
MORALITY:
THE BAD NEWS

IF YOU ARE LIKE MOST PEOPLE CONTEMPLATING scientism, it's the persistent questions about morality and mortality that grip you, not trivial topics like the nature of reality, the purpose of the universe, or the inevitability of natural selection.

After all, the trouble most people have with atheism is that if they really thought there were no God, human life would no longer have any value, they wouldn't have much reason to go on living, and even less reason to be decent people. The most persistent questions atheists get asked by theists are these two: In a world you think is devoid of purpose, why do you bother getting up in the morning? And in such a world, what stops you from cutting all the moral corners you can? This chapter and the next deal with the second question. The first one, about the meaning of our lives, we'll take up in Chapter 9.

Scientism seems to make the unavoidable questions about morality even more urgent. In a world where physics fixes all the facts, it's hard to see how there could be room for moral facts. In a universe headed for its own heat death, there is no cosmic value to human life, your own or anyone else's. Why bother to be good?

We need to answer these questions. But we should also worry about the public relations nightmare for scientism produced by the answer theists try to foist on scientism. The militant exponents of the higher superstitions say that scientism has no room for morality and can't even condemn the wrongdoing of a monster like Hitler. Religious people especially argue that we cannot really have any values—things we stand up for just because they are right—and that we are not to be trusted to be good when we can get away with something. They complain that our worldview has no moral compass. These charges get redoubled once theists see how big a role Darwinian natural selection plays in scientism's view of reality. Many of the most vocal people who have taken sides against this scientific theory (for instance, the founders of the Discovery Institute, which advocates “intelligent design”) have frankly done so because they think it's morally dangerous, not because it lacks evidence. If Darwinism is true, then anything goes!

You might think that we have to resist these conclusions or else we'll never get people to agree with us. Most people really do accept morality as a constraint on their conduct. The few who might agree privately with Darwinism about morality won't do so publicly because of the deep unpopularity of these views. “Anything goes” is nihilism, and nihilism has a bad name.

There is good news and bad news. The bad news first: We need to face the fact that nihilism is true. But there is good news here, too, and it's probably good enough to swamp most of the bad news about nihilism. The good news is that almost all of us, no matter what our scientific, scientific, or theological beliefs, are committed to the same basic morality and values. The difference between the vast number of good people and the small number of bad ones isn't a matter of whether they believe in God or not. It's a difference no minister, imam, vicar, priest, or rabbi can do much about. Certainly,

telling people lies about what will happen to them after they die has never done much to solve the problem of morally bad people. In addition to not working, it turns out not to be necessary either. By the same token, adopting nihilism as it applies to morality is not going to have any impact on anyone's conduct. Including ours.

There is really one bit of bad news that remains to trouble scientism. We have to acknowledge (to ourselves, at least) that many questions we want the "right" answers to just don't have any. These are questions about the morality of stem-cell research or abortion or affirmative action or gay marriage or our obligations to future generations. Many enlightened people, including many scientists, think that reasonable people can eventually find the right answers to such questions. Alas, it will turn out that all anyone can really find are the answers that they like. The same goes for those who disagree with them. Real moral disputes can be ended in lots of ways: by voting, by decree, by fatigue of the disputants, by the force of example that changes social mores. But they can never really be resolved by finding the correct answers. There are none.

WHO ARE YOU CALLING A NIHILIST, ANYWAY?

Nihilism was a word that was thrown around a lot in the nineteenth century. It often labeled bomb-throwing anarchists—think of Joseph Conrad's *Secret Agent* and people like Friedrich Nietzsche who rejected conventional morality. Today, nobody uses the word much. We can use it for a label, though we need to sharpen it up a bit, given the misuses to which it has also been put.

What exactly is nihilism? It's a good idea to start with what it isn't. Nihilism is not moral relativism. It doesn't hold that there are lots of equally good or equally right moral codes, each one appropriate to its particular ethnic group, culture, time period, or ecological niche. Nihilism doesn't agree with the relativist that capital punishment is okay at some times and places and not okay at other ones. It's also not moral skepticism, forever unsure about which among competing moral claims is right. Moral skepticism holds that capital punishment is definitely right or definitely wrong, but alas we can't ever know which.

Nor does nihilism claim that everything is permitted, that nothing is forbidden. Still less does it hold that destructive behavior is obligatory, that figurative or literal bomb throwing is good or makes the world better in any way at all. These are all charges made against nihilism, or at least they are often attributed to people the establishment wants to call nihilists. But they don't stick to nihilism, or at least not to our brand.

These charges fail because nihilism is deeper and more fundamental than any moral code. Nihilism is not in competition with other codes of moral conduct about what is morally permitted, forbidden, and obligatory. Nor does it disagree with other conceptions of value or goodness—what is good or the best, what end we ought to seek as a matter of morality.

Nihilism rejects the distinction between acts that are morally permitted, morally forbidden, and morally required. Nihilism tells us not that we can't know which moral judgments are right, but that they are all wrong. More exactly, it claims, they are all based on false, groundless presuppositions. Nihilism says that the whole idea of "morally permissible" is untenable nonsense. As such, it can hardly be accused of holding that "everything is morally permissible." That, too, is untenable nonsense.

Moreover, nihilism denies that there is really any such thing as intrinsic moral value. People think that there are things that are intrinsically valuable, not just as a means to something else: human life or the ecology of the planet or the master race or elevated states of consciousness, for example. But

nothing can have that sort of intrinsic value—the very kind of value morality requires. Nihilism denies that there is anything at all that is good in itself or, for that matter, bad in itself. Therefore, nihilism can't be accused of advocating the moral goodness of, say, political violence or anything else.

Even correctly understood, there seem to be serious reasons to abstain from nihilism if we can. Here are three:

First, nihilism can't condemn Hitler, Stalin, Mao, Pol Pot, or those who fomented the Armenian genocide or the Rwandan one. If there is no such thing as "morally forbidden," then what Mohamed Atta did on September 11, 2001, was not morally forbidden. Of course, it was not permitted either. But still, don't we want to have grounds to condemn these monsters? Nihilism seems to cut that ground out from under us.

Second, if we admit to being nihilists, then people won't trust us. We won't be left alone when there is loose change around. We won't be relied on to be sure small children stay out of trouble.

Third, and worst of all, if nihilism gets any traction, society will be destroyed. We will find ourselves back in Thomas Hobbes's famous state of nature, where "the life of man is solitary, mean, nasty, brutish and short." Surely, we don't want to be nihilists if we can possibly avoid it. (Or at least, we don't want the other people around us to be nihilists.)

Scientism can't avoid nihilism. We need to make the best of it. For our own self-respect, we need to show that nihilism doesn't have the three problems just mentioned—no grounds to condemn Hitler, lots of reasons for other people to distrust us, and even reasons why no one should trust anyone else. We need to be convinced that these unacceptable outcomes are not ones that atheism and scientism are committed to. Such outcomes would be more than merely a public relations nightmare for scientism. They might prevent us from swallowing nihilism ourselves, and that would start unraveling scientism.

To avoid these outcomes, people have been searching for scientifically respectable justification of morality for least a century and a half. The trouble is that over the same 150 years or so, the reasons for nihilism have continued to mount. Both the failure to find an ethics that everyone can agree on and the scientific explanation of the origin and persistence of moral norms have made nihilism more and more plausible while remaining just as unappetizing.

PLATO'S PROBLEM FOR SERMONIZING ABOUT MORALITY

The problem of justifying morality—any morality—is a serious one for everybody—theist, atheist, agnostic. The difficulty of grounding ethics is one Plato wrote about in his very first Socratic dialogue—the *Euthyphro*. Ever since, it's been a thorn in the side of Sunday sermon writers.

In the *Euthyphro*, Plato invites us to identify our favorite moral norm—say, "Thou shalt not commit abortion" or "Homosexual acts between consenting adults are permissible." Assume that this norm—either one—is approved, sanctioned, even chosen by God.

Now, Plato says, if this norm is both the morally correct one and also the one chosen for us by God, ask yourself the following question: Is it correct because it is chosen by God, or is it chosen by God because it is correct? It has to be one or the other if religion and morality have anything to do with each other. It just can't be a coincidence that the norm is both the right one and chosen for us by God. Theism rules out coincidences like this.

Consider the first alternative: what makes the norm correct is that God chose it. If he had chosen the opposite one (a pro-choice or anti-gay norm), then the opposite one would be the morally right one.

Now consider the second alternative: God chose the pro-life or pro-gay norm, or whatever your

favorite moral norm is, because it is the right one, and he was smart enough to discern that fact about it.

Pretty much everybody, including the dyed-in-the-wool theists among us, prefers the second alternative. No one wants to admit that our most cherished moral norms are the right, correct, true morality solely because they were dictated to us by God, even a benevolent, omniscient God. After all, he is omniscient, so he knows which norm is right; and he is benevolent, so he will choose the norm he knows to be morally best for us. If it's just a matter of whatever he says goes, then he could have made the opposite norm the right one just by choosing it. That can't be right.

It must be that the norms God imposes on us were ones he chose because they are the right ones. Whatever it is that makes them right, it's some *fact* about the norms, not some fact about God. But what is that fact about the right moral norms that makes them right? All we can tell so far is that it was some fact about them that God in his omniscient wisdom was able to see.

Atheists and agnostics, too, will make common cause with the theists in seeking this right-making property of any moral norm that we all share in common (and there are many such norms). Plato's argument should convince us all that finding this right-making property of the moral norms we accept is an urgent task.

There is one way for the theist to avoid this task, but it's not one people will have much intellectual respect for. The theist can always say that identifying the right-making property of the morally right norms is a task beyond our powers. They could claim that it is beyond the powers of any being less omniscient than God. So, we had better just take his word for what's morally right. We ought neither question it nor try to figure out what the right-making fact about the right morality is, the one that God can see but we can't. This is a blatant dodge that all but the least inquisitive theists will reject. After all, morality isn't rocket science. Why are its grounds beyond the ken of mortal humans?

What does Plato's problem for sermonizing about morality have to do with scientism, or nihilism for that matter? Two things. First, Plato's argument shows that our moral norms need to be justified and that religion is not up to the job. Second, it turns out that scientism faces a coincidence problem just like the one troubling the theists. No theist can accept that it's just a coincidence that a moral norm is the right one and that God just happened to choose it for us. One of these two things must explain the other. Similarly, scientism is going to be faced with its own intolerable coincidence. But unlike theism, it's going to have a solution: nihilism.

Nihilism maintains that there isn't anything that makes our moral norms the right ones or anyone else's norms the right ones either. It avoids the challenge Plato set for anyone who wants to reveal morality's rightness. Nihilism instead recognizes that Plato's challenge can't be met. But the nihilist doesn't need to deny that almost all people share the same core moral norms, theists and nihilists included. Ironically, almost everyone's sharing a core morality is just what nihilism needs to show that no morality can be justified. What's more, the reasons that make nihilism scientifically and scientistically unavoidable also reveal that it doesn't have the disturbing features scientific people worry about. Public relations problem solved.

TWO EASY STEPS TO NIHILISM

We can establish the truth of nihilism by substantiating a couple of premises:

- First premise: All cultures, and almost everyone in them, endorse most of the same core moral principles as binding on everyone.

- Second premise: The core moral principles have significant consequences for humans' biological fitness—for our survival and reproduction.

Justifying the first premise is easier than it looks. It looks hard because few moral norms seem universal; most disagree somewhere or other. Some moral values seem incommensurable with each other—we don't even know how to compare them, let alone reconcile them. Moral norms are accepted in some localities but not in others: some cultures permit plural marriage, while others prohibit it; some require revenge taking and honor killing, while others forbid these acts. Adultery, divorce, abortion, and homosexual relations go in and out of moral fashion. At first glance, it looks like there is a lot of moral disagreement. It's enough to make some anthropologists into moral relativists (moral codes bind only within their cultures). It makes some philosophers into moral skeptics (there may be an absolute moral truth, but we can't be certain or don't know what it is).

On the other hand, in human culture there has long been a sustained effort to identify a core morality—one shared by the major religions, one that cuts across political differences and diverse legal codes. It is a core morality that has held constant or been refined over historical epochs. Soon after the founding of the United Nations in the late 1940s, there was enough consensus on core morality that it was enshrined in the Universal Declaration of Human Rights. Roughly the same core morality has been written into many more international conventions on rights and liberties that have been adopted by varying coalitions of nations ever since. Even countries that decline to be parties to these treaties have at least paid lip service to most of the norms they enshrine. These agreements don't prove that there is a core morality, but they're good evidence that most people think there is.

Thus, there is long-standing evidence for the existence of core morality in the overlap of ethical agreement by the major religions and in the lip service paid by international agreements. But wait, there's more. Neuroscience, in particular the techniques of fMRI (functioning magnetic resonance imaging of brain activity), increasingly shows that people's brains react the same way to ethical problems across cultures. This is just what the existence of core morality leads us to expect. If you are really worried about whether there is a core morality, you can jump to the next chapter, in which its existence is established beyond scientific (if not scientific) doubt, and then come back.

At the base of the diverse moral codes out there in the world, there are fundamental principles endorsed in all cultures at all times. The difficulty lies in actually identifying the norms that compose this core morality. What is the difficulty? These almost universally agreed-on norms are so obvious that they are easy to miss when we set about trying to identify them. Instead we think of interesting norms such as "Thou shalt not kill" and immediately realize that each of us buys into a slightly different and highly qualified version of the norm. For some of us, "Thou shalt not kill" excludes and excuses self-defense, military activity, perhaps capital punishment, euthanasia, other cases of mercy killing, and killing other great apes, primates, mammals, and so forth. When we consider how pacifists, opponents of capital punishment, proponents of euthanasia, and so many others disagree on some or all of these qualifications, it's tempting to conclude that there is no core morality we all share, or else it's too thin to have any impact on conduct.

A more accurate way to think about core morality begins by recognizing those norms that no one has ever bothered to formulate because they never come into dispute. They might even be difficult to formulate if they cover every contingency and exclude all exceptions. If we set out to express them, we might start out with candidates like these:

Don't cause gratuitous pain to a newborn baby, especially your own.
Protect your children.

If someone does something nice to you, then, other things being equal, you should return the favor if you can.

Other things being equal, people should be treated the same way.

On the whole, people's being better off is morally preferable to their being worse off.

Beyond a certain point, self-interest becomes selfishness.

If you earn something, you have a right to it.

It's permissible to restrict complete strangers' access to your personal possessions.

It's okay to punish people who intentionally do wrong.

It's wrong to punish the innocent.

Some of these norms are so obvious that we are inclined to say that they are indisputably true because of the meaning of the words in them: having a right to something is part of what we mean by the words "earn it." Other norms are vague and hard to apply: when exactly are "other things equal"? And some norms could easily conflict with each other when applied: would you really treat your children and other people the same? This shows how difficult it is to actually tease out the norms we live by, to list all their explicit exceptions, establish their priorities over other norms, and show how they are to be reconciled when they come into conflict. Almost certainly, the actual norms we live by but can't state will be somewhat vague, will have a list of exceptions we can't complete, and will conflict with other equally important norms. But most of the time, none of these problems arise to bedevil the application of the norms of core morality.

The next step in understanding moral disagreement involves recognizing that such disagreements almost always result from the combination of core morality with different factual beliefs. When you combine the uncontroversial norms of the moral core with some of the wild and crazy beliefs people have about nature, human nature, and especially the supernatural, you get the ethical disagreements that are so familiar to cultural anthropology. For example, Europeans may deem female genital cutting and/or infibulation to be mutilation and a violation of the core moral principle that forbids torturing infants for no reason at all. West and East African Muslims and Animists will reject the condemnation, even while embracing the same core morality. They hold that doing these things is essential to young girls' welfare. In their local environments, some genital cutting makes them attractive to potential future husbands; some sewing up protects them from rape. The disagreement here turns on a disagreement about factual beliefs, not core morality.

Even Nazis thought themselves to share core morality with others, including the millions they annihilated, as the historian Claudia Koonz documents in *The Nazi Conscience*. Outside of the psychopaths among them, Nazis were right to think that to a large extent they shared our core morality. It was their wildly false *factual* beliefs about Jews, Roma, gays, and Communist Commissars, combined with a moral core they shared with others, that led to the moral catastrophe of the Third Reich. You may be tempted to reply that it couldn't be that the Nazis just got their facts wrong, since there was no way to convince them they were mistaken. You can't reason with such people; you just have to confine them. True enough. But that shows how difficult it is to pry apart factual beliefs from moral norms, values, and the emotions that get harnessed to them. It's what makes for the appearance of incommensurability of values we so often come up against.

In any case, the argument that we're developing here doesn't require that every part of the core morality of every culture, no matter how different the cultures, be exactly the same. What we really need to acknowledge is that there is a substantial overlap between the moral cores of all human cultures. The principles in the overlapping core are among the most important ones for regulating

human conduct. They are the ones we'd "go to the mat for," the ones that are justified if any moral norms are justified. But as noted, they are also the ones that occasion the least argument, even across different cultures, and so hardly ever get explicitly articulated.

Along with everyone else, the most scientific among us accept these core principles as binding. Such norms reveal their force on us by making our widely agreed-on moral judgments somehow *feel* correct, right, true, and enforceable on everyone else. And when we are completely honest with ourselves and others, we really do *sincerely endorse* some moral rules we can't fully state as being right, correct, true, or binding on everyone. Scientism is not going to require that we give up that feeling or withdraw the sincerity of the endorsement. In a fight with these emotionally driven commitments, they'll win every time, for reasons that will become clear in Chapter 6. But scientism does recognize that emotionally driven commitment is no sign of the correctness, rightness, or truth of what we are emotionally committed to. As we'll see, add that recognition to science's explanation of where the shared norms come from, and you get nihilism.

The second premise we need to acknowledge is that core morality, the fundamental moral norms we agree on, has serious consequences for survival and reproduction. This will be especially true when those norms get harnessed together with local beliefs under local conditions. The connection between morality and sex is so strong, in fact, that each may be the first word you think of when you hear the other. The connection makes for a lot of good jokes. (Did you hear the one about Moses descending from Sinai with the two tablets bearing the ten commandments? "I have good news and I have bad news," he says. The children of Israel ask, "What's the good news?" Moses replies, "I argued him down from 38." They ask, "And the bad news?" Moses replies, "Adultery is still one of them.")

The idea that the moral core has huge consequences for survival and reproduction should not be controversial. Any long-standing norm (and the behavior it mandates) must have been heavily influenced by natural selection. This will be true whether the behavior or its guiding norm is genetically inherited, like caring for your offspring, or culturally inherited, like marriage rules. That means that the moral codes people endorse today almost certainly must have been selected for in a long course of blind variation and environmental filtration. Because they had an impact on survival and reproduction, our core moral norms must have been passed through a selective process that filtered out many competing variations over the course of a history that goes back beyond *Homo erectus* to our mammalian ancestors.

Natural selection can't have been neutral on the core moralities of evolving human lineages. Whether biological or cultural, natural selection was relentlessly moving through the design space of alternative ways of treating other people, animals, and the human environment. What was that process selecting for? As with selection for everything else, the environment was filtering out variations in core morality that did not enhance hominin reproductive success well enough to survive as parts of core morality. (You'll find much more on how Darwinian processes operate in cultural evolution in Chapter 11.)

There is good reason to think that there is a moral core that is almost universal to almost all humans. Among competing core moralities, it was the one that somehow came closest to maximizing the average fitness of our ancestors over a long enough period that it became almost universal. For all we know, the environment to which our core morality constitutes an adaptation is still with us. Let's hope so, at any rate, since core morality is almost surely locked in by now.

If you are in any doubt about this point, you are in good company. Or rather, you were until the last 50 years of research in behavioral biology, cognitive social psychology, evolutionary game theory, and paleoanthropology. Until the recent past, no one thought that core morality was selected for. Now

science knows better. If you are in doubt about this matter, the next chapter shows how natural selection made core morality inevitable. It's the key to taking the sting out of the unavoidable nihilism that this chapter is about to establish.

**IF CORE MORALITY IS AN ADAPTATION,
IT MUST BE GOOD FOR US.
SO WHY NIHILISM?**

Grant the two premises—moral norms make a difference to fitness and there is a universal core morality. Then the road to nihilism becomes a one-way street for science and for scientism.

The core morality almost everyone shares is the correct one, right? This core morality was selected for, right? Question: Is the correctness of core morality and its fitness a coincidence? Impossible. A million years or more of natural selection ends up giving us all roughly the same core morality, and it's just an accident that it gave us the right one, too? Can't be. That's too much of a coincidence.

Remember Plato's problem. God gave us core morality, and he gave us the right core morality. Coincidence? No. There are only two options for the theist. Either what makes core morality right is just the fact that God gave it to us, or God gave it to us because it is the right one.

Of course, there is no God, but science faces a very similar problem. Natural selection gave us morality, and it gave us the *right* morality, it seems. So, how did that happen? The question can't be avoided. We can't take seriously the idea that core morality is both correct and fitness maximizing and then claim that these two facts have nothing to do with each other. That's about as plausible as the idea that sex is fun and sex results in reproduction, but these two facts have nothing to do with each other.

Is natural selection so smart that it was able to filter out all the wrong, incorrect, false core moralities and end up with the only one that just happens to be true? Or is it the other way around: Natural selection filtered out all but one core morality, and winning the race is what made the last surviving core morality the right, correct, true one.

Which is it?

It can't be either one. The only way out of the puzzle is nihilism. Our core morality isn't true, right, correct, and neither is any other. Nature just seduced us into thinking it's right. It did that because that made core morality work better; our believing in its truth increases our individual genetic fitness.

Consider the second alternative first. Natural selection filtered out all the other variant core moralities, leaving just one core morality, ours. It won the race, and that's what made the last surviving core morality, our core morality, the right, correct, true one. This makes the rightness, correctness, truth of our core morality a result of its evolutionary fitness. But how could this possibly be the answer to the question of what makes our core morality right? There doesn't seem to be anything in itself morally right about having lots of kids, or grandchildren, or great grandchildren, or even doing things that make having kids more likely. But this is all the evolutionary fitness of anything comes to.

The first alternative is the explanation for the correlation that we'd *like* to accept: core morality is the right, binding, correct, true one and that is why humans have been selected for detecting that it is the right core morality. But natural selection couldn't have been discerning enough to pick out the core morality that was independently the right, true, or correct one. There are several reasons it had little chance of doing so.

First, there is lots of evidence that natural selection is not very good at picking out true beliefs,

especially scientific ones. Natural selection shaped our brain to seek stories with plots. The result was, as we have been arguing since Chapter 1, the greatest impediment to finding the truth about reality. The difficulty that even atheists have understanding and accepting the right answers to the persistent questions shows how pervasively natural selection has obstructed true beliefs about reality.

Mother Nature's methods of foisting false beliefs on us were so subtle they have only recently begun to be detected. By the sixteenth century, some people were beginning to think about reality the right way. But natural selection had long before structured the brain to make science difficult. As we'll see in later chapters, the brain was selected for taking all sorts of shortcuts in reasoning. That was adaptive in a hostile world, but it makes valid logical reasoning, and especially statistical inference, difficult. Without the ability to reason from evidence, getting things right is a matter of good luck at best.

Second, there is strong evidence that natural selection produces lots of false but useful beliefs. Just think about religion, any religion. Every one of them is chock full of false beliefs. We won't shake any of them. There are so many, they are so long-lasting, that false religious beliefs must have conferred lots of adaptive advantages on believers. For example, it's widely thought that religious beliefs are among the devices that enforce niceness within social groups. The hypothesis that organized religion has adaptive functions for people and groups is backed up by a fair amount of evolutionary human biology. It couldn't have done so except through the false beliefs it inculcates. Of course, all the religious beliefs that natural selection foisted on people made acquiring scientific truths about the world much more difficult.

There is a third reason to doubt that natural selection arranged for us to acquire the true morality. It is really good at producing and enforcing norms that you and I think are immoral. It often selects for norms that we believe to be morally wrong, incorrect, and false. In fact, a good part of the arguments against many of these "immoral" beliefs rests on the fact that natural selection can explain why people still mistakenly hold them.

There are lots of moral values and ethical norms that enlightened people reject but which Mother Nature has strongly selected for. Racism and xenophobia are optimally adapted to maximize the representation of your genes in the next generation, instead of some stranger's genes. Consider the almost universal patriarchal norms of female subordination. They are all the result of Darwinian processes. We understand why natural selection makes the males of almost any mammalian species bigger than the females: male competition for access to females selects for the biggest, strongest males and so makes males on average bigger than females. The greater the male-male competition, the greater the male-female size difference. We also know that in general, there will be selection for individuals who are bigger and stronger and therefore impose their will on those who are weaker—especially when it comes to maximizing the representation of their genes in the next generation. But just because the patriarchy is an inevitable outcome of natural selection is no reason to think it is right, correct, or true.

In fact, once we see that sexism is the result of natural selection's search for solutions to the universal design problem of leaving the most viable and fertile offspring, some of us are on the way to rejecting its norms. We can now explain away sexism as a natural prejudice that enlightened people can see right through. In different environments, natural selection has produced other arrangements; consider, for example, the social insects, where the top of the hierarchy is always a female.

Natural selection sometimes selects for false beliefs and sometimes even selects against the acquisition of true beliefs. It sometimes selects for norms we reject as morally wrong. Therefore, it can't be a process that's reliable for providing us with what we consider correct moral beliefs. The

fact that our moral core is the result of a long process of natural selection is no reason to think that our moral core is right, true, or correct.

Scientism looks like it faces a worse problem than the one Plato posed for theism. At least the theist can admit that there must be something that makes our moral core the right one; otherwise God would not have given it to us. The theist's problem is to figure out what that right-making property is. Faced with the widespread belief that the moral core is correct and faced with the fact that it was given to us by natural selection, scientism has to choose: did our moral code's correctness cause its selection, or did its selection make it the right moral core? As we've seen, scientism can't take either option.

Scientism cannot explain the fact that when it comes to the moral core, fitness and correctness seem to go together. But neither can it tolerate the unexplained coincidence. There is only one alternative. *We have to give up correctness.* We have to accept that core morality was selected for, but we have to give up the idea that core morality is true in any sense. Of course, obeying core morality is convenient for getting our genes copied in the next generation, useful for enhancing our fitness, a good thing to believe if all you care about is leaving a lot of offspring. If core morality is convenient, useful, good for any fitness-maximizing creature in an environment like ours to believe, then it doesn't matter whether it is really true, correct, or right. If the environment had been very different, another moral core would have been selected for, perhaps even the dog-eat-dog morality Herbert Spencer advocated under the mistaken label of social Darwinism. But it wouldn't have been made right, correct, or true by its fitness in that environment.

SCIENTISM STARTS WITH the idea that the physical facts fix all the facts, including the biological ones. These in turn have to fix the human facts—the facts about us, our psychology, and our morality. After all, we are biological creatures, the result of a biological process that Darwin discovered but that the physical facts ordained. As we have just seen, the biological facts can't guarantee that our core morality (or any other one, for that matter) is the right, true, or correct one. If the biological facts can't do it, then nothing can. No moral core is right, correct, true. That's nihilism. And we have to accept it.

Most people want to avoid nihilism if they can. And that includes a lot of people otherwise happy to accept scientism. Anti-nihilists, scientific and otherwise, may challenge the two premises of this chapter's argument for nihilism: the notion that there is a core morality and the idea that it made a difference for survival and reproduction. The irony is that together these two premises, whose truth implies nihilism, also take the sting out of it. The next chapter sketches enough of what science now knows about human evolution to underwrite both premises and so make them unavoidable for scientism. At the same time, it shows that nihilism is nothing to worry about.



THE GOOD NEWS:
NICE NIHILISM

TO NAIL DOWN NIHILISM, WE NEED TO BE VERY confident about how and why core morality is adaptive. We need to show why over the course of 3.5 billion years, relentless selection for fitness-maximizing creatures should have produced people with an almost universal commitment to a core morality. For evolutionary biology, this is the problem from heaven. Or at least it's the reason why more people prefer God's dominion to Darwin's hell of cutthroat competition.

For a long time after Darwin wrote *On the Origin of Species*, the existence of core morality made Darwinian natural selection apparently irrelevant to humanity's history, its prehistory, or its natural history for that matter. No one gave natural selection a chance to explain how our core moral code is possible, let alone actual. On the surface, core morality looks harmful to individual fitness. Think about all the cooperation, sharing, and self-sacrifice it enjoins. Consider its commitment to fairness, justice, equality, and other norms that obstruct looking out for number one and number one's offspring.

For over 100 years, Darwin's own difficulty explaining how core morality is even a possible result of natural selection was the single greatest obstacle to his influence in the social and behavioral sciences. Except for some racist kooks, no one took Darwin seriously in sociology, politics, economics, anthropology, psychology, or history.

All that has now changed. In the last 45 years or so, we have come to understand how natural selection gave us core morality. It turns out that doing so was integral and essential to the only adaptation that could have ensured our survival as a species. In fact, the adaptation that core morality constitutes enabled us to break out of that evolutionary bottleneck to which we were confined on the African savanna. We'll trace that process here in a little detail because of its attractive consequences for the nihilism to which we are committed. It is crucial to remember, however, that in this story, all the links in the chain are forged by the laws of nature that Darwin discovered. It is these that make the explanation convey real understanding.

WHEN IT COMES TO HUMANS, NATURE HAD TO SELECT FOR NICENESS

Some components of core morality are easy explanatory targets for natural selection. Pretty much any norm that encourages good treatment of children and discourages bad doesn't need a very complicated evolutionary explanation. It's the components that encourage or require being nice to people genetically unrelated to you that has resisted evolutionary explanation. The most characteristic feature of human affairs is at the same time the one most recalcitrant to Darwinian explanation. The fact is, genetically unrelated people are nice to one another in many different ways, often without any guarantee or prospect of reward or reciprocation.

From the point of view of fitness maximization, this doesn't seem to make any sense. Of course, in small family groups of very closely related members, being nice to others will be selected for

because it favors one's own (selfish) reproductive fitness; the people you are nice to will share your genes, and so their reproductive success will be yours, too. But once our hominin ancestors began to expand beyond the nuclear family, any tendency to niceness beyond the family should have been stamped out by natural selection, right?

This is not just a theoretical problem for human evolution. It was a concrete, life-or-death “design problem” faced by our hominin ancestors once they found themselves on the African savanna a few million years ago. They had been pushed out of the shrinking jungle rain forest by competition or because it was no longer capable of supporting them, or both. Our closet living relatives, the chimpanzees and gorillas, are still there, though hardly thriving themselves. It almost certainly was not intelligence or foresight that forced our ancestors onto the cooling, drying African plain. That's the lesson of Lucy, the 3-million-year-old fossil hominin found in Ethiopia in 1974; she was bipedal, but her brain was no larger than that of other primates. Survival on the savanna started to select for upright walking, both to see threats and keep cool, long before it started to select for increased intelligence.

More than a million years ago, our ancestors were beginning to differ from the other primates in at least three respects: they were living longer; they were having more offspring, owing to a much reduced interval between births compared to other primates; and the offspring were dependent for longer and longer periods than other primate babies. These changes, of course, had been gradual, probably caused by a change in diet that had other consequences. On the savanna, there was meat to scavenge from the carcasses that the megafaunal top predators were bringing down. It was mainly to be found encased by bone, in marrow and brain that the top predators couldn't get to. We could. Like chimps today, our ancestors were already using simple stone tools. Breaking bone was possible if they could only get to the carcasses quickly enough, then scare off the predator and keep it away. There wasn't much chance of that—but there must have been a little.

Over time, the increased access to protein would by itself have produced the three problems our ancestors faced. Increased protein enhances fertility, reduces birth intervals, and increases longevity. It also selects for reduction in the size of the digestive system of animals like us, since it is far easier to extract nutrients from marrow and brain than from vegetable matter. The same forces changed the shape of our molar teeth, the jaw, and jaw muscles needed to grind away at cellulose. Eventually, we ceased to require the ruminant capacities of bovines and other herbivores. No longer needing to build big guts for ruminant digestion released energy to make larger brains. The persistent increase in brain size led to longer childhood dependence. Recall the brain fontanel discussed in Chapter 4. To survive the narrow birth canal, the infant has to postpone most of its brain growth till after birth; whence the long period of childhood dependence.

This package of changes produced the design problem from hell: more mouths to feed over longer periods, but mothers prevented from providing for older offspring by the demands of younger ones; males living longer and so having still more offspring, putting further strains on available resources; and all those offspring needing literally years of protection and nourishment before they could fend for themselves. What a nightmare.

Some obvious solutions to this design problem are not attractive. Reducing maternal and paternal care will solve the resource problem by reducing infant survival. But it will expose lineages, groups, the species to extinction. Besides, it's not an available option, having been selected against among mammals for hundreds of millions of years. Increasing the birth interval may also have the same population-reducing risk. Besides, it's difficult to do. The only available method aside from abstinence (fat chance) is nursing the children for longer intervals. Lactation prevents menstruation,

but it also puts further pressure on nutrition. Adults dying off earlier and offspring becoming self-supporting earlier are not available options.

Now, add in the presence of those megafauna predators on the savanna. In addition to the design problems already identified, our hominin ancestors also find themselves well down toward the very bottom of the carnivorous food chain. They were prey, not predators.

What could have solved the design problem faced by *Homo erectus*? Somehow the population explosion has to be put to use to increase resources instead of consuming them. Groups of adults might be able to protect populations against predators or even scare predators away from animal corpses with meat left to scavenge. Females could take turns caring for several offspring of relatives, freeing them to provide for the family group. Besides scaring megafauna away from carcasses, males could even work together to bring game down themselves. Older adults could start learning more about the environment, and with the kids' extended childhood, the adults would have enough time to teach them the most efficient and productive technologies. Obviously, an ever-increasing brain size was needed for all these things. But something else was needed, too. The minimum size of a group was going to have to increase beyond the nuclear family for any of these strategies to work. Somehow our ancestors had to acquire the social tolerance of other hominins that our primate relatives conspicuously lack.

Once they could actually stay calm in larger groups, the next step was to form some semblance of organization. To do this, somehow our ancestors needed to break out of the straitjacket of selfishness that natural selection had hitherto imposed. People needed to start cooperating big time. They needed to learn to behave according to norms of core morality. They need to slip the leash of "look out for number one" fitness maximization. And somehow they must have done so or we wouldn't be here to defend scientism. How did it happen?

Darwin tried to deal with this puzzle but failed. He argued that besides operating on lineages of individuals and their traits, natural selection also operated on lineages of groups of people. When these groups competed with megafauna and with one another for territory or prey, the winning groups would have to be the ones whose members worked together, who cooperated with one another to fight off predators, kill prey, and eventually kill off other groups of people. In the long run of competition between groups, those groups composed of people nice to one another, even when unrelated, would have been selected for. Eventually, the only groups left would be those composed of people nice to each other. Thus, according to Darwin, core morality triumphs as fitness enhancing for groups, despite its immediate fitness costs for the individual members of groups.

The trouble with this theory is the arms race. Once a fit enough strategy is in place, Mother Nature will start searching through design space to find new strategies that take advantage of the original one. In this case, the groups composed of nice individuals are always vulnerable to variation in the behavior of new members. A mutation can arise at any time among the offspring of nice people that makes one or more of them start to exploit the niceness of everyone else, to feather their own nests without being nice in return. Through their exploitation of the nice people around them, they will secure disproportionate resources. This improves their survival odds and increases the number of their offspring. After enough generations of such within-group competition, the original group of nice people has become a group of not very nice people. Worse than nihilists, natural selection has produced Thatcherite Republicans. And what is worse, they would subsequently face the original design problem of living too long and having too many demanding kids to survive in small family groups.

If it wasn't group selection, how did blind variation and environmental filtration hit on niceness, on

norm following, as a solution to the design problem the ancestors of *Homo sapiens* faced? Mother Nature had one thing to work with that we did share with other primates: big brains and the “social intelligence” that goes with them. Like a few other species, our ancestors had already managed to acquire a fair ability to predict and so to exploit and to protect themselves from the tactics of other animals—predators, prey, other primates. Ethologists who study these species credit this ability to a “theory of mind.” We’ll adopt their label for it. It’s an ability hominins and apes share with whales, dolphins, and elephants—other big-brained creatures.

If our distant ancestors were like present-day chimpanzees, then they engaged in precious little cooperation. Even in the best of circumstances, chimps don’t share easily. Food transfer from parent to child is pretty much tolerated theft. Chimps almost never spontaneously help one another and never share information, something even dogs do with their masters. By comparison, contemporary human babies do all these things, and with strangers, well before they can do much more than crawl. They help, they share food, they even convey information before they have language and certainly without knowing anything about norms or niceness. That means that human babies have both an ability conveniently called a theory of mind and an inclination to cooperate; both are hardwired or quickly and easily learned with little experience at a very early age. If these traits are learned early, quickly, and reliably, then they don’t have to be innate. But all the cognitive skills we need to acquire them will have to be hardwired, along with some sort of bias toward learning them. How did humans acquire so much social intelligence and so amiable a disposition to use it?

Here is a theory based on comparative behavioral biology. There is another trait that we share with a few of the species that also have a fair grip on other minds—elephants, some whales, tamarin monkeys, and dogs: cooperative child care. This is a trait absent among chimpanzees and gorillas. The fact that this trait—sharing the care of the offspring of others—arose independently more than once suggests that it, or whatever triggers it, is a fairly common blind variation. Selection for sharing child care demands improvements in theory-of-mind abilities by species that already have that ability, such as the primates. Now we have a coevolutionary cycle in which shared child rearing selects for improvements in theory of mind, and improvement in theory of mind enhances the fitness of those practicing cooperative child care, and so on, until you get people almost born to cooperate.

These abilities are being selected for in nuclear families or among very closely related ones. But survival requires larger groups of less closely related people to cooperate. Why suppose that there was strong enough selection for cooperation to trump the short-sighted selfish fitness maximization that threatens to unravel the first stirrings of niceness? After all, the nuclear family is not a group large enough to enable us to move up the food chain.

BEAUTIFUL MINDS: CHARLES DARWIN MEETS JOHN VON NEUMANN

This is where evolutionary biology makes common cause with modern economics to identify the design problem’s solution that saved humanity and saves Darwinian theory.

Game theory was invented by Hungarian physicist John von Neumann in the late 1920s. Never heard of him? He designed the first electronic computer, figured out how to build the H-bomb, and invented the U.S. nuclear deterrent strategy during the cold war. How did he manage to remain so anonymous? Mainly because his life was not as dramatic as John Nash’s, another of the founders of the field. Nash’s path from mathematics through madness to the Nobel Prize was such a good story that it was made into *A Beautiful Mind*, a movie that won four Academy Awards. Von Neumann’s mind was at least as beautiful and his contributions to game theory were at least as important. He started the whole

enterprise, and without it we wouldn't have the solution to Darwin's puzzle about how core morality is even possible.

The label "game theory" does von Neumann and Nash's branch of economics a profound disservice; it's hard to take seriously a bit of science devoted to understanding games. Despite the name, game theory isn't really about games. It's about all strategic interactions among people, about how people behave when they have to figure out what other people are going to do so that they can scheme together with them or protect themselves from the schemes of others. Notice that to scheme with or against someone else requires that you have a theory of mind. You have to be able to put yourself in the other player's shoes, so to speak, and act as if he or she has a mind that makes choices in order to figure out what you should do. Game theory should have been called the theory of strategic interaction, but it's too late to change names.

Once invented, game theory had rather limited influence outside mathematical economics until it was taken up by evolutionary biologists. Their interest was sparked by the fact that game theory enables us to make explicit the design problem posed by the long-term benefit of cooperation and the immediate cost in individual fitness of doing so.

In economics, game theory assumes that the competitors in the game are rational agents. Of course, evolutionary game theory can't assume that individual animals are rational agents—it doesn't even assume humans are. Nor does evolutionary game theory assume that animals' behavior is genetically hardwired. It requires only that behavior respond to environmental filtration. If hardwired, then natural selection keeps selecting for more and more optimal behavior by filtering the genes that hardwire the behavior.

If the behavior is learned, the learning device is hardwired into the animal's brain. It filters for the fittest strategic behaviors by operant reinforcement, the process B. F. Skinner discovered and named. Operant reinforcement is just natural selection operating over an animal's lifetime. Behavior is randomly, blindly emitted; the environment rewards some of these behaviors and punishes others. In Skinner's lab, the rewards were usually food pellets. In nature, rewards include obvious outcomes like successful feeding foraging, fighting, and reproduction. Punishments are unsuccessful outcomes in these activities. Reinforcement results in the repetition, persistence, or strengthening of the initially randomly emitted behavior. Punishment has the reverse result. Skinner devoted most of his research life to showing how much behavior, including what looks to us to be intended, consciously planned behavior, is the result of this process of blind variation and environmental filtration.

In the case of hardwiring, nature usually has millions of years to fine-tune behavior to its environment; in the case of learning, it often has time only to hit on minimally satisfactory strategies.

Let's set up the problem Darwin tried to solve with group selection in the way that game theorists think about it. Often, they start with a particular model, the famous *prisoner's dilemma* game (PD for short). Suppose you and I set out to rob a bank and we are caught with our safecracking tools before we can break in. We are separated and informed of our rights as criminal suspects and then offered the following "deals." If neither of us confesses, we'll be charged with possession of safecracking tools and imprisoned for one year each. If we both confess to attempted bank robbery, a more serious crime, we'll each receive a five-year sentence. If, however, only one of us confesses and the other remains silent, the confessor will receive a suspended one-year sentence in return for his or her confession, and the other will receive a ten-year sentence for attempted bank robbery. The question each of us faces is whether or not to confess.

As a rational agent, I want to minimize my time in jail. If I think you're going to confess, then to minimize my prison sentence, I had better confess, too. Otherwise, I'll end up with ten years and

you'll just get a suspended one-year sentence. But come to think of it, if I confess and you don't, then I'll get the suspended one-year sentence. Now it begins to dawn on me that whatever you do, I had better confess. If you keep quiet and I confess, I'll get the shortest jail sentence possible. If you confess, then I'd be crazy not to confess as well. If I don't confess I might get the worst possible outcome, ten years. So, I conclude that no matter what you do, the only rational thing for me to do is to confess. Game theorists call this the dominant strategy.

Now, how about your reasoning process? Well, it's exactly the same as mine. If I confess, you'd be a fool to do otherwise, and if I don't, you'd still be a fool to do otherwise. You have a dominant strategy—the most rational under the circumstances—and it's the exact same strategy as mine.

The result is that we both confess and both get five years in the slammer. Where's the dilemma? Well, if we had cooperated, we would both have gotten one year. Looking out for our own interests leads us to a "suboptimal" outcome, one less desirable than another that is attainable. Rationality leads to suboptimality; there is an outcome that both rational egoists prefer but can't reach. Hence, the dilemma.

Why should this model of strategic interaction have any bearing on an evolutionary account of social cooperation? Because prisoner's dilemmas are all over the place in social and biological life, and they have been since before there were humans. They must have been especially frequent and important among scavenging hominins on the African savanna a million years ago. Any way you look at it, two hungry human scavengers coming on a recently killed carcass face several distinct prisoner's dilemmas. They know that a fresh kill means that there is a serious predator around, one they can't deal with. They know also that the fresh kill will attract other scavengers they can't deal with either—hyena packs. If they both scavenge, they'll get some food, but the chances of being surprised by predators or more powerful scavengers are high, and the outcome would be fatal. One could watch for predators while the other eats, but chances are the watcher gets nothing while the scavenger feasts. If they each scavenged for a few minutes and then stopped and listened for a few minutes, they'd get some food and avoid a fatal encounter with a predator. But the moment one of them stops to watch, the other will have a clear field to gorge himself at the look-out's expense. What to do? The dominant strategy is to scavenge like mad and don't bother looking out for threats. If the other guy splits his time between eating and watching, you're better off, and if the other guy doesn't, your watching for predators lets him take the best cuts. One thing they can't do is agree with one another to alternate eating and watching. Even if they are smart enough to figure this out and have enough language to bargain to an agreement, there is no way they can enforce the agreement on one another. That would take time from eating from the carcass and watching for predators. What is more important, neither will be prepared to enforce on the other the promise each makes to alternate eating and watching. That would just take time from eating or watching. Their mutual promises will be cheap talk, empty words. A classic PD.

Of course, watching and eating are not the only strategies available to each scavenger. Another option is to just kill the other guy the moment he turns his back and starts to scavenge. That will leave everything to the killer, at least till a stronger predator arrives. Of course, if they both left the other alone and began to eat, they'd each do better than dying. Neither can afford to risk that. All they can do is warily stalk each other, getting hungrier all the time, while the vultures take all the meat on the carcass.

The PD problem kicks in even before there is a chance to feed. Suppose two scavengers come upon a carcass being eaten by one lone hyena. If they both expose themselves, making noise and throwing rocks, there is a good chance they'll drive the hyena away and have a decent meal. But there is a

chance one or both might be injured by the surprised hyena. If one hides while the other tries to scare the hyena away, he'll do better no matter what happens—less risk, more eating time. Each hangs back in the undergrowth, waiting for the other to charge the hyena—that's the dominant strategy. As a result, neither of them gets anything. But they can't enforce on each other the mutually beneficial strategy of both of them scaring the hyena. Another PD.

In everyday life today, we still find ourselves in PD situations—for example, every time we go shopping. Consider the last time you purchased a soft drink at a convenience store just off the highway in a region of the country you'll never visit again. You have a dollar bill in your hand and want a drink; the salesperson behind the counter has the drink and wants the bill in your hand. He proffers the bottle, with the other hand held out for the money. *His* best strategy is to take your bill and hang on to the drink. If you complain, he'll simply deny you paid him. You won't call the police. You simply don't have time and it's not worth the trouble, in spite of your moral indignation; you're better off just going to the next convenience store on the road. *Your* best strategy is to grab the bottle, pocket your bill, and drive off. Will the salesperson call the police? If he did, would they give chase just for a bottle of Coke? The answer to each of these questions is no. It's not worth their trouble.

Knowing all this, neither of you does the rational thing. Thoughtlessly, irrationally, you both cooperate, exchanging the dollar bill for the drink. There are infinite examples like this. Consider the last time you left a tip, flushed a public toilet, or added some change to a street musician's take. People find themselves in PDs constantly and almost never choose the dominant egoistical strategy. The economists need an explanation for why we are nice when it doesn't look like the rational strategy. The biologist needs an explanation for why we are nice when it doesn't look like the fitness-maximizing strategy.

Here is the evolutionary game theorist's explanation, and it's just what Darwinian theory needs to show that niceness can enhance fitness and thus get selected for. Fortunately for humans, in addition to imposing a small number of single "one-shot" games on our ancestors, nature also imposed on them (and still imposes on us) other quite different strategic choice problems. These problems are made up of a large number of repetitions of the same simple game played over and over with different people. Nowadays, you and a lot of other people go to your local convenience store and purchase something from a number of different people behind the counter. Each transaction is one round in an iterated PD game you and everyone else—shopper and salesperson—is playing. Similarly, in the Pleistocene, scavenging repeatedly presented our ancestors with the same repeated strategic choice problems over and over again. That's what tribal life is all about.

In an iterated PD, there are many possible strategies players can adopt: always defect; always cooperate; cooperate on even-numbered games, defect on odd-numbered games; cooperate five times in a row, then defect five times in a row; cooperate until the other player defects, defect thereafter; defect until the other guy cooperates, then cooperate thereafter. The number of possible strategies is vast. Are there better strategies than always to defect? Is there a best strategy?

Under fairly common circumstances, there are several far better strategies in the iterated PD game than "always defect." In iterated PD, always taking advantage of the other guy is almost never the dominant strategy. In many of these iterated PDs, the best strategy is a slightly more complicated strategy called "tit for tat." Start out cooperating in round 1; then, in round 2, do whatever your opponent did in the previous round. Round 1 you cooperate. If the other player also cooperates, the next time you face that same player, cooperate again. If the other player cooperates in round 2, you know what to do in round 3. If in round 1 the other player defects, trying to free-ride on you, then in round 2 against that player, you have to defect as well. Following this rule after round 1, you will

have to continue cooperating until the other guy defects, and vice versa. If you have both been cooperating through 33 rounds and in his 34th game against you the other player suddenly switches to defect, trying to free-ride on your cooperation, then in round 35 against him, you have to defect. If in round 35, your opponent switches back to cooperate, then in round 36, you should go back to cooperation. That's why the strategy is called tit for tat.

Why is tit for tat almost always a better strategy in iterated PD than always defect? In a competition in which tit-for-tat players face each other, they cooperate all the time and so accumulate the second best payoff every time. That's usually a lot better outcome than playing always defect: you get the successful free-rider's top payoff the first time and only the third best payoff all the rest of the time. The more rounds in the iterated PD, and the more chances you have to play with tit-for-tat players and get the second highest payoff every time, the better the outcome for cooperating and the worse for defecting. Provided that the number of individual rounds in the iterated game is large enough, that the chances of playing against the same person more than once are high enough, and that the long-term payoff to cooperating in future games against other cooperators is high enough, it's always best to play tit for tat, especially if you are relentlessly looking out only for number one.

It's easy to set up one of these iterated PD games either among people or in a computer simulation. Give each of two test subjects two cards, one with a D for defect, the other with a C for cooperate. Each subject plays one card at the same time. If both play D, each gets \$2; if both play C, both get \$5; if one plays D and the other plays C, the D player gets \$10 and the C player gets nothing. Now, let a lot of subjects play one another many times. It's also easy to program a computer with a large number of strategies for playing the D card or the C card. Instead of money payoffs, you assign point payoffs to the various outcomes. Play the game or run the simulation 100 or 1,000 times. Add up the payoffs. Among human subjects and computer simulations, tit for tat usually comes out on top.

What's so good about tit for tat? Three things. First, this strategy always starts out being nice—that is, cooperating. Second, a player won't be taken for a sucker; a player retaliates when taken advantage of. Third, a player doesn't have to play against a tit-for-tat player many times to figure out what strategy he or she is using. And it's easy to switch to tit for tat and to start making the same gains.

The message that emerges from human experiments and computer simulations is also clear and also nice. Under a wide variety of conditions in which people face the strategic problem posed by an iterated PD, there will be selection for certain kinds of strategies. The fittest strategies are all variants on tit for tat: they are initially nice—players using them start out cooperating and are quick to forgive and minimally retaliatory; if opponents switch from defecting to cooperating, the tit-for-tat players switch to cooperation; and the strategies are transparent to other players, so it's easy to figure out how to do best when playing other tit-for-tat players—just do what they do.

If many or even just the most crucial interactions in our evolutionary past were such iterated games, then there would have been strong selection for nice strategies. If blind variation could impose tit for tat or some other nice strategy on the members of one or more groups of hominins, then those groups would flourish. Their members would solve the huge design problem imposed on them, and groups of such cooperators would be immune to the sort of subversion from within that undermined Darwin's group selection idea for how niceness could be selected for.

It's a nice story of how niceness could have emerged, but is it any more than a story? Is there evidence that our hominin ancestors actually faced an iterated prisoner's dilemma often enough over the time scale needed for cooperative strategies like tit for tat to be selected for? There is no direct evidence. There was no currency (natural or otherwise) to measure the payoffs, no way to measure

the frequency and numbers of interactions, and even if there were, strategic interaction problems don't fossilize well. On the other hand, when the choice was between extinction or cooperation, the fact that we are here strongly suggests that cooperative strategies must somehow have been selected for.

Moreover, repeated PD games are not the only ones people face and are not the only ones that select for niceness. Experiments focusing on other games that people (and computers) play and that we know our ancestors faced strongly suggest that fitness and niceness really do go together.

Consider this strategic interaction problem, called "cut the cake": Two players who don't know each other and can't communicate are each asked to bid for some portion of an amount of money, say, \$10. Each is told that if the other player's bid and theirs total more than \$10, neither gets anything, and if the total of their two bids is equal to or less than \$10, each receives what they bid. In this one-shot game, most people spontaneously bid an amount somewhere close to \$5. In this case, rationality does not by itself point to any particular strategy. So why should one be predominant? What is more, the predominance of fair bids is cross-culturally constant. Across a broad range of Western and non-Western, agricultural, pastoral, slash-and-burn, nomadic, and even hunter-gatherer societies, the fair offer is the predominant one.

Consider another game, this one a little more complicated. In this game, one player gets the first move; call him the proposer. He proposes how much of the \$10 the other player will receive and how much the proposer will keep. If the player going second, the disposer, agrees, then each party gets what the proposer decided. If the disposer rejects the proposal, neither party gets anything. Since this is a "take it or leave it" game, it's called the ultimatum game. In this game, it would obviously be irrational ever to decline even a very unequal split, since getting something is better than getting nothing. But across the same broad range of cultures, more often than not participants in the ultimatum game offer the other player \$4 or more, and often \$5. They also reject anything much less than a fair split. And people do this even when a tiny fraction of the total to be divided is a significant amount of money. In a lot of the cultures where this experiment has been run, \$1 will keep you alive for a day, yet many people will reject proposals that leave them with less than \$4 or so.

When asked to explain their strategies, most people identify a norm of equality as dictating their choice in the first game and a commitment to fairness as doing so in the second. In "cut the cake," people rarely ask for more than 5 units even when they think they can get away with 6. When asked why, they say they do so out of a sense of fairness. They also say that they get angry when the other player makes choices that strike them as unequal. When they play the disposer role in the ultimatum game, people describe their satisfaction at rejecting offers they think are too low. The feeling of satisfaction is evidently worth more to them than accepting a low offer of some amount of money rather than nothing. These games tap into strong feelings that lead people to do things that reflect important parts of core morality—norms of reciprocity, fairness, and equality.

In the experiments, the "cut the cake" and ultimatum games are one-shot games, played just once by people who don't know each other. Consider what happens when you program computers to play these two games over and over using many different strategies over a wide range of payoffs in individual games: demand 9, accept anything; demand 5, accept 4 or more; demand 4, accept 3 or more; and so on. Program a little evolution by natural selection into this simulation: If you start out with a lot of different competing strategies, have the simulation filter out the least rewarded strategies every ten or a hundred or a thousand or a million rounds. Replace each player using an eliminated strategy with a new one using one of the surviving strategies. It doesn't even have to be the best one. Program the computer to make a random choice. What you find is that if you play enough rounds, the

strategies that do best overall are very often the ones that are “fair” or “equal.” In “cut the cake,” asking for half is a winning strategy most of the time. In the ultimatum game, asking for half and refusing anything much less does very well.

What does best in the long run doesn't always do best in the short run, but human evolution was (and is) a long-run process. These computer simulation results, along with similar ones for repeated prisoner's dilemma games, strongly suggest one thing: if we evolved in circumstances that had these kinds of payoffs, there would have been strong selection for core morality. There would have been strong selection for anything that made people adopt norms of fairness, equity, and cooperation.

But how does natural selection get people to adopt such norms? How does it shape such adaptations? What is the quick and dirty solution to the design problems that arise in situations of iterated strategic choice? This problem looks like it's too hard to be solved by genetically based natural selection. Maybe if there were genes for playing tit for tat, they would be selected for. But at least in the human case, if not in animal models, such genes seem unlikely. The solution has to involve some cultural natural selection. It can't, however, help itself to very much culture. After all, we are trying to explain core morality. And to get much culture, you must already have a fair amount of core morality.

IRRATIONAL LOVE, UNCONTROLLABLE JEALOUSY, POINTLESS VENGEFULNESS, AND OTHER HEALTHY EMOTIONS

Mother Nature is ever on the lookout for cheap solutions to design problems. In the rain forest, the design problem was figuring out how to do best, or at least how to avoid doing worst, in a lot of one-shot strategic games. Here selfishness maximized both individual gains and fitness. To this day, chimpanzees won't cooperate to secure food that they then have to share. Real cooperation in the jungle never had a chance. Nice guys finished last. But then hominins found themselves on the savanna, with too many young mouths to feed for too long a time. There they found nothing much to eat except berries, nuts, and what they could scavenge from the megafauna's latest meal, provided they could scare it away and keep other scavengers away.

It's pretty obvious that on the savanna there were new strategic interaction problems in which selfishness and fitness maximization came apart. Under these changed circumstances, nature will start searching through the inevitable variations in behavior for quick and dirty solutions to the new design problem—the iterated strategic game. It will sift through strategies like “be nice only to offspring and mate,” “be nice only to close relatives,” “be nice to everyone in the tribe all the time,” or “be nice to people who were nice to you before.” It will filter many other strategies, too.

Until this point in human evolution, the be-nice-to-nonkin strategies were all being remorselessly selected against. But now, on the savanna, the payoffs in strategic interaction began to take on the values of iterated prisoner's dilemma or repeated “cut the cake” or the ultimatum game played again and again. Under these new circumstances, Mother Nature could exploit the ability to read minds that our ancestors shared with the other primates. She could even exploit social intelligence (the theory of mind) and the tendency to share that cooperative child care made mutually adaptive. Hominins already disposed by selection to make cooperative opening moves in a PD encounter would already be acting on a rule that is “nice.” Now it only needed to be shaped by operant conditioning—the natural selection in learned behavior Skinner discovered—into one that is retaliatory and clear. A variation in our traits that did those two things would itself strongly select for niceness in others and strongly select against being mean as well.

What kind of a device could nature have hit on in the course of our evolution that could guarantee to others that we will act in accordance with norms of niceness, fairness, equity, and much of the rest of the moral core?

It would have had to be a device that overrides the temptation to cheat, cut corners, free-ride when the opportunity occurs—and temptation can't be resisted, as we all know only too well. Psychologists long ago identified and explained a tendency that humans share with other creatures to prefer the smaller immediate payoffs that free-riding or cheating offers over the larger long-term payoffs that cooperating provides. We tend to have trouble postponing gratifications. From a Darwinian point of view, that used to be a good thing. Nature lacks foresight and can't afford it. It can't select traits whose payoff is very far down the road. By the time it gets down the road, we will probably be dead already. So, for a long time nature selected for immediate gratification over long-term gain. By the time hominins arrived on the scene, this preference for short-term over long-term gain had been firmly hardwired in us. But now all of a sudden, on the African savanna, it became maladaptive. Once payoffs to cooperation became high enough, Mother Nature needed to find a device to get us to resist temptation. Otherwise our posterity wouldn't have been around to secure the benefits of our self-control.

Go back to the reports of people who play the ultimatum game and “cut the cake,” and you'll see that Mother Nature already had something in her tool kit that fits this bill very nicely: *emotion*.

Long ago, nature provided mammals with emotions. Darwin was among the first naturalists to study them carefully. One of his last books was titled *The Expression of the Emotions in Man and Animals*. He recorded there the remarkable similarity between human, primate, and canine expressions of emotion. Emotions must go back very far into our evolutionary tree. In humans, natural selection has co-opted emotions to enforce norms. These norms that emotions enforce are often good for our genes but not for us. We honor them in spite of the disservice they do to our interests. Love and jealousy are good examples. Both are strong emotions and are often very high in short-term costs for guys and highly beneficial to the maximization of their (long-term, multigenerational) genetic fitness.

Let's start with love and the design problem it solves for males. A male won't get sexual access to a female unless the male can convince her that he'll be around to share some of his resources with her and the kids he is going to produce. Since females have been selected for not being fooled by mere expressions of fidelity, they demand stronger assurances before they will allow males to have their way with them. As the Hollywood producer Samuel Goldwyn noted, a verbal contract is not worth the paper it is written on. A male's promise is unenforceable. Females can't rely on it because for a male it would be irrational to keep. With millions of sperm, the male's best strategy is to promise, get sexual access, and renege. The mammalian female has only a few hundred eggs and a limited number of ovulatory cycles. She can't afford to guess wrong about a reliable mate. What will reliably guarantee unenforceable promises about the future when it would be irrational for any male to keep them? One thing that would do it is a sign of irrational commitment to the female and to her interests that could not be faked.

Why must the sign signal irrational commitment? Because females recognize that it's irrational of males to commit resources to one female. So the sign the male sends the female really has to be one of irrational commitment. Why must the sign be unfakable? Because a fakable sign of commitment is just that, fakable, and therefore not credible. Love is irrational and unfakable, by males at any rate. In nature's search through design space for a strategy that will secure males' sexual access, the emotion of love looks like it will just do the trick.

Irrational love does not fully solve the male's design problems. After pairing up, the male faces another issue: the uncertainty of paternity. To convey resources to his mate's offspring, he needs assurance that the kids are really his. This is an uncertainty problem females don't have (unless kids get switched after birth). The male needs to reduce the uncertainty as much as possible. One way to do this is to pose a credible threat to anyone suspected of taking advantage of any absence from his partner's bed. To make this threat credible, the male must be motivated to carry it out even when it is crazy to do so. And often it is crazy, since it's the strong, the powerful, and the rich who usually try to take advantage of the weaker. The emotion of uncontrollable jealousy fits the bill perfectly. Revenge must be a credible threat; males must convince everyone that they will take measures to punish cheating wives and/or their lovers no matter how great the cost to themselves. Overpowering jealousy does the job, though it makes the occasional male actually sacrifice his own short-term and long-term interests. In the overall scheme, the fact that every male is prone to feel such emotions maintains a norm among men and women that effectively reduces the uncertainty of paternity and so enhances most males' fitness. (Of course, female jealousy isn't selected for reducing the uncertainty of maternity. There is little to reduce. But the emotion's unfakable and irrational force deters other females from shifting her partner's resources to their offspring.)

Emotions are hardwired by genes we share and presumably share with other primates and indeed other mammals, as Darwin himself noticed. In us, of course, they get harnessed together with our highly developed theory-of-mind ability and with norms adaptive in our environments. They motivate enforcement of the norms they get paired up with, on others and on ourselves. Some of these norms solve design problems common to humans in all the environments we inhabit. These are parts of the moral core we all share. Others will not be part of core morality but will be locally restricted to the different ecologies that different groups inhabit. Some examples will illustrate how this works.

Anthropologists have uncovered how different ecologies select for differences in local moral norms. For example, there will be quite different norms in pastoral versus agricultural cultures. Shepherds and herders need to solve the problem of asserting ownership over flocks that may spread far beyond their immediate supervision. Among pastoral people, the norm of going to any lengths to punish rustling will be highly adaptive. So there will be strong selection for emotions of anger, revenge, disdain that accompany the norms prohibiting theft of stock animals. Where norms against rustling are not linked with these emotions, the motivation to act on them is reduced. And that lowers the fitness of the pastoralists who endorse the norms but don't have the right emotions. Again, the process is the same as the one that linked sex and pleasure. Nature was selecting among variations for fitness. Of course, in the emotion/norm combination, the linkage is culturally established and transmitted.

By contrast with pastoralists, sedentary farmers can keep an eye on their plows, even when others borrow them without asking. Among cultivators, there will be selection for norms that allow a certain amount of informal borrowing without permission. The emotions that go along with their norms will be different or weaker than among sheep and cattle herders.

The differences in which behaviors are fitness enhancing dictate which norms persist. Variation and selection harness preexisting emotions to enforce them. Not much genetic hardwiring is required for the Darwinian mechanism that finds and enforces moral norms, sometimes globally and sometimes locally. The emotions are bodily feelings and they are hardwired, subject to the usual Darwinian range of variation. But which norms they get harnessed to depends on what is fitness enhancing in an environment. Different environments make different packages of norm and emotion fitness enhancing. Some packages are fitness enhancing in all the environments we have occupied. These are the ones

that form our moral core. They are the ones packaged with the same emotions universally, or nearly so. All societies and cultures have words for anger, disdain, shame, and guilt, and in each, they name roughly the same set of bodily feelings. In each culture, they are linked to common norms—the ones in core morality. They are also linked to different norms—ones that reflect local culture and local environmental filtration.

If you think about how, for example, shame and guilt work, you will see that they are emotions practically designed to solve nature's problem of getting us to do the right thing, to be nice, at least for the sake of our genes' futures.

In brief, shame is the unpleasant feeling we get when others catch us not abiding by the norms almost everyone (including us) embraces. Guilt is the unpleasant feeling that overcomes us when we aren't caught. They are both strong enough by themselves to solve the problem nature created for itself when it made us prone to immediate gratification and resistant to long-term benefits.

Long ago in human cultural evolution, parents began drumming tit-for-tat norms into kids already prone to niceness. Shame at being caught free-riding on others' niceness and shame at being taken for a sucker by others' free-riding is easy to instill: "Aren't you ashamed of yourself?" The hard-to-fake tears, red face, and other bodily signs make shame a credible signal that we have learned our lesson and others can start to be nice to us again. Guilt may be even more important, since its unpleasantness, even in anticipation, is often enough to get us to resist the immediate payoff that comes from defecting once in an iterated PD. Those lucky enough to feel guilt when tempted to cheat are fitter than those who don't. Why? Most hominins underestimate the chances of getting caught by others. It's hardwired into us by Mother Nature because that was the fittest strategy before other people as smart as us came on the scene. Once we appeared, the chances of getting caught free-riding increased, but the temptation remained strong. It needed to be counterbalanced by something that would make honesty the best policy. Feelings of guilt fit the bill.

NICE NIHILISM IS ALL WE NEED

Some emotions enforce parts of core morality, and some enforce norms and actions we reject as not being part of core morality or even compatible with it. What do the parts of core morality that emotions enforce have in common with nonmoral or even immoral norms, like honor killing, that emotions also enforce? They are fitness enhancing in their respective local environments.

If we were selected for niceness, how come there are so many SOBs in the world, and still worse, serial killers, moral monsters, and Adolf Hitlers? Biology has the answer. Remember, perhaps the most profound of Darwin's observations was that there is always some variation in most heritable traits in every generation. A distribution of variations—often a normal, bell-curve distribution—is the rule and not the exception. Traits like niceness and packages combining nice norms and emotions that motivate them are the result of the interaction of the long childhood and a large number of different genes that program development. Slight differences in some of them, often together with slight differences in environments, will produce a good deal of variation in every generation. In every generation, there are going to be a few people who are too nice and get walked on and a few people who are not nice at all—saints and sociopaths.

When we add in variations in social skills, intelligence, and other traits, it's inevitable that a certain percentage of any population will turn out to be grifters, thieves, and other victimizers of us nice people. But we'll also have some saints, martyrs, and Samaritans. And of course, in every generation of a population that is large enough, there will be extreme cases—serial murderers like Jeffrey

Dahmer or Peter Sutcliffe and worse, charismatic monsters like Hitler and bureaucratic ones like Stalin.

What is the maximum proportion of egoistic free-riders that a society of nice guys can tolerate? Is it 10 percent or one-tenth of 1 percent? It depends on the number of times we end up playing a PD or some other strategic interaction with a free-rider, and it depends on the stakes. When the number of free-riders gets to the point that to protect ourselves from them we stop cooperating with everyone, something has to give. Either society unravels, or we act together to reduce the number of free-riders. Until we reach that point, we nice people will have to tolerate a few sociopaths and psychopaths as inconveniences in our daily lives. What we nice people won't tolerate is Hitlers, Stalins, Mao Zedongs, Pol Pots, and Osama bin Ladens, at least not forever. The trick is to detect them before it's too late.

But wait. Where do we scientific types get off condemning purdah and suttee, female genital mutilation and honor killing, the Hindu caste system, the Samurai Code of Bushido, the stoning of women who have committed adultery, or the cutting off of thieves' right hands? Isn't the high dudgeon we want to effect in the face of this sort of barbarism flatly inconsistent with nihilism—any kind, nice or not?

Not only do we condemn the blinkered morality of intolerant religions and narrow-minded cultures, nowadays, we condemn some moral norms that we ourselves used to embrace. We think of this change as progress, improvement, and enlightenment in our moral conscience. A good example is the palpable change in the attitude toward homosexuality. The Monty Python crew wouldn't be caught dead making fun of gays in the way they did in the early 1970s, even knowing as they did that one of their number was a homosexual.

Once it's saddled with nihilism, can scientism make room for the moral progress that most of us want the world to make? No problem.

Recall the point made early in this chapter that even most Nazis may have really shared a common moral code with us. The qualification "most" reflects the fact that a lot of them, especially at the top of the SS, were just psychopaths and sociopaths with no core morality. Where most Nazis "went wrong" was in the idiotic beliefs about race and a lot of other things they combined with core morality, resulting in a catastrophe for their victims and for Germany. The same goes for Stalin and his hatchet men, although most of them were motivated simply by fear in addition to having false beliefs. We can equally well find the false factual beliefs behind purdah, suttee, honor killing, and our own now repudiated sexism, racism, and homophobia. Think about how American Christianity managed to reconcile core morality with African-American slavery for over 250 years. All it had to do was combine core morality with lots of false beliefs about African people. Scientism allows for moral "improvement." It's a matter of combining the core morality that evolution has inflicted on us with true beliefs vouched safe for us by science. It's the failure to be scientific and scientific that leads from core morality to moral mistakes and even moral catastrophe. (More on this in Chapter 12.)

About the only thing that there is to worry about with nihilism is the name. Most people are nice most of time, and that includes nihilists. There is no reason for anyone to worry about our stealing the silver or mistreating children in our care. As for moral monsters like Hitler, protecting ourselves against them is made inevitable by the very same evolutionary forces that make niceness unavoidable for most of us. There is nothing morally right about being nice, but we are stuck with it for the foreseeable future.

Scientism has to be nihilistic, but it turns out to be a nice nihilism after all.

IN WALT DISNEY'S VERSION of *Pinocchio*, the Blue Fairy advises the boy-puppet that when it comes to right and wrong, "always let your conscience be your guide." Then Jiminy Cricket volunteers to be his conscience. If Pinocchio had been human instead of a puppet, he would not have needed the advice or for that matter a grasshopper to help him act on it. Humans have been selected for heeding their consciences—their moral censors. About the only ones who can't do so are the sociopaths and psychopaths who lack them altogether. Scientism shows us that letting our consciences be our guides enhances our fitness, but that doesn't make us morally right, or morally wrong for that matter. Instead it shows that there is no such thing as either morally right or wrong.

In the next chapter, we'll see that when it comes to the conscious, as opposed to the conscience, the situation is quite different. A great deal of what consciousness tells us is just plain false, and provably false. The mistakes that consciousness leads us into are so egregious that the only conclusion scientism can draw is to never let your conscious be your guide. And once we take this moral on board, we'll be ready to really see how radical, how bracing, how breathtakingly revolutionary a worldview scientism really is.