facts but that, *in fact*, innocent mistakes are always excused: that if the information available to us did not suffice to reveal a moral obligation's existence, morality never condemns us for failing to meet that obligation.³

However, I am inclined to reject all of these views, and to say that it *is* possible to act wrongly without knowing it. Suppose, plausibly enough, that some Nazis sincerely believed—not out of cynicism, willful ignorance, or anything like that, but simply as a product of their

² Reflect on Thoreau (1849)'s statement: “Under a government which imprisons unjustly, the true place for a just man is also a prison.”

³ For example, I take Zimmerman (2008) to be defending such a view in sect. 4.2, pp. 173–193.

tragically-skewed life experiences—that they were doing the right thing and that their actions would somehow improve the human condition in the long run.⁴ While such sincerity might perhaps mitigate the amount of punishment deserved for their actions, I feel that it could not change the brute fact that their lives were, morally speaking, spectacularly unsuccessful. Call it an innocent mistake, bad luck, or what-have-you; the danger of living a life like theirs is to be faced with trepidation, not equanimity.

Incidentally, I consider this to be the case regardless of whether their mistakes were, at root, moral or factual: someone who advocated genociding Jews because he had the false moral belief that genocide was the just and proper punishment for the “crime” of rejecting Jesus was not particularly more guilty than someone who advocated genociding Jews because he had the false factual belief that all Jews belonged to a conspiracy aimed at overthrowing civilization. Both anti-Semites were engaged in the same morally-catastrophic project.

So—henceforth this shall be an assumption of my article, since my remarks are intended primarily for my fellow moral realists—it is possible to suffer a moral catastrophe unknowingly. Hence it is possible that *we*, even those of us whose roads are paved entirely with good intentions, are suffering one as well. But I want to argue for a stronger claim than merely “it is possible.” I am not a historian, so this is largely just my own impression, but: I suspect that it is *common* for the doers of grave wrongs to view their own behavior in a positive light. Consider the crusaders, inquisitors, and conquistadors who killed, tortured, and maimed in the name of a God whom they called *good* and *just*. Consider the soldiers and colonizers of the age of imperialism, many of whom were *proud* to be serving mighty empires and spreading “civilization”—or, for that matter, the “noblemen,” clergy, slave-owners, and so on, who were *proud* of the wealth and status signified by their lives of parasitic idleness. Consider the men in various patriarchal societies, some of which continue to exist today, who stone to death their wives or daughters, or in more subtle ways deny them a fair shot at life, while framing the oppression as a requirement of *decency* and *honor*.

On one hand, no adequate moral theory would condone such behavior. On the other hand, I find it hard to believe that there could be *that* many people in history whose expressed concern with justice, pride, and honor was mere pretense. No doubt some of the wrongdoers did recognize the contradictions inherent in their value systems, and espoused those values anyway out of cynical self-interest—but I suspect that many others genuinely believed in what they were doing. Their moral failing was not due to ill intent, but rather to horrific moral blind spots. Therefore it is possible for human beings to have horrific moral blind spots, and therefore it is possible that *we* have some too.⁵

But, again, it is not just the mere *possibility* that worries me. I think it is *probable* that we have serious blind spots. After all, just about every other society in history *has* had them. Show me one society, other than our own, that did not engage in systematic and oppressive discrimination on the basis of race, gender, religion, parentage, or other irrelevancy, that did not launch unnecessary wars or generally treat foreigners as a resource to be mercilessly exploited, and that did not sanction the torturing of criminals, witnesses, and/or POWs as a matter of course. I doubt that there is even one; certainly there are not many. Even our own society has only very recently, and incompletely, begun showing respect to all human beings: my parents grew up in an era when racial segregation was taken for granted, so if our present

⁴ For some vivid examples of how moral beliefs can lead one into evil, see Bennett (1974). Of particular pertinence is the discussion of Heinrich Himmler on pp. 127–129.

⁵ This basic point appears in Singer (1974), p. 103: “One should always be wary of talking of ‘the last remaining form of discrimination’. If we have learnt anything from the liberation movements, we should have learnt how difficult it is to be aware of latent prejudice [....]”

values are the right ones then my generation is the *first* to view those values as commonsensical rather than innovative.

Is it credible that my generation could be so special? Literally hundreds of generations have *thought* that they had the right moral values. Two thousand years ago, the Romans—the *imperialistic, crucifying, slave-owning* Romans—were congratulating themselves on being “civilized,” because unlike the “barbarians” they had abolished human sacrifice. This *was* genuine progress, but what they did not realize was that thousands of years’ additional progress remained to be made. We are in the same position: we know how much progress is embodied in our values, but not how much progress remains to be made in the future. This, then, is the Inductive Worry: most cultures have turned out to have major blind spots in their moral beliefs, and we are in much the same epistemic situation as they are, so we will probably also turn out to have major moral blind spots.

I do have to say “probably,” not “certainly.” It is entirely possible that the induction will eventually break down, that some generation *will* have the privilege of living at the exact historical moment when the true values have finally evolved. However, I think it would be naïve, even un-Copernican, to assume that ours is the one generation living at that privileged moment in history. If anything the assumption would be *more* naïve for us than it was for the Romans, since we are living in a time of accelerating social and technological change whereas they were at more of a plateau. The first generation to get everything right will most likely have parents who got *nearly* everything right—which our generation emphatically does not.

2.2 The Disjunctive Worry

The reader may be an activist, already convinced that some specific moral catastrophe is taking place, and doing everything he can to put an end to it. However, so as not to obscure my main point about *unidentified* catastrophes, I ask the reader to set *known* catastrophes aside; let him imagine that all of his favorite political causes triumph, and society becomes organized exactly as he thinks best. I hope to convince him that even in such a scenario, a moral catastrophe would *still* probably be taking place. My reason is this: there are so *many* different ways in which a society—whether our actual one or the one of the reader’s dreams—could be catastrophically wrong that it is almost impossible to get *everything* right.⁶ I shall now describe a few of those possible ways in which our policies could be seriously victimizing on the order of a million people per year, or could be otherwise wrong to an extent morally comparable, in both scale and severity, to such victimization. Do note the scale: a million people per year is roughly the number of Jews killed per year during the Holocaust. Of course, the Holocaust was not *merely* a mass murder but was evil in many other ways as well, e.g., as an attempted genocide, so I do not mean to imply that all of these possible mistakes are fully on a par with it; nevertheless I think evoking it is useful to give a sense of the severity of the worry.⁷

To begin, first note the many different *kinds* of mistakes we might be making. We could be wrong about *who* has moral standing, for example if it turns out that aborting human fetuses is morally as serious as murdering adult people. We could have correct moral beliefs about who has standing, but incorrect factual beliefs about what sort of treatment is healthy for them, for example if Dawkins (2006) turns out to be right that religious indoctrination can be so damaging to children’s minds that it qualifies as child abuse. We could be right about basic issues of harm and benefit, but wrong about other obligations: for example, suppose that duties

⁶ “It is possible to fail in many ways [...], while to succeed is possible only in one”—Aristotle (350 BCE), II.6.

⁷ For a more detailed comparison of a hypothesized present-day moral catastrophe to the Holocaust, and a defense of making such comparisons, see Szybel (2006).

of fidelity turn out to be of comparable importance to duties of non-maleficence, and observe that every year over eight hundred thousand American couples, having previously made solemn and elaborate promises to the effect that they would stay together until parted by death, break those promises by divorcing. We could be right about obligations to individuals, but wrong about the just way to weigh different individuals' concerns against each other and/or against other desiderata such as freedom—consider the debate about Affirmative Action, or for that matter the debate about what constitutes a fair level of taxation and wealth redistribution. We could be right about what is just, but wrong about how to prevent and respond to injustice: there are at least one and a half million prisoners behind bars in the United States who would be free if our incarceration rate were as low as culturally- or geographically-similar countries such as the U.K., Canada, or Mexico.⁸ We could be wrong about which morally *bad* situations are morally *wrong* for us to ignore, and which are not our responsibility: for example, there are over a billion desperately-poor or horribly-oppressed people in the world, to whom we could provide much more aid to them than we are currently offering.⁹ Or we could be wrong about the extent of our obligation to build a good future for our successors, and/or about the best way to do so—an issue with particularly high stakes, since the choices we make today might be the difference between humanity spreading out across the stars and forming billions of colonies of billions of happy people which last for millions of years, or humanity being extinguished by some natural disaster within the next few centuries.¹⁰

Within any one of these categories, horrible mistakes are possible in more than one direction. Consider the “who has moral standing?” issue. If the Pro-Lifers are correct and all human organisms—or even just those which have potential futures like our own¹¹—have as much right to life as people do, then the U.S. has tolerated the equivalent of about a million murders every year for the past four decades. On the other hand, if it turns out that what matters is something's sentience, or even mere aliveness, rather than its humanness, our treatment of billions of food animals is probably catastrophic. For example, pigs are not as bright as humans, but if they are even a hundredth as conscious as humans are, and the wrongness of killing something turns out to be proportional to its degree of consciousness, the slaughter of a hundred million pigs each year in the U.S. is a moral catastrophe.¹² Or perhaps we are lucky, and neither fetuses nor nonhumans are morally significant—in which case, consider the million or so women who accidentally become pregnant each year but choose to gestate and raise the babies anyway. They are accepting life-changing burdens, frequently allowing their career, education, and marriage plans to be disrupted. Presumably they do so because our culture has raised them to believe that it is the right thing to do. If our culture has taught them falsely or is pressuring them into accepting unnecessary burdens, we bear responsibility for the derailment of those millions of life plans.

So within any given area of moral uncertainty, there are multiple traps for a culture to stumble into—a society seeking to avoid catastrophe must sail a fine line between the Scylla of too much redistribution, punishment, foreign intervention, or whatever, and the Charybdis of

⁸ I produced this figure with a back-of-the-envelope calculation, using data from the International Centre for Prison Studies (2012).

⁹ See Singer (1972)'s classic article demanding humanitarian aid to the poor.

¹⁰ Even if our future turns out not to be quite *that* fragile, Bostrom (2005) has remarked that a one-year delay in finding a cure for old age would result in an extra number of human deaths greater than the population of Canada.

¹¹ The classic anti-abortion argument appears in Marquis (1989). In a nutshell: given that it is wrong to kill an unconscious person who would wake up later if left alone, why is it not wrong to kill an embryo which would wake up later if left alone?

¹² Worse, instead of rights being proportionate to consciousness, there might be some single threshold at which rights are achieved—and this threshold might be low enough to include pigs. For example, see the classic arguments for an “experiencing subject of a life” test in Regan (1983).

too little. I listed seven such areas above and unpacked one of them; if the reader peruses a few syllabi for applied ethics courses, he will find that my list is far from exhaustive. There is also a non-negligible possibility that we will turn out to be catastrophically wrong about some issue that nobody has even raised yet, a genuine blind spot.¹³

The large number of possible catastrophic mistakes creates a serious worry, since the *disjunction* of them must, mathematically, be viewed as much more credible than any particular such mistake. For example, if we have 95 % confidence in each of 15 different independent propositions—that legal abortion is not morally catastrophic, that the high incarceration rate is not morally catastrophic, that the world poverty problem is not morally catastrophic, and so on—then the appropriate level of confidence to place in the claim that we are not catastrophically wrong about *any* of the items on the list will be 95 % raised to the power of 15, which is to say about 46 %. We could be confident in each individual policy, while still finding it more likely than not that one or another would turn out to be wrong.

In fact, those figures are very conservative. I suspect that a full list of ways in which our practices could be catastrophically wrong would have far more than 15 items on it; some would have probability much higher than 5 %; and many of the possible catastrophes—e.g., the possibility that we are punishing criminals more harshly than they deserve and the possibility that we are punishing them more leniently than their victims deserve—would be mutually exclusive rather than merely independent. So even if we were reasonably confident in the moral acceptability of each individual practice that we support, it would nevertheless also be appropriate to regard it as *highly probable* that at least one of those practices is catastrophically wrong. The only way to escape this reasoning would be to postulate some *systematic* reason why moral catastrophes are impossible, e.g., by endorsing moral non-realism, and thereby to defeat all of the disaster scenarios at once. But since I do not find it credible to claim that moral catastrophes are impossible—see my earlier discussion of the moral unacceptability of sincere Nazism—the rest of this paper will proceed under the supposition that we *are* suffering some unidentified moral catastrophe.

3 What Should We Do About It?

So, let us stipulate that we are suffering an unspecified moral catastrophe. What should we do about it? Well, obviously, we need to compensate our victims, or at least stop continuing to victimize them. But how can we do this when we do not know *which* of our policies is the catastrophic one, nor *who* its victims are?

When faced with uncertainty, a natural approach is to “hedge” against individual risks. For example, if one is driving down the highway and is unsure whether one has enough gas to reach one’s destination, it makes sense to go ahead and refill one’s tank when the opportunity presents itself, rather than to risk running out partway and becoming stranded. A morally-charged equivalent of this would be for someone who is unsure whether nonhuman animals

¹³ For an example of a possible catastrophe which the reader has probably never even contemplated, consider the possibility that the function of the corpus callosum in a human brain is not to unite the two hemispheres in the production of a single consciousness, but rather to allow a dominant consciousness situated in one hemisphere to issue orders to, and receive information from, a subordinate consciousness situated in the other hemisphere. For the classic philosophical discussion of this two-minds-in-one-body idea, see Nagel (1971); for a review of the neurological evidence in favor of it, see Bogen (1986). The worst-case scenario here is that there are 300 million human slaves in America, whose frustration, boredom, and oppression we have not made any effort to ameliorate. Our failure to notice their existence would perhaps be a partial excuse, but may not be completely exculpatory—especially now that the evidence discussed by Nagel and Bogen has been available for decades.

have moral significance to adopt a vegan lifestyle to avoid the *risk* of sharing responsibility for their exploitation.¹⁴ This can make sense regardless of whether one's uncertainty is, at heart, factual or moral: e.g., whether it is due to being unsure what degree of sentience the animals in question have, or due to being unsure what degree is necessary to qualify as an object of significant moral concern.¹⁵

However, the strategy of hedging against each particular risk faces serious limits. Firstly, as discussed above, on many issues there are multiple—opposing—ways to be catastrophically wrong. For example, it is possible that we are fighting too many foreign wars and causing unjustified bloodshed, but it is also possible that we are fighting too few and failing in a duty to rescue oppressed peoples from tyrants. There is no morally-safe option on that sort of issue. Relatedly, hedging against any given possible catastrophe is expensive and to some extent precludes hedging against other possible catastrophes: we could do as much as possible to conserve the environment, or we could do as much as possible to help the world's starving masses, but we cannot do both since conserving the environment will cost us resources that could have been used for feeding the starving. Lastly, we must not forget the “we are catastrophically wrong about some issue that we have not even considered” possibility; how could one hope to hedge against that?

What is needed is a way to deal with multiple possible catastrophes at once, rather than approaching each individually—a way to hedge against a *generic* moral catastrophe rather than specific ones. To see how this could work, consider how any possible catastrophe on the list, if it indeed turned out to be catastrophic, would have to be resolved. Regardless of the details of the catastrophe, two major elements will be necessary: first, we need to be able to *recognize* what we are doing wrong; and second, we need to be able to *implement* the appropriate changes in order to stop doing it. There are things that we can do now, before we know what the catastrophe is, to make each of these steps easier.

3.1 Recognition of Wrongdoing

If a moral catastrophe is taking place, the first step to putting an end to it is for us to recognize it as morally bad. Otherwise we would simply be in the position of making random social changes that were as likely to be for the worse as for the better. I suppose we could adopt a strategy of trying to minimize our impact on the world: engage only in activities absolutely necessary for survival, keep those activities as non-diverse as possible, avoid actions with long-term repercussions—even deliberately allow ourselves to go extinct. Doing these things, minimizing our impact on the world, might reduce our chances of being guilty of any analog of murder or slavery. However, some possible catastrophes involve violations of positive duties rather than negative ones; if we avoided impacting the world, then in particular we would not be making the world *better*. Also, while living such narrow and cautious lives, we would not be making any progress that would allow us to improve our behavior in the future—our children would be as ignorant as we are. So this is a recipe for *eternally* failing to fulfill our positive responsibilities.

A better plan, one which in the long run would make it possible for our society to fulfill both its negative obligations *and* its positive ones, is to try to *figure out* what catastrophic wrongs we are committing. Once we have identified the moral catastrophe that we are

¹⁴ Singer (1993) takes this line on p. 119: “if you see something moving in the bushes and are not sure if it is a deer of a hunter, don't shoot!”

¹⁵ For a discussion of how to formulate hedging in the case of moral uncertainty, and of some of the problems that can arise while trying to do so, see Lockhart (2000), ch. 4.

suffering, we will be able to correct our policies in an intelligent manner, rather than stumbling around blindly.

Identifying our mistakes will require progress not just in moral theory, but also in the scientific fields that allow moral theory to be applied: for example, if it is true that it is always wrong to kill self-aware beings unnecessarily, discovering this moral fact would be important—but so would, say, discovering that dolphins are self-aware, discovering that certain types of fishing nets kill dolphins, and discovering a way to feed people without using those nets. So I am not just making a narrow plea for philosophical progress here; I suspect that a great deal of human intellectual activity—most science, some engineering, even a significant chunk of art and literature—can be described in terms of “improving our understanding of what options we have and what those options involve,” and so all of it is valuable to the project of someday identifying and choosing non-catastrophic options. That said, philosophical progress is definitely an essential part: increased technical knowledge, if unaccompanied by increased moral wisdom, can be seriously harmful—imagine what would have happened if Genghis Khan had known how to build nuclear missiles.¹⁶

I should also clarify what it means for “us” to make intellectual progress. It is not sufficient that each relevant bit of information come to be known by someone, somewhere. Returning to the dolphin example: if a philosopher knew that conscious beings ought to be protected, and a scientist knew that dolphins were conscious, and an engineer knew how to design fishing nets that would not harm dolphins—no good would be achieved unless they combined those bits of knowledge with one another. So it is important not just that there be progress within each discipline, but also that there be dialog between disciplines. Furthermore, even if the intellectuals talk to one another and put together all of the pieces of the puzzle, it still does no good unless the conclusion they reach can then be transferred to the fishermen, or at least to some legislative body with the power to regulate the fishermen.

In short, we need a marketplace of ideas in which *good* ideas—correct moral principles, true scientific facts, useful inventions, etc., whichever those turn out to be—will outcompete bad ones and eventually spread across the entire population, or at least the portion of the population involved in social decision-making. I fear that our society does not presently possess such a marketplace. This is perhaps most obvious in the case of scientific theories: despite near consensus among experts, a very significant proportion of laypeople—and of the elected legislature—continue to disbelieve scientific claims such as evolution by natural selection or anthropogenic climate change, and continue to believe debunked claims such as the alleged link between childhood vaccination and autism. Undoubtedly we would see the same with respect to *moral* theories, were it not for the fact that we have been unable to achieve a consensus even among experts.

This is something we can work on *now*. We do not know exactly which new discovery will prove crucial for ending the present moral catastrophe, but we do not *need* to know that in order to recognize that the catastrophe will end sooner if we can create a social environment in which new discoveries of every kind—but especially discoveries in ethics—can be made and publicized. If there are features of our society which are making it harder for new ideas to arise, for good ones to distinguish themselves from bad ones, and/or for the good ones to then gain wide currency, those features must be identified and changed.

A detailed examination of *how* to maximize our society’s rate of intellectual progress is outside the scope of this paper. Presumably it would involve a variety of reforms to education, media, and perhaps even the basic social hierarchy.¹⁷ What I want to emphasize here, however,

¹⁶ For a more extensive survey of ways in which some types of knowledge can be dangerous, see Bostrom (2011).

¹⁷ For an in-depth examination of what policies might hasten moral progress, see Buchanan (2002).

is the *importance* of building a rapidly-progressing society. It was worth fighting wars to end institutionalized slavery or to stop the Holocaust; since the unrecognized catastrophe is likely to be of comparable scope and severity, also involving the victimization of millions or something morally comparable to such victimization, halting it *also* warrants substantial, war-scale sacrifice. Call it a War on Our Own Backwardness. For example, measures to contemplate, if we had reason to believe that they would be effective at hastening progress, could include: accepting more austere lifestyles, so that more resources could be directed into science and education; accepting more compromises of privacy, to facilitate collection of social and psychological data; refocusing our medical system on keeping people productive for as long as possible, rather than on prolonging the lives and/or meliorating the suffering of the elderly or dying; refocusing our educational system on identifying the most talented children and nurturing them to their full potential, not on trying to shepherd every last mediocre student into a four-year college; and so on.

To be clear: I would not recommend such measures if I were merely weighing the *recognized* benefits of intellectual progress against the value of prosperity, privacy, life, prevention of suffering, equality, and so on. But I have argued in this paper that the *unrecognized* benefits of progress—the social changes which we have not yet identified as necessary, but which we will be made once we understand their importance—are morally very urgent. That alters the equation. Antebellum Southerners were not wrong when they complained that abolishing slavery would damage their economy and disrupt their way of life—what they were wrong about was thinking that those *genuine* costs could outweigh the importance of halting their ongoing moral catastrophe.

3.2 Implementation of Improved Values

After we figure out what we are doing wrong, we will still need to fix it: change our laws and norms in whatever ways are necessary to end the moral catastrophe, and then make any restitution or penance which is appropriate. The specific changes that will be necessary will depend on *which* catastrophe we are suffering, but even now—before we have identified our mistake—we can start laying the groundwork for those changes. We can, and should, try to build a flexible society which is capable of taking *whatever* actions it deems necessary, so that the catastrophe can be resolved as soon as possible after being identified.

There are many ways for a society to fail to be flexible, such that even after it discovers that it ought to make a change, it will be unable to make it in a timely manner. We should seek to avoid those failures.

One failure is to be impoverished—a society too busy struggling for survival to be able to afford much else will also tend to be inflexible. Social changes almost always involve transition costs: dismantling the institutions that were immoral, building replacement institutions in their stead. For example, if we discovered that an “animal rights” scenario were true and decided to transition to a vegan society, that would entail massive disruptions as the half-million workers in the meat-processing industry retrained for other jobs, farmers shifted their planting from feed crops to vegetable crops, households and restaurants adapted their menus to comply with the new rules while still providing adequate nutrition, and so on. Even if vegan agriculture turned out to be more efficient, in the long run, than animal-based agriculture, the transition could still only be made by a society that was willing and able to absorb short-term losses. The same is true for the other kinds of changes we might need to make. Of particular note: if the changes that need to be made include paying compensation to some group which has been wronged, a prosperous society will also be in a better position to pay it. Just as we stockpile strategic resources against hypothetical military emergencies, we should also be

saving up for hypothetical *moral* emergencies, rather than running up huge debts on the assumption that we will be able to pay them off in a rosy future.

Inflexibility may not always be felt as poverty. For example, many American towns have been designed so as to be navigable only by car: homes and businesses are too spread out to make walking practical; infrastructure for pedestrians and cyclists is absent; residential streets are arranged into cul-de-sacs that no bus could efficiently traverse; etc. Since everybody in those towns has a car, there is no problem. But if we were to discover a need to abolish private automobiles—e.g., due to safety, energy conservation, pollution, and/or land use considerations—these towns would have to be almost totally rebuilt. In light of the possibility of an ongoing moral catastrophe, this sort of inflexibility should be regarded as a critical design flaw, and avoided in the future as we build new cities and/or renovate old ones. Indeed, when designing *anything*, from a large city to a small consumer appliance, flexibility of use and ease of modification should be seen as major desiderata—which must be weighed alongside, and will sometimes outweigh, other important desiderata such as environmental impact and public safety.

Even if society is materially capable of making a morally necessary change, its decision-making apparatus might make that change easier or harder. Obviously a society dominated by some foreign power, or by some despotic government whose only aim is to remain in power, will not be in a position to restructure itself in reaction to moral revelations. The same is also true of a society with an overly-rigid Constitution or other basic law. For an example of what *not* to do, consider the “Corwin Amendment” to the U.S. Constitution. Passed by Congress in 1861, if it had been ratified by the States then it would have denied the federal government the power to abolish slavery, and altered the amendment process to prevent any further changes on the matter (Bryant 2003). It is understandable why Congress wanted to do this—states were seceding out of fear that slavery would be abolished, so making it impossible to abolish slavery was a move that could have prevented the Civil War—but in retrospect that would have been too high a price to pay for peace. Even the pro-slavery advocates should have recognized that they *might* turn out to be wrong, that slavery *might* be a moral catastrophe, and that fighting a disastrous war was preferable to irrevocably locking the country into a policy that might later have turned out to be morally abhorrent.

Even those Constitutional amendments which do not alter the amendments process should still be viewed with suspicion. For example, one wonders what Congress was thinking when it passed the Eighteenth Amendment establishing Prohibition. Instead of merely giving the federal government the *power* to control alcohol via legislation, the amendment *itself* prohibited alcohol. This amounts to, “Even if a majority of both houses and the President—who will know more about the effects of Prohibition than we do, since they will have experienced it first-hand—later come to believe that Prohibition is a mistake, we think they should be stuck with it anyway until they can muster a supermajority sufficient for a new amendment.” What hubris they must have had, to believe that they were so much wiser than those better-informed counterparts would be! I would recommend that, rather than trying to make their policies as hard to change as possible, legislatures instead try to make it easy: e.g., by including sunset clauses that will cause automatic policy shifts unless people who have lived under those policies, for enough time to gauge their effectiveness, make an affirmative choice to keep them in operation.

Ultimately, of course, the question of *how* to build a flexible society—like the question of how to build a rapidly-progressing society—is one for sociologists and is outside the scope of this paper. My comments above are meant only as examples of the kinds of things we should be considering. The essential philosophical point is simply that building a flexible society is *important* and *urgent*, in much the way that stopping the Nazis was important and urgent. The

more flexible we are, the sooner any current Holocaust-scale moral catastrophe will be halted after being identified.

4 Concluding Remarks

I have argued that it is possible, and indeed probable, that we are suffering a severe, ongoing moral catastrophe. Our descendants may well view us with the same repugnance as we view our slave-owning forefathers. I have also suggested that the fastest way to end this catastrophe, so that our wrongdoing will stop sooner rather than later, is to build a society that makes rapid intellectual progress and is flexible enough to take decisive action when the need is recognized. It may be too late to save our own moral admirability, but perhaps it is not too late to save our children's or grandchildren's.

On that note, it should be understood that setting up the prerequisites for ending an ongoing catastrophe is not the same thing as actually ending it. It is not as though we can make ourselves blameless simply by putting society on track to recognize and correct its errors. The errors must actually *be* corrected before the ongoing moral catastrophe will be over. So there is a bit of a consequentialist premise embedded in my argument for building the kind of society I think we should build. Our reason to do it is not that doing so is morally right in itself, but rather than it will foreseeably lead to moral improvement in the future. Indeed, if we take extreme enough measures for the sake of progress and flexibility—e.g., at the cost of some degree of freedom and happiness—those very measures may turn out to be among those needing future reversal. This risk notwithstanding, I think such measures may be justified, in the short run, by the need to deal with other ongoing moral catastrophes.

Despite its consequentialist, ends-justify-the-means type nature, my claim, that we ought presently to take actions which will foreseeably result in better behavior over the long run, should not be *too* controversial. Consider the case of a person who is gradually sinking into mental illness. Suppose that he is experiencing homicidal ideation, and knows that it is getting worse over time; he foresees that he will commit a terrible murder unless he seeks psychiatric help. Nevertheless, suppose that his present options are to seek psychiatric help or to ignore his problem and go shopping. A strict deontologist could reason that murder is wrong but shopping is not: if the agent goes shopping, foreseeing that this will lead to him committing a murder in the future, he has still not done anything wrong until he actually commits the murder. However, I think most of us—even most deontologists—believe that there is an obligation, not just to refrain from taking wrongful actions in the present, but also to try to prevent oneself from taking wrongful actions in the future. And that is the sort of obligation that grounds my discussion here: we should build the kind of society I have described, not because doing so is necessary or sufficient for present moral adequacy, but rather because it will foreseeably make our future behavior much less wrong.

From this assumption, that taking the necessary steps to prevent future wrongdoing is of comparable importance to avoiding present wrongdoing, it follows that it is *very* important than any ongoing moral catastrophe we are suffering be brought to an end sooner rather than later. Putting ourselves on a path to make rapid progress, and to take decisive action once we deem such action necessary, should be a major social priority.

Acknowledgments I am especially grateful to Larry Temkin for his many helpful suggestions on this and related material. I also wish to thank Jeff McMahan, the Future of Humanity Institute at Oxford, the various audiences which have given me feedback on this article, and an anonymous reviewer.

References

- Aristotle (350 BCE) *Nicomachean ethics*. Trans by D. P. Chase (1911)
- Bennett J (1974) The conscience of huckleberry Finn. *Philosophy* 49(188):123–134
- Bogen J (1986) Mental duality in the intact brain. *Bull Clin Neurosci* 51:3–29
- Bostrom N (2005) The fable of the dragon-tyrant. *J Med Ethics* 31(5):273–277
- Bostrom N (2011) Information hazards: a typology of potential harms from knowledge. *Rev Contemp Philos* 10: 44–79
- Bryant AC (2003) Stopping time: the pro-slavery and “irrevocable” thirteenth amendment. *Harv J Law Publ Policy* 26(2):501–549
- Buchanan A (2002) Social moral epistemology. *Soc Philos Policy* 19(2):126–152
- Dawkins R (2006) *The God delusion*. Houghton Mifflin, New York
- International Centre for Prison Studies (2012) World prison brief. <http://www.prisonstudies.org/world-prison-brief>. Accessed 20 January 2014
- Lockhart T (2000) *Moral uncertainty and its consequences*. Oxford University Press, New York
- Marquis D (1989) Why abortion is immoral. *J Philos* 86(4):183–202
- Nagel T (1971) Brain bisection and the unity of consciousness. *Synthese* 22(3):396–413
- Regan T (1983) *The case for animal rights*. University of California Press, Berkeley
- Singer P (1972) Famine, affluence, and morality. *Philos Publ Aff* 1(3):229–243
- Singer P (1974) All animals are equal. *Philos Exch* 1(5):103–116
- Singer P (1993) *Practical ethics*. Cambridge University Press, New York
- Szybel D (2006) Can the treatment of animals be compared to the holocaust? *Ethics Environ* 11(1):97–132
- Thoreau HD (1849) *Resistance to civil government*
- Zimmerman M (2008) *Living with uncertainty: the moral significance of ignorance*. Cambridge University Press, New York