

Our Moral Fate

Allen Buchanan

© 2020 The Massachusetts Institute of Technology

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher.

Library of Congress Cataloging-in-Publication Data

Names: Buchanan, Allen E., 1948– author.

Title: Our moral fate : evolution and the escape from tribalism / Allen Buchanan, The MIT Press.

Description: Cambridge, Massachusetts : The MIT Press, 2020. | Includes bibliographical references and index.

Identifiers: LCCN 2019025836 | ISBN 9780262043748 (hardcover) | ISBN 9780262357876 (ebook)

Subjects: LCSH: Ethics—History. | Progress—Moral and ethical aspects.

Classification: LCC BJ71 .B83 2020 | DDC 170—dc23

LC record available at <https://lcn.loc.gov/2019025836>

Failed Attempts to Solve the Big Puzzle: Cooperation Is Not Enough

When I began this detective story, I lamented that in spite of all their other contributions, scientists haven't paid enough attention to using evolutionary principles to explain how moralities have changed over time—and in particular that they haven't provided or even tried to provide evolutionary explanations, whether biological or cultural, of the Two Great Expansions. And so they haven't tried to determine how evolutionary thinking can illuminate the possibilities for moral progress. Now I want to examine attempts to solve the Big Puzzle by people from another discipline, philosophy.

Philosophical Thinking about the Big Puzzle

A few contemporary philosophers have been impressed with how the circle of moral regard has expanded, *and* tried to explain it with an eye toward what evolutionary scientists say about the origins of human moralities. In this chapter, I focus on those thinkers, in particular Philip Kitcher, Peter Railton, and Peter Singer. I think they've all made valuable contributions to understanding human moralities. Yet I think that none of them has succeeded in explaining how the Two Great Expansions can be squared with the standard view about the evolutionary origins of moralities. In fact, they don't explain the Two Great Expansions, even if we set aside the question of how they could have come about, given the standard evolutionary origin story. My purpose here isn't criticism for its own sake. Understanding where these attempts to explain the Two Great Expansions go wrong will put us in a position to get the explanation right.

Solutions That Focus on How Morality Facilitates Cooperation

I'll start with Kitcher's and Railton's views, because they have a lot in common. I will argue that neither of these formidable thinkers can explain either of the Two Great Expansions because they make the same mistake: they try to explain how the circle of moral regard expands by showing how human cooperation changes. Kitcher (2011, 131–138) seems to commit the error I flagged in the introduction: thinking that we can explain expanding the circle of moral regard as resulting from expanding our cooperation to include more people. Railton (1986, 200) has a squeaky-wheel explanation for how some important forms of moral regard get extended to people who previously were disregarded. He thinks that under certain conditions, when a society neglects the important interests of a significant portion of its population, those people may mobilize and disrupt the cooperative scheme, leading to a new form of cooperation in which they are treated more equally.

What these two accounts have in common is that they are still shackled to the Cooperation Dogma, the assumption that we can explain human moralities and whatever changes have occurred in them by focusing exclusively on how moralities facilitate cooperation. As we'll see, sticking with that assumption renders both accounts incapable of explaining either of the Two Great Expansions.

Kitcher's Understanding of What Morality Is (and All It Is)

Kitcher's rich and ambitious book, *The Ethical Project*, makes it clear that he believes that expanding the circle of moral regard is a good thing, in fact perhaps the most important kind of moral progress. And he wants to explain it in evolutionary terms. He starts out as evolutionary scientists do, expounding on how our remote ancestors became moral beings because changing in that way enabled them to cooperate successfully, where success means reproductive success. To his credit, he doesn't think of human morality as a fait accompli, something that was completed at some time back in the middle to late Pleistocene. He thinks it is a work in progress, a *project* that humans are still engaged in (hence the title of his book).

Kitcher thinks that moralities first appeared as adaptations that helped early human societies solve or avoid “altruism failures” — selfish behaviors that disrupted cooperation. Even if natural selection produced humans who were *somewhat* altruistic (at least toward their kin and perhaps other members of their group as well), altruism was limited or, what amounts to the same thing, humans didn’t evolve to automatically always succeed in inhibiting the pursuit of their own interests when doing so was necessary to benefit others. So Kitcher thinks there were “altruism failures,” resulting in conflicts, free riding, and other behavior that impaired cooperation and thus reduced reproductive fitness. Groups that developed moralities—rules, social practices, and moral responses that coped well with “altruism failures”—survived and reproduced; those that didn’t went under. So eventually all human groups had moralities.

Kitcher doesn’t explore the possibility that morality nowadays does more than cope with “altruism failures.” In other words, Kitcher is a clear example of thinking of morality in purely functional terms, of assuming that morality is *constituted* by the function it originally evolved to perform, namely, facilitating cooperation and thereby contributing to reproductive fitness. Thinking that way ignores the possibility that morality nowadays, for some people at least, has features that are not functional in any sense that connects with evolutionary thinking. That’s a mistake, as my example of the human shoulder joint makes clear. If you think that morality (nowadays) is nothing more than a device for facilitating cooperation, then you have only one option for trying to explain moral progress in the dimension of inclusion: you’ll have to say that the circle of moral regard expanded because the circle of cooperation expanded.

That strategy doesn’t work. It can’t explain why people would extend moral regard to humans they will never cooperate with (either because those individuals are so disabled that they can’t cooperate or because one simply won’t ever *need* or even *want* to cooperate with them, either to maximize our reproductive fitness or for any other purpose). So it can’t even begin to explain the First Great Expansion.

Nor can it explain the fact that nowadays many people think that even in the case of humans we do or may cooperate with, their

capacity to cooperate is not the *only* basis of our moral regard for them, that instead they have a high moral status just because they are human beings. (Remember, the idea of human rights is that of rights you have by virtue of your humanity, not your ability to participate in cooperation.) Even more obviously, no cooperation-facilitator view of morality can explain why we should give a damn, morally speaking, about nonhuman animals that we'll never cooperate with.

Another problem arises for any view that characterizes morality as being constituted by the function of facilitating cooperation (in Kitcher's case, by overcoming altruism failures): it can't explain why for some people morality includes *intrapersonal* ideals of excellence, commitments to living in certain ways that have nothing to do with cooperation or with any interpersonal relations at all. I have in mind something I mentioned earlier: some central features of the account of virtues in Stoic philosophy and the core of Buddhist thought, and many other examples as well, including Christian and Hindu ascetic ideals. These moralities provide prescriptions for how a person's soul or psyche should be ordered that are independent of placing a value on any social effects such an ordering might have. In fact, some of them have advocated separation from society, dropping out from all cooperative schemes. An evolutionary understanding of morality that cleaves to the dogma that morality is all about cooperation can't explain the existence of these sorts of moralities.

Expanding the Circle of Cooperation Doesn't Produce the Two Great Expansions

The idea that the circle of regard has expanded because the circle of cooperation has expanded can't even begin to explain the Second Great Expansion. The vast majority of nonhuman animals simply aren't capable of being included in any form of cooperation that humans engage in. Yet many of us think that—and act as if—those animals have moral standing.

Most of the animals that people now think should be treated less cruelly than they traditionally have been just don't have what it takes, cognitively or motivationally, to cooperate with us in any meaningful sense. They can neither benefit us by cooperating with

us nor disrupt the cooperation we engage in. When we recognize animals as having any moral standing at all, we are operating with an understanding of the basis of moral standing that is nonstrategic, not cooperation based.

Even in the cases where it makes sense to say that humans and animals are cooperating (as in the interactions between sheepdogs and shepherds or bomb-sniffing dogs and their handlers), the power asymmetry is so great that the nature of the cooperation doesn't require the human side of the partnership to treat the animal side as equals or even to treat them in ways that avoid gratuitous suffering on their part. Sadly, anthropologists have documented that even people who depend on dogs for their survival, including Inuits who rely heavily on their sled dogs, routinely treat them in extremely cruel ways, inflicting pain on them, not because the pain is necessary to motivate the animals to perform the tasks assigned to them, but apparently just because they find it fun (Edgerton 1992, 20). And even if people treat animals on whom they depend with some minimal constraint, if they do so only for instrumental reasons, that is, because such behavior is necessary to reap benefits from those animals, this is not the same as recognizing that animals are worthy of moral consideration in their own right.

Successfully cooperating with humans or nonhuman animals doesn't require treating them with proper moral regard; it doesn't even require *thinking* that you ought to treat them well just because of what they are like, independently of whether treating them well advances your own interests. But remember, the big idea that the Two Great Expansions have in common is the idea that some beings are worthy of moral regard in their own right, simply because of what they are like, independently of whether treating them well advances our interests, through cooperation or in any other way.

What is more, in many cases of human–nonhuman animal relationships, it's wrong to speak of "interaction" at all, much less interaction that could meaningfully be called "cooperation." Instead, humans act, and animals are merely acted on. Where one party is an agent and the other only a patient, no cooperation occurs.

So it simply isn't true that for many people the circle of moral regard has expanded to encompass nonhuman animals because

human beings have made a corresponding expansion of their cooperative schemes to include them. Most of the animals we treat badly and should treat better are not even potential cooperators, much less actual cooperators, with us. The idea that the circle of moral regard expands because the circle of cooperation expands is a non-starter if your goal is to explain the Second Great Expansion.

How do people, like Kitcher, who hold that morality is constituted by its function of facilitating cooperation, try to explain the First Great Expansion? They suggest that as global institutions, communication, and travel technologies have linked all humans together, it is becoming vital for successful cooperation that all or at least most of us show some basic regard for all humans. Because human cooperation has gone global, they think, it is now prudentially rational for practically everyone, including the world's richest and most powerful people, to ensure that there is at least significant movement in the direction of greater equality. People who paint this rosy picture don't usually go so far as to proclaim that modern global cooperation requires acknowledging *equal* basic moral status for all. If they did, their claim would be highly implausible. But they do at least suggest that as global cooperation develops, there will be significant pressure to move in the direction of according equal basic moral status to everyone. I wish that were true, but I don't think it is.

Global Cooperation without Equal Regard

Why am I skeptical that further globalization will expand the circle of moral regard sufficiently to move us toward the First Great Expansion? Because it's unlikely that the most powerful people in our world, the people who largely control the structure of global cooperative networks, will suffer, in terms of reproductive fitness or in any way that matters to them, unless they do a lot to reduce the enormous inequalities that now exist. In fact, those people seem to be able to reap plenty of benefits from global cooperation by treating some people—millions of them—as if they don't count at all, or at best as if they count a whole lot less than they themselves do. Sustaining the cooperative schemes from which the best-off benefit

doesn't require moving toward deep inclusion; it doesn't necessitate the First Great Expansion or even make it probable.

So I don't think the expansion of the circle of cooperation explains the expansion of the circle of moral regard, even in the case of humans, leaving aside the case of nonhuman animals. It certainly can't do so for the case of humans who are unable to participate in global cooperative networks either because of debilitating diseases like malaria or because of severe mental retardation or psychological disorders or lack of literacy or numeracy.

It also can't explain something more fundamental: if every human has equal basic moral status, then we ought to try to construct our cooperative schemes so as to ensure that more people *can* participate. In other words, people who say the circle of moral regard expands because the circle of cooperation expands, put the cart before the horse: they think that the expansion of the cooperative scheme to include more people drives the extension of basic equal moral status, when in fact, from the point of view of an inclusive morality of the sort they say they embrace, recognition of the basic equal moral status of all people should drive the expansion of our cooperative schemes to include as many of them as we reasonably can. If you don't see that, you're missing a major message of the disability rights movement: the idea that society ought to make its cooperative schemes more inclusive. For those who have acknowledged the First Great Expansion, the point isn't that if you can cooperate with us, you have equal basic moral status; it's that because you have equal basic moral status, we should try to ensure that you are able to participate in cooperation with us.

The depressing fact of the matter is that a handful of the most powerful nations, led (at least so far) by the United States, has created a global economy that, viewed purely as a smooth-running cooperative scheme, is quite impressive but at the same time highly inequalitarian, profoundly hierarchical. Moreover, inequality seems to be increasing, not decreasing, as the global economy expands, even though absolute poverty is declining. And presumably at least some of those growing inequalities, if they become great enough, are incompatible with genuinely acknowledging in practice the idea that all human beings have the same basic equal moral status.

The big problem for any attempt to explain the expansion of the circle of moral regard as resulting from expansions in cooperation is that even though cooperation has expanded to include most of the world's population in one way or another (at least if you have a thin enough notion of "cooperation"), it isn't true that this expanded cooperation doesn't run smoothly unless everybody is treated as having equal basic moral status. In fact, the existing global cooperative scheme's running smoothly doesn't even seem to require that most people *think* that everyone deserves to be treated as having equal basic moral status. And it certainly doesn't require that the best-off take seriously in their actual behavior the idea of realizing human rights for all.

I'll offer one more obvious and compelling reason to reject the thesis that the circle of regard expands because the circle of cooperation expands, even if we restrict this claim to the First Great Expansion and don't worry about its painfully obvious failure to explain the Second. The circle of cooperation has expanded many times in human history. Two examples come to mind: the Roman Empire, at its peak of power and influence, created a sophisticated cooperative network that encompassed around 70 million people and all the lands bordering the Mediterranean; and in the two decades preceding World War I, what many historians call the first truly global economy appeared. Yet in neither case do we see a clear expansion of the circle of moral regard, at least not anything approaching the First Great Expansion. In fact, both of these dramatic expansions of the circle of cooperation were marked by extreme inequalities: pervasive slavery in the former and Western colonial domination in the latter. Historically, then, we can observe no clear correlation between the expansion of cooperation and the First Great Expansion.

So expanding the circle of cooperators can't by itself explain expanding the circle of moral regard in anything like the strongly egalitarian way the First Great Expansion requires. At most, pointing to the fact that the circle of cooperation has gone global could help explain why human morality would have developed a moral psychology that allows humans to have *shallowly inclusive moralities*, to accord the limited sort of recognition that is required for effective cooperation among strangers with whom it is advantageous for

them to cooperate—something akin to the sort of limited moral recognition that was required for long-distance trade or military alliances in the distant past or perhaps a bit more than that, but surely far less than full recognition of basic equal moral status. At present, the prediction that the further development of global cooperation will decrease inequalities in wealth and political power, much less contribute to the realization of the First Great Expansion, is not supported by evidence about how the global economy works. It might happen, but it might not.

The Deeper Flaw in the Functionalist Understanding of Morality

That's not really the point, however. The point is that even if such a hopeful prediction comes true, it wouldn't explain the First Great Expansion, much less the Second. However much global cooperation expands, it won't include every human being who is worthy of equal basic respect, because there will always be some humans who aren't able to participate in the global economy or indeed in any form of interaction that could be called "cooperation" in anything like the sense in which evolutionary theorists use that term. Nor will it ever include most nonhuman animals.

The mistaken idea that morality is constituted by its original function of facilitating cooperation gives rise to a correspondingly limited understanding of how human beings regard moral status and moral standing: namely, that they restrict them to those beings who are or can be cooperators (or who can disrupt cooperation).

Fortunately, many people nowadays don't have that limited view of what confers moral status and standing. The fact that they don't will remain utterly mysterious so long as one operates with the limited view of morality, namely, that it is just a device whose function is to facilitate cooperation.

The Squeaky-Wheel Model of Moral Progress

Let's consider another way that people who assume that everything interesting about morality can be explained by understanding how it facilitates cooperation try to explain moral progress in the

direction of inclusion. They appeal to a kind of friction or squeaky-wheel theory of progressive change.

The renowned philosopher Peter Railton offers a prime example. He thinks that when the social practices and governmental structures of a society give short shrift to the important interests of some substantial portion of its members, the oppressed sometimes shake things up—disrupt cooperation—and that this process can lead to a new equilibrium, a new state of affairs where the interests that were previously disregarded are at least to some extent realized. Because Railton believes that a morally good society is one in which everyone's interests are taken seriously, he believes that establishing a new equilibrium that serves previously disregarded interests is an instance of moral progress. But there's a catch: he frankly admits that this kind of progressive change will only occur if the people whose interests were previously disregarded *are able to mobilize effectively enough to disrupt the existing cooperative arrangements* (Railton 1986, 200).

When the Wheel Can't Squeak Loudly Enough

The oppressed aren't always able to disrupt the status quo. A large body of literature in political science explains why the oppressed are often not able to solve collective-action problems needed for effective mobilization and consequently can't exert sufficient pressure on the best-off to shake things up and make society work better for them.

The squeaky-wheel or friction model is plausible as a theory for some cases of successful social revolution. It's a far cry, however, from explaining the Two Great Expansions. It clearly tells us nothing about how the Second Great Expansion could ever occur. Barring some enormous evolutionary leap (something like the deeply flawed scenario of *Planet of the Apes*), the animals that humans have usually treated so badly aren't likely to mobilize and disrupt our cooperative schemes. (As much as I'd like to, I can't imagine chickens going on strike in a factory farm, much less a horde of angry pigs storming Hormel's corporate headquarters; can you?)

A Deeper Flaw of Squeaky-Wheel Views

Certain historical cases do fit the squeaky-wheel model—wars of liberation from colonialism come immediately to mind. That is, in some instances, oppressed people have been able to mobilize effectively enough to create a disruption that leads to a new equilibrium where their interests are taken more seriously and to that extent have moved their society in the direction of equality.

Yet it's simply not true that the origination or the spread of the idea of the basic equal moral status of all human beings has always, or even most of the time, been the *result* of the oppressed being able to shake things up sufficiently to disrupt social cooperation. The 800,000 people who were liberated when slavery was abolished in the British Empire in 1834 didn't accomplish that feat by disrupting Britain's economy. Emancipation did disrupt the British economy, but emancipation didn't occur because the slaves disrupted the economy, which would have had to be the case for the squeaky-wheel model to apply. (In later chapters, I'll have a lot to say about the abolition of Atlantic slavery and how it illustrates my theory of the interaction between the moral mind and certain recently constructed human niches.)

The squeaky-wheel theory can explain why power elites should be willing to grant concessions that reduce inequalities *to those who are capable of disrupting cooperation from which the elites benefit*—but *only* to those people. It can't explain why the elites or anybody else would or should care about the moral status of beings, whether human or nonhuman, who aren't capable of disrupting cooperation.

It's also quite clear that sometimes those who achieve progressive moral change do *not* do so by disrupting cooperation. For example, I believe it is evident that the successes of the animal liberation movement have been due more to persuasion (through appeals both to reason and to emotion) than to the limited and probably counter-productive disruptive actions of organizations like PETA.

In fairness to Railton, I should note that at one point he states that progress in the direction of inclusion can occur when those who *represent* the interests of the oppressed (but aren't themselves

oppressed) disrupt cooperation. This is a purely ad hoc move, however, because he does nothing to explain why they would be motivated to do so and how such motivation is compatible with the standard evolutionary moral origins story. This move doesn't explain how thinking of morality simply as something that facilitates cooperation can accommodate the First or Second Great Expansions.

Two Different Understandings of the Basis of Moral Regard: Strategic and Nonstrategic

Views that understand morality solely in terms of how it facilitates cooperation have something in common: they conceive of actual human morality as a way of thinking and acting that confers moral status on beings strictly on the basis of their *strategic capacities*—more precisely, their ability either to contribute to or to disrupt cooperation. As long as you continue to think of morality in that way, as something that is at bottom purely strategic, you will never be able to explain the Two Great Expansions. And if the Two Great Expansions are an important aspect of human morality, then as long as you proceed under that limitation, your understanding of morality will be seriously incomplete.

The Two Great Expansions constitute a radical shift in understanding what the basis of moral standing and equal high moral status is. Instead of grounding moral regard for an individual, whether human or animal, in that being's strategic capacities, that is, his ability to contribute to or disrupt cooperation with us, people who have embraced the Two Great Expansions implicitly ground moral regard in a different way. When moral standing or high moral status are conferred on a strategic basis, they are conferred only as the product of a kind of implicit bargain: I will show moral regard for you if you will show moral regard for me.

When the Two Great Expansions occur, people abandon the purely strategic point of view. Instead of assuming that a being's strategic capacity is what counts, people focus on other characteristics that they believe to be morally significant. For Kantians, the relevant nonstrategic characteristic is the capacity for practical reason

or, as contemporary Kantians usually put it, responsiveness to reasons, the capacity to engage with others in the practice of giving reasons for how we ought to act. For Utilitarians, the relevant non-strategic characteristic is sentience, the ability to experience pleasure and pain.

Sentience and the capacity for practical reason are nonstrategic characteristics, in Utilitarian and Kantian moral theories respectively, in the sense that they are supposed to confer moral standing (or, in the case of practical reason, the highest moral status) on a being independently of whether that being can benefit or harm us, and hence independently of whether it is advantageous for us to recognize or not recognize that being as worthy of moral regard. Any understanding of morality that implies that the basis of moral regard is purely strategic cannot account for the Two Great Expansions. It also can't account for why there are any Kantians or Utilitarians. The transition from a purely strategic to a subject-centered understanding of the grounding of moral standing is, in my judgment, a major instance of moral progress. You can't begin to understand it if your mind is shackled to the Cooperation Dogma.

Does Reducing Inequality Promote More Efficient Cooperation?

I don't want to give short shrift to the idea that movements toward equality or at least toward reductions of the grosser forms of inequality can *sometimes* facilitate better cooperation, in particular cooperation that is more efficient because it depends on the strong motivation of participants to do their bit voluntarily, rather than only on their being threatened with coercion if they don't. The exceptionally creative and systematic evolutionary anthropologist Peter Turchin endorses this idea: he says that equality promotes cooperation—and worries that as inequalities in wealth and power increase in contemporary societies, cooperation may break down (Turchin 2015, ch. 10).

Clearly Turchin is on to something important here; yet I think the claim that equality is required for efficient cooperation, unless it is highly qualified, is clearly false. After all, as I've already emphasized, history shows many examples where expanding cooperation

brought about great gains in the efficiency of production; yet there was great inequality. These historical examples of highly successful yet deeply inegalitarian cooperative schemes didn't last forever (what does?), but it is not clear that the cause, or the main cause, of their demise was inequality.

It may be true, however, that under certain conditions—conditions that *now exist for the first time in history and only in some parts of the world*—some kinds of inequalities can reduce the efficiency of cooperation. That might be true, for example, in societies in which the ideology of equal citizenship for all compatriots has taken root. If people have internalized this ideology, and if some citizens believe they're being denied the rights that constitute equal citizenship, their motivation to cooperate might decline. If there are enough of them, and if there is enough of a decline in their motivation to cooperate voluntarily, that might adversely affect cooperation. But then we are back to Railton's squeaky-wheel theory, and we've already seen that it can't explain either of the Two Great Expansions. Nevertheless, someone like Turchin might insist that the emergence of the ideal of equal citizenship is an important step toward the First Great Expansion.

Enthusiastic Cooperation without Equality

Perhaps; but I think there are two problems with this variant of the idea that a move toward equal status for all human beings comes about because successful cooperation requires it. First of all, people can enjoy the formal status of being equal citizens—possessing on paper, as it were, the same basic rights as other members of the nation-state's community—and still not be accorded full equal basic moral status. Yet they may still cooperate.

Furthermore, I'm aware of no evidence that the ideal of equal citizenship reliably becomes transformed into or leads to the notion that all humans—not just all members of one's own polity—have the same basic moral status. After all, scholars of nationalism, including Eugen Weber in his classic *Peasants into Frenchmen* and more recently Andreas Wimmer in his outstanding *Nationalist Exclusion and Ethnic Conflict*, have argued that the modern idea of

equal citizenship has historically been a force not just for inclusion but also for exclusion—exclusion from basic political rights and access to economic and cultural opportunities for members of the nation-state who are not part of the dominant ethnic group in whose image the concept of the nation is shaped (Weber 1976, 486; Wimmer 2002, 58).

For example, the extension of equal citizenship rights in France after the French Revolution didn't fully apply to some minorities (such as Bretons), or if it did, it required them to abandon their distinctive ethnic or national identity, including their language. In fact, so-called nation building, along with its nationalistic conception of equal citizenship, has historically been nation destroying—forging a modern nation-state meant obliterating all other national identities within a territory in order to privilege only one; and the understanding of what it is to be an equal citizen followed suit. Moreover, nationalism can and often does extend recognition of equal status to all who are regarded as members of the nation while at the same time relegating people of other nations to an inferior status. Nationalism is often a kind of tribalism.

Second, in modern societies, people sometimes believe that they are being accorded equal citizenship status when they aren't. For example, many Americans believe that they are all equal citizens, where this includes having equal political rights, and that they live in a democracy, understood as a society in which all have an "equal say" in how we are governed. Yet in fact they are increasingly living in something approaching an oligarchy, a polity in which the very rich exercise their "equal" political rights much more effectively than the poor and even the middle class and wield greater political power in other ways, behind the scenes of the official political processes. Where a society allows those with such great resources to use them to disproportionately influence political outcomes, it is absurd to say that we all have an equal say in how we are governed. (Do you think your "say" is anywhere near equal to that of the Koch brothers or George Soros?)

Yet this state of affairs apparently hasn't led to a reduction in cooperation, much less a breakdown of it. Ideological thinking can prevent people from seeing just how deep inequalities run and

hence from being motivated to change things to achieve greater equality. I'll have a lot more to say about the obstacles ideologies pose for moral progress in chapter 8.

Perhaps even more importantly, modern consumer culture can make people believe that their material well-being is what matters most and that their lives are improving and will continue to improve in that dimension. When that happens, they may not care so much about whether their equal citizenship is substantial, not merely formal. Contemporary China may be a good example. The majority of the population seems quite willing to participate enthusiastically in social cooperation even though the society is obviously deeply hierarchical and the idea of equal citizenship is patently delusional. Apparently people can contribute enthusiastically to impressive social cooperation in extremely inegalitarian societies, so long as they believe that they, or at least their children, will continue to be better off in material terms. So successful cooperation, even in societies that feature the idea of equal citizenship, doesn't require the sort of robust equality that the First Great Expansion encompasses.

The Bottom Line: Morality Is Not (Now) All about Cooperation

All those complications, however, don't really get at my main point, which is this: even some version of the thesis that successful cooperation requires movement in the direction of equality is true, that still wouldn't account for either of the Two Great Expansions. Not the Second, because most of the animals we now think have moral standing can never be part of our cooperative arrangements; not the First, because it includes the idea that everyone is entitled to recognition of basic equal status *independently of any role they might play in cooperation and hence independently of whether they are able to make their cooperation conditional on being treated with equal regard*. Neither of the Two Great Expansions can be explained as being requirements of cooperation or of efficient cooperation. That conclusion is only surprising if you mistakenly assume that morality, including the basis of equal basic moral status and moral standing, is all about cooperation.

At this point, I want to avoid a misinterpretation of what I'm saying. I'm happy to acknowledge that certain developments in the

ways human beings cooperate were a *necessary* condition for the emergence of the First Great Expansion. More specifically, humans learned to cooperate in ways that resulted in surplus reproductive success—and that success enabled the Great Uncoupling of moralities from the maximization of reproductive success. I have a lot to say about that process in chapters 5 and 6. I’m also sympathetic to the idea that features of the moral mind that evolved because they facilitated cooperation, such as sympathy that reaches beyond kin, and perspective taking, eventually—under the influence of cultural changes—came to play an important role in the First Great Expansion.

Did Moral Consistency Reasoning Cause the Two Great Expansions?

Now let’s consider a different and more promising attempt to explain the Two Great Expansions: the thesis that they came about through the exercise of *moral consistency reasoning*. That explanation is quite popular; among the most prominent thinkers who subscribe to it are the provocative and influential philosopher Peter Singer (Singer 2011, 115–116, 142) and the formidable cultural evolutionary thinkers Hugo Mercier and Dan Sperber (Mercier and Sperber 2017, 314).

The basic idea is that human beings have a capacity for rationally grounded empathy—an ability to take seriously the well-being of other creatures (whether they be humans from other groups or non-human animals) because they have the ability to detect inconsistencies in their moral beliefs and then resolve those inconsistencies in a way that leads them to change their moral judgments and emotional responses.

Here’s an example: you acknowledge that one ought not to inflict pain on other human beings *because pain is bad*. But if that’s why one ought not to inflict pain on humans, then one shouldn’t inflict it on animals, either. That a creature isn’t a human is not a morally relevant difference; what matters is that it feels pain. According to what may be the most sophisticated analysis of moral consistency reasoning available so far, offered by the creative philosophical team of Campbell and Kumar, moral consistency reasoning in this case

involves ironing out a conflict between one's emotionally charged judgment that inflicting pain is wrong, on the one hand, and one's belief that it is acceptable to inflict pain on animals, on the other (Campbell and Kumar 2012, 276). A broader conception of moral consistency reasoning would also cover cases where the conflict is between moral principles, that is, general moral rules. What is common to the narrower and broader conceptions of moral consistency reasoning is that they both help explain how the desire to achieve consistency in one's moral attitudes and beliefs can be an engine for moral change.

It's quite right to emphasize that moral consistency reasoning played a significant role in both of the Two Great Expansions, at least for some people who have achieved that moral reorientation. Yet it's clear that this explanation of how the circle gets expanded is seriously incomplete, because it doesn't explain why moral consistency reasoning sometimes occurs and sometimes doesn't, or why, when it does occur, it sometimes leads toward the Two Great Expansions and sometimes doesn't come anywhere near them. It's a depressing fact that moral consistency reasoning can lead either to progress or to regression. A proper explanation needs to identify the conditions under which moral consistency reasoning occurs and becomes socially and politically potent enough to reorient our conception of moral standing or equal basic moral status.

When Does the Right Kind of Moral Consistency Occur?

Presumably human beings have been capable of moral consistency reasoning for as long as there have been humans or at least as long as humans have had moralities. Or at least they've been capable of moral consistency reasoning much earlier than the beginnings of the Two Great Expansions. For example, in 2 Samuel 12, New International Version of the Bible, we find a clear instance of moral consistency reasoning. Nathan uses a parable about a rich man stealing a poor man's prized possession to trick King David into moral consistency reasoning that leads the king to conclude that he has violated his own moral principles by arranging to have Uriah the Hittite killed so that he can take the man's wife.

That episode of moral consistency reasoning supposedly took place in the late Bronze Age, almost three thousand years ago, but among a people who thought *they* (and nobody else) had an especially high moral status (being God's chosen people). And the story is told in a book that includes several instances in which God orders the utter destruction of other human groups (a.k.a. genocide), including noncombatant women and children, as if they had no moral standing at all, a book that also prescribes unnecessarily painful methods of killing food animals.

So even though some humans used moral consistency reasoning in the remote past, it didn't issue in either of the Two Great Expansions until very recently. Those two enlargements of the circle of moral regard, at least so far as they have reached a social scale and were not limited to a morally precocious minority, are much more recent. To solve our mystery, then, it isn't nearly enough to point out that humans are capable of moral consistency reasoning or that they have sometimes exercised that capacity. We need to know why it's only recently that enough of them have *exercised it in a way that leads to the Two Great Expansions*. The story turns out to be a complicated one; I tell it in chapters 5 and 6.

Here it's worth noting that though certain strains of Hindu culture have long placed significant constraints on the treatment of some animals, this practice doesn't seem to have been the result of moral consistency reasoning. Nor were these moral compunctions grounded in the belief that nonhuman animals had moral standing on their own account.

Although many people today in that tradition have made the Second Great Expansion, the idea originally wasn't that you shouldn't kill or inflict pain on nonhuman animals because they had moral standing in their own right. Rather, the reason you were supposed to avoid treating them badly was that if you did, you'd be reincarnated as one of them. It's one thing to avoid torturing a goat because you might come back as one; quite another to think it's wrong to torture goats because pain is bad and they feel it much as we do and therefore have some sort of moral standing in their own right. The results of people acting on the belief that they should avoid mistreating animals if they want to avoid a bad outcome in the next

reincarnation may have been good from a moral point of view, in the sense that it led to less cruel treatment of nonhuman animals, but it doesn't imply that there was a transition to the Second Great Expansion among the people who thought in that way. This isn't a case of moral consistency reasoning, much less of moral consistency reasoning leading to the conclusion that some nonhuman animals have moral standing in their own right, independently of how our treating them affects us.

What's the point of these excursions into the history of moralities? The point of the example of Nathan's moral jujitsu in the book of Samuel is simply that even if humans not only had the capacity for moral consistency reasoning for a long time but also sometimes actually exercised it, that doesn't explain the First Great Expansion, because that enlargement of the circle of moral regard came a lot later than the first recorded uses of moral consistency reasoning. The point of the Hindu tradition example is that the fact that some cultures for a very long time have included in their moralities some constraints on how nonhuman animals are treated doesn't show that the Second Great Expansion had already come about way back then, if these constraints were not based on the idea that animals count morally in their own right, independently of how our treating them affects our well-being. My main point, however, is that although moral consistency reasoning may well have been necessary for both of the Two Great Expansions to occur, merely pointing to the fact that humans can engage in it isn't sufficient to explain why significant numbers of humans have engaged in it *only rather recently in ways that contributed to the Two Great Expansions*.

Failures of Moral Consistency Reasoning

Moral consistency reasoning has two problems. First, it often doesn't come into play when it should; second, it frequently goes awry. In some ways, the second problem may be even worse than the first. The difficulty is that moral consistency reasoning is often distorted, circumscribed in ways that are arbitrary even from the moral standpoint of those exercising it. When that happens, moral

consistency reasoning can contract the circle of moral regard rather than enlarge it.

Here are a couple of examples. Wealthy folks who associate exclusively with people like themselves are likely to be quick to realize that if something is painful to them, it is also likely to be painful to other people *like them*. But they may be remarkably unaware of, or unempathetic toward, the suffering of poor people. Similarly, Serbs who have been brought up to hate Croats know that it is wrong to torture their fellow Serbs (and most other people, as well); but they may think it is perfectly acceptable to torture Croats, because they think Croats are not at all like them, at least not in ways that are relevant to whether it's permissible to torture them.

Distorted Moral Consistency Reasoning: Garbage In, Garbage Out

Here's one last example of how moral consistency reasoning can facilitate tribalistic morality rather than overcome it. The historian Claudia Koonz, in her exceptionally valuable book *The Nazi Conscience*, notes that during the Third Reich public school teachers were told by Nazi Party officials how important it was to instill in their pupils the right moral values, including the Golden Rule—but with the proviso that it only applied to “racial comrades” (Koonz 2003, 10). The Golden Rule is a marvelous mental trick for expanding the circle of moral regard and a shining example of moral consistency reasoning. Yet engaging in the role-reversal experiment it recommends only produces inclusive results if you *already* have certain beliefs about which differences and similarities are morally relevant and hence should be thought about in a consistent way. If “non-Aryans” lack equal moral status, then there is nothing inconsistent in treating your fellow “Aryans” as equals while treating “non-Aryans” as dangerous, unclean beasts. Simply treating like cases alike doesn't give the right answer if you are mistaken about what the morally relevant likenesses are.

So if you don't believe that the other is like you in morally relevant ways so far as basic moral status is concerned, then you can

engage in moral consistency reasoning till the cows come home and still get answers that do nothing to expand your circle of moral regard. In fact, you may engage in a lot of moral consistency reasoning and act faithfully on its results and never come close to making either of the Two Great Expansions.

So what we need is an account of when and why moral consistency reasoning of a particular sort (the kind that expands the circle of moral regard) occurs and when and why it becomes widespread enough and powerful enough to cause the Two Great Expansions. We need to know why and how moral consistency reasoning is sometimes restricted in certain ways, foreclosing particular results. We also need to know what motivates people to engage in moral consistency reasoning in the first place.

Moral Consistency Reasoning by Cognitively Flawed Beings

More specifically, we need a scientifically based account of how moral consistency reasoning sometimes leads to inclusion, sometimes to exclusion—and why sometimes people don’t engage in it at all, even though they hold inconsistent moral beliefs. We need to learn, from the best available research on errors of reasoning that all normal human beings are prone to and from experiments that confirm the ubiquity of implicit racial bias, how people can fail to see that, according to their own moralities, morally relevant similarities exist between themselves and other creatures, whether human or nonhuman. We also need to understand why people sometimes act on the results of their moral consistency reasoning and sometimes don’t. We need to know how cultural constructions, rooted in our evolved moral nature, can lead us to conceive of the Other in ways that can restrict the scope of moral consistency reasoning.

Above all, if we care about moral progress, we need to know how cultural innovations can help remove those restrictions (for example, how someone who previously thought the Golden Rule only applied to fellow Aryans could come to think it applies more widely). In chapters 5 and 6, I explain how cultural innovations have only recently produced human-made niches in which significant numbers of people have become motivated to engage in moral

consistency reasoning that leads all the way to the conclusions that all human beings have an equal basic moral status and that animals have moral standing—and to act on that conclusion, if only imperfectly.

Evolved Limitations on Who Is Seen as Being Worthy of Moral Regard

I mentioned in the introduction that some evolutionary psychologists have done fascinating work on how very young children learn to cooperate morally or, if you will, how they become good at cooperation by becoming moral. Among the most fruitful researchers of this sort are Michael Tomasello and his graduate students who have gone on to become outstanding researchers in their own right. Unfortunately, as I suggested earlier, from a moral perspective that cherishes inclusion, this research has a negative side: one of their chief findings is that children learn to cooperate morally, or to be good cooperators because they are moral, by picking and choosing with whom to cooperate on the basis of just the sort of superficial resemblances that many people now regard as morally irrelevant and that can lead moral consistency reasoning astray. These are also the same sort of resemblances that, if missing, led our early human ancestors to react with fear, distrust, or aggression toward other human beings, if the standard evolutionary moral origins story is largely correct.

Recall that, according to the standard moral origins story, in the EEA it was a matter of life or death to be able to detect whether another being of human form was one of Us or one of Them, and detection was generally based on at-a-glance similarities and differences, like clothing, hairstyle, hair texture, language, and skin tone. Well, Tomasello's team finds that young children today use the same sorts of cues.

The good news is that in Tomasello's lab, as in the EEA, those sorts of cues often provide reliable indicators of whether trying to cooperate with somebody is likely to turn out well. The bad news is that those cues have nothing to do with rightly determining whether a being has equal basic moral status. From that same evaluative

perspective, the even worse news is that when people don't distinguish between what is relevant to basic moral status and what isn't, those cues can lead them to act in exclusive (tribalistic), rather than inclusive, ways.

If you think that a being with dark skin is a different kind of creature, morally speaking, from light-skinned folks like you, then your moral consistency reasoning about dark-skinned people will not lead you to conclude that you should treat them as your moral equals. And your ability to identify with their suffering, to empathize with them in the way you would with someone you regard as your moral equal, may be blocked. Your moral consistency reasoning won't lead you to expand your circle of moral regard. Garbage in, garbage out.

Tomasello and company don't claim to be providing an explanation of how human moralities could have become deeply morally inclusive. Like other evolutionary scientists, they haven't pretended to tackle that problem; it simply isn't on their research agenda so far. In fact, their work so far doesn't even fully explain how *shallowly inclusive moralities* could have come about, because it only explains inclusiveness that is restricted to coethnics broadly understood, people of the same cultural group who look and speak alike or resemble each other in various superficial ways. The psychological mechanisms that these researchers observe in small children could have been at work in the EEA, enabling limited moral recognition of strangers, if they were coethnics. Something more would be needed, however, to explain the transition to shallowly inclusive morality, which involves extending the circle of moral regard to include people of other ethnicities or, in Tomasello's terms, different cultural groups.

Progress through Failures: How Detective Stories Work

Finding out what *doesn't* solve the Big Puzzle, which is what this chapter has done, is useful. So before moving on to a more productive line of investigation, let's review what we've learned about what doesn't work. One thing that doesn't work is to proceed by relying on the Cooperation Dogma, the assumption that morality is—not

just was originally—all about facilitating cooperation. Another is simply pointing to the fact that moral consistency reasoning *can* help expand the circle of moral concern. That's not enough: you also have to provide an account of the kind of the moral consistency reasoning that leads to the Two Great Expansions and an explanation of why humans sometimes engage in it and often don't.

We need an even more fundamental kind of rethinking if we are to make much headway: it's time to go back to the moral origins story that generated the Big Puzzle in the first place. What if we got that story wrong? What if the moral mind as it evolved in the EEA wasn't as hostile to inclusive moralities as a lot of people think? Some of the best detective stories feature a point in the plot where the investigator backtracks, revisiting the initial assumptions she made when she first began the investigation.

Things can go wrong in two ways if you rely on a story about the origins of morality as a guide to what morality is now like and what it can become. The first is by getting the origins story wrong. The second is by getting it right but failing to realize that things can happen later that make moralities quite different from what they were originally. My aim is to expose and to avoid both sorts of errors. In the next chapter, I show that the origins story as presented in chapter 2 is seriously inaccurate—and in ways that open up possibilities for the Two Great Expansions. In chapters 5 and 6, I step through that open door, setting out the key elements of an explanation of how some human moralities became (at a rather late date) deeply inclusive, by pointing out some powerful new developments in human culture and psychology that have only occurred in the last three hundred years or so.

