

V. VALUE

The idea of teleology seems to imply some kind of value in the result toward which things tend, even if teleology is separated from intention, so that the result is not the goal of an agent who aims at it. But let me set these cosmological speculations aside for the moment, and take up the more mundane existence and role of value in our lives. Real value – good and bad, right and wrong – is another of those things, like consciousness and cognition, that seem at first sight incompatible with evolutionary naturalism in its familiar physicalist form. In the previous chapter I have commented briefly on how evolutionary theory might be reconciled with realism about value, on the model of a supposed reconciliation of scientific realism with evolutionary theory. I expressed skepticism about the adequacy of such a reconciliation in either case, but I now want to take up the problem of value and reasons for action in more detail.

The problem of the place of value in the natural world includes but goes beyond the problems of the place of consciousness and of cognition in general. It is clear that the existence of value and our response to it depend on consciousness and cognition, since so much of what is valuable consists in or involves conscious experience, and the appropriate responses to what is good and bad, right and wrong, depend on the cognitive recognition of reasons for and against. Practical reasoning is a cognitive, largely conscious process. I have argued so far that the reality of consciousness and cognition cannot be plausibly reconciled with traditional scientific naturalism, either constitutively or historically. I believe that value presents a further problem for scientific naturalism – that even against the background of a world view in which consciousness and cognition

are somehow given a place in the natural order, value is something in addition, and it has consequences that are comparably pervasive.

To explain this claim it is necessary to explain what is meant by the reality of value – something far less transparent than the reality of consciousness, and also less clear than the reality of cognition or reason in general. One of the difficulties in this area is to describe a form of value realism that is not weighed down by metaphysical baggage, but that is still clearly distinguishable from a sophisticated subjectivist conception of value. Value realism is highly controversial because subjectivist accounts of value are not inherently implausible in the way that subjectivist accounts of science or physicalist accounts of consciousness are. So let me try to explain the kind of value realism that presents the problem, and how it is to be distinguished from sophisticated forms of subjectivism in the Humean tradition, which allow for antirealist versions of moral objectivity, interpersonal justification, and practical reason.

In simple terms, the subjectivist position that I will contrast with realism is that evaluative and moral truth depend on our motivational dispositions and responses, whereas the realist position is that on the contrary, our responses try to reflect the evaluative truth, and can be correct or incorrect by reference to it. For the sake of intellectual and discursive economy, I am going to treat subjectivism as a view about truth and truth conditions. But what I say about it, and about the contrast with realism, should be taken to apply also to the equally common expressivist forms of subjectivism, derived from Hume, according to which value judgments express certain kinds of attitudes or feelings of the subject, rather than being true or false in virtue of such attitudes. Such views can usually be put in terms of assertability conditions rather than

truth conditions. Subjectivism about values and morality comes in many varieties, and I hope it will be possible to treat the basic issue without distinguishing among them.

The contrast between subjectivism and realism with regard to the dependence of truth on our responses does not apply to the whole of value. In particular, realists can agree with subjectivists that the value of basic experiences of pleasure and pain, for example, is inseparable from our natural responses of attraction and aversion to them. In these cases, for a realist, appearance and reality coincide. To the question, “Would pain be bad even if we didn’t mind it?” the answer is “No” (in fact it wouldn’t even be pain). It is only when we move to the evaluation of absent experiences – those in the future or those of other subjects – or to judgments about how to deal with possibilities involving multiple experiences, perhaps with conflicting values, or to judgments about the value of things other than experience, that realism and subjectivism will give clearly divergent accounts. The subjectivist position is that the right answer depends on our attitudes and dispositions; the realist position is that our judgments attempt to identify the right answer, and to bring our attitudes into accord with it, whether the question is about pain or anything else.¹

The most plausible forms of subjectivism rely on some variant of Hume’s conception of the passions, including the moral sense. On this type of view, value judgments in general and moral judgments in particular are grounded in aspects of the motivational system more sophisticated and reflective than the basic appetites and

¹ I leave aside the large topic of contingent normative truths, which both realists and subjectivists about value can agree are dependent on our collective responses. Anyone who has lived long enough to watch the language drift away from him will know what I mean. “Disinterested” is coming to mean “uninterested,” “precipitous” is coming to mean “precipitate,” “criteria” is becoming a singular, and “I can’t help but think” and “I could care less” are ubiquitous. Reading with daily exasperation canonical publications like *The New York Times*, I have to recognize my unshakeable conviction that these are mistakes as the product of a contingent set of linguistic norms internalized years ago, which I cannot get rid of (and wouldn’t even if I could, since they are of course correct).

instinctive emotions. Prudential judgments are the manifestation of a calm passion of temporally impartial self-interest that generates an equal desire or aversion for future and for present benefit and harm to oneself. Moral judgments manifest a sentiment of impartial benevolence, or in more complex cases an attachment to practices or institutions that advance the general welfare, or the good of everyone. The details of the moral sentiment may be complicated, and it may be subject to social modification, but the essential point is that value judgments are nothing more than the expression of such a sentiment: They can be said to be correct or incorrect only by reference to it, together with the nonevaluative facts which are the object of the judgment. What determines the truth of a value judgment is what would result from the calm, reflective, and fully informed operation of these motivational dispositions.

A realist position, by contrast, denies that the truth depends on our dispositions (above the level of immediate feelings like pleasure and pain), and holds that when our value judgments are correct, it is because our dispositions are in accord with the actual structure and weight of values in the case at hand. A judgment that one should not impose serious harm on someone else for the sake of slight benefit to oneself, for example, is based on the recognition that the reason against imposing the harm is much stronger than the reason for pursuing the benefit, and that the fact that the harm would be suffered by someone else is not a reason to disregard it.

2. But what makes it the case that the interests of others provide us with reasons for action, if not that we are disposed on reflection to be motivated by some degree of benevolence? That question is the natural expression of the subjectivist's skepticism

about realism. It is also the first step on the path to a fatal misinterpretation of realism – the misinterpretation that turns it into a metaphysical theory. The question implies that something other than value must make value judgments true or false – something in the world in addition to the empirical facts that count for and against action, such as the fact that if I don't step on the brakes I will run over a dog. Subjectivists think this question for the realist is reasonable because their own theory has an answer to it: They find the ground for the truth of value judgments in psychological facts about human motivational dispositions – something more fundamental than values. (For expressivists the question is avoided by denying that value judgments have truth value, but they too find the grounds for these judgments in motivational dispositions.)

This kind of explanation of value is precisely what realists deny. More important, in my conception of moral and evaluative realism, realists do not replace the psychological answer to the question posed by subjectivists with a different answer. Realism is not a metaphysical position. It is not even a naturalist position. It does not maintain that value judgments are made true or false by anything else, natural or supernatural. Of course natural facts are what make value judgments true in the sense that they are the facts that provide reasons for and against action.² In that sense the fact that you will run over a dog if you don't step on the brakes makes it the case that you should step on the brakes. But the general moral truth that licenses this inference – that it counts in favor of doing something that it will avoid grievous harm to a sentient creature – is not made true by any fact of any other kind. It is nothing but itself.

² See T. M. Scanlon, *What We Owe to Each Other* (Cambridge, Mass.: Harvard University Press, 1998) pp. 56-57.

The issue between realism and subjectivism is not about the contents of the universe. It is an issue about the order of normative explanation. Realists believe that moral and other evaluative judgments can often be explained by more general or basic evaluative truths, together with the facts that bring them into play (the fact about the dog). But they do not believe that the evaluative element in such a judgment can be explained by anything else. That there is a reason to do what will avoid grievous harm to a sentient creature is, in a realist view, one of the kinds of things that can be true in itself, and not because something of a different kind is true. In this it resembles physical truths, psychological truths, and arithmetical and geometrical truths.

Instead of explaining the truth or falsity of value judgments in terms of their conformity to the reflective deliverances of our motivational dispositions or moral sense, as the subjectivist does, the realist explains our moral sense as a faculty that aims to identify those facts in our circumstances of choice that count for and against certain courses of action, and to discover how they combine to determine what course would be the right one, or what set of alternatives would be permissible or advisable and what others ruled out. The facts that provide these reasons and justifications are just the familiar ones about what would happen if one did this or that, who would benefit, who would be harmed, and so forth.

But although realism does not add anything to what a subjectivist believes to exist in the world, it does hold that certain truths that subjectivists think have to be grounded in something else do not have to be so grounded, but are just true in their own right. After all, whatever one's philosophical views, so long as there is such a thing as truth there must be some truths that don't have to be grounded in anything else. Disagreement over

which truths these are defines some of the deepest fault lines of philosophy. To philosophers of an idealist persuasion it is just as self-evident that physical facts can't just be true in themselves, but must be explained in terms of actual or possible experience, as it is self-evident to those of a physicalist persuasion that mental facts can't just be true in themselves, but must be explained in terms of actual or possible behavior, functional organization, or physiology. The issue over moral realism is of the same kind. Someone who finds an unanswerable challenge in the question, "But what *is* a reason?" thereby reveals a restrictive assumption of the largest scope about what kind of basic truth there can be.

In every area of thought we must rely ultimately on our judgments, tested by reflection, subject to correction by the counterarguments of others, modified by the imagination and by comparison with alternatives. Antirealism is always a conjectural possibility: The question can always be posed, whether there is anything more to truth in a certain domain than our tendency to reach certain conclusions in this way, perhaps in convergence with others. For that reason the intuitive conviction that a particular domain, like the physical world, or mathematics, or morality, or aesthetics, is one in which our judgments are attempts to respond to a kind of truth that is independent of them, may be impossible to establish decisively. Yet it may be very robust all the same, and not unjustified.

To be sure, there are competing subjectivist explanations of the appearance of mind-independence in the truth of moral and other value judgments. One of the things a sophisticated subjectivism allows us to say when we judge that infanticide is wrong is that it would be wrong even if none of us thought so, even though that judgment is still

ultimately grounded in our responses. However I find those quasi-realist, expressivist accounts of the ground of objectivity in moral judgments no more plausible than the subjectivist account of simpler value judgments. These epicycles are of the same kind as the original proposal: They deny that value judgments can be true in their own right, and this does not accord with what I believe to be the best overall understanding of our thought about value.

There is no crucial experiment that will establish or refute realism about value. One ground for rejecting it, the type used by Hume, is simply question-begging: If it is supposed that objective moral truths can exist only if they are like other kinds of facts – physical, psychological, or logical – then it is clear that there aren't any. But the positive support for realism can come only from the fruitfulness of evaluative and moral thought in producing results, including corrections of beliefs formerly widely held, and the development of new and improved methods and arguments over time. The realist interpretation of what we are doing in thinking about these things can carry conviction only if it is a better account than the subjectivist alternative, and that is always going to be a comparative question, as it is about any other domain, whether it be mathematics or science or history or aesthetics.

3. All this is a prelude to the larger question: What are the implications for the natural order of different conceptions of value? My denial that realism about value is a metaphysical postulation of extra entities or properties might suggest that realism has no implications for the natural order. However, that is not the case. In essence, I agree with Sharon Street's position that moral realism is incompatible with a Darwinian account of

the evolutionary influence on our faculties of moral and evaluative judgment.³ Street holds that a Darwinian account is strongly supported by contemporary science, so she concludes that moral realism is false. I follow the same inference in the opposite direction: Since moral realism is true, a Darwinian account of the motives underlying moral judgment must be false, in spite of the scientific consensus in its favor.

But the implications are much larger than that. Since we are evidently the products of evolution, and ultimately of a cosmic process that led to the development first of unicellular organisms and then of conscious agents, before eventually producing intelligent beings capable of value judgments, the conception of the natural order that made this process possible must be expanded. An adequate conception of the cosmos must contain the resources to account for how it could have given rise to beings capable of thinking successfully about what is good and bad, right and wrong, and discovering moral and evaluative truths that do not depend on their own beliefs. This is analogous to the previously defended implications for the natural order of the existence of consciousness and cognition, but it goes further.

Some comment is called for on the strange, category-jumping nature of this dispute. Street's argument relies on an empirical scientific claim to refute a philosophical position in metaethics. I, perhaps even more strangely, am relying on a philosophical claim to refute a scientific theory supported by empirical evidence. But I do not think the movement of thought is inappropriate in either case. Value judgments and moral reasoning are part of human life, and therefore part of the factual evidence about what

³ Throughout the discussion, I will continue to speak in very general terms of Darwinian explanation, even though Darwinian evolutionary theory is immensely complex and the subject of serious internal controversies. I believe that for present purposes the complexities and disagreements within Darwinian theory do not matter.

humans are capable of. The interpretation of faculties such as these is inescapably relevant to the task of discovering the best scientific or cosmological account of what we are and how we came into existence. What counts as a good explanation depends heavily on an understanding of what it is that has to be explained.

Let me begin with an outline of Street's argument that moral realism is incompatible with Darwinism. This is offered in her paper, "A Darwinian Dilemma for Realist Theories of Value,"⁴ which is quite intricate and addresses more than one form of realism. But the argument is explicitly applied to what she calls *non-naturalist* versions of value realism, according to which evaluative facts or truths are not reducible to any kind of natural facts, and do not require any mysterious "extra" properties in the world, but are irreducibly normative facts or truths. That is the kind of realism I have described above.

Street points out that if the responses and faculties that generate our value judgments are in significant part the result of natural selection, there is no reason to expect that they would lead us to be able to detect any mind-independent moral or evaluative truth, if there is such a thing. That is because the ability to detect such truth, unlike the ability to detect mind-independent truth about the physical world, would make no contribution to reproductive fitness. It is not at all implausible that the characteristically evaluative motivational dispositions of human beings, and their capacities for practical reasoning and for interpersonal convergence on practices and forms of justification, can be to a significant extent explained, either directly or indirectly, by Darwinian natural selection. But so understood, it is completely irrelevant

⁴ *Philosophical Studies* (2006) 127:109-166.

whether these faculties enable us to detect mind-independent moral truth, should there be such a thing, or whether they lead us radically astray.

Street observes that the natural Darwinian explanation of the motives and dispositions that form the starting points of our value judgments, and which we can then modify through the process of reflective equilibrium, is that they have contributed to reproductive fitness not only by aiding individual survival but by promoting the nurture and care of children, deterring aggression, and making social cooperation possible. The mind-independent truth of the operative judgments has no role to play in the Darwinian story: So far as natural selection is concerned, if there were such a thing as mind-independent moral truth, those judgments could be systematically false.

The same cannot be said about our factual judgments. If there is a mind-independent physical world, the systematic inability to detect the truth about it would be disastrous for our reproductive fitness. Realism about the physical world is a fundamental aspect of the Darwinian explanation of our perceptual and cognitive faculties, as well as of our motives and capacities for action. But realism about value is irrelevant.

This requires me to revisit something I said in the previous chapter. There I offered a way in which moral realism could be combined with a Darwinian account of our cognitive faculties, including the faculty of practical reason. I said that if we can understand our prereflective impressions of value – instinctive attractions and aversions, inclinations and inhibitions – as appearances of real value, then the cognitive process of discovering a systematic and consistent structure of general reasons and practical and moral principles can be thought of as a way of moving from appearance to reality in the

normative domain. To that extent it would be analogous to the Darwinian account of scientific reason, realistically construed. I then went on to make an argument of a different kind against the acceptability of a Darwinian account of reason; but let me set that aside for the moment, because I want to concentrate on a specific problem for value realism that is due to a difference between the prerational data from which factual reasoning and practical reasoning begin. For convenience I will limit the discussion to pleasure and pain, though there are other examples.

The big disanalogy between the two cases from a Darwinian standpoint is between the prereflective appearances of value and the prereflective appearances of the physical world.⁵ A realistic conception of the latter, for example of visual perception, is essential to the Darwinian story: Vision confers reproductive fitness because it enables us to see what is out there, which is necessary for all kinds of successful functioning. By contrast, the real badness of pain and the ability to recognize that badness are completely superfluous in a Darwinian explanation of our aversion to pain. The aversion to pain enhances fitness solely in virtue of the fact that it leads us to avoid the injury associated with pain, not in virtue of the fact that pain is really bad. So far as natural selection is concerned, pain could perfectly well be in itself good, and pleasure in itself bad, or (more likely) both of them might be in themselves valueless -- though we are naturally blind to the fact.

A realist, unlike a subjectivist, cannot dismiss these possibilities as meaningless but must instead believe that they are false – just as a realist about the physical world believes that Descartes’ evil demon hypothesis about visual perception is not meaningless but false. A Darwinian account of visual perception entails that it gives us information

⁵ Cf. Street, “A Darwinian Dilemma,” pp. 123-125 & Fn. 35.

about the external world and that the evil demon hypothesis is false. A Darwinian account of the origin of our basic desires and aversions, by contrast, has no implications as to whether they are generally reliable perceptions of judgment-independent value, or whether indeed there is such a thing.

4. So I am in agreement with Street that, from a Darwinian perspective, the hypothesis of value realism is superfluous – a wheel that spins without being attached to anything. And if that is true for our most basic responses, it is also true for the entire elaborate structure of value and morality that is built up from them by practical reflection – just as scientific realism would be undermined if we abandoned a realistic interpretation of the perceptual experiences on which science is based. Even a system based on the maintenance of coherence or consistency among one’s responses does not need the idea of mind-independent truth about value (as opposed to logic), if the responses that provide its original content refer to nothing beyond themselves.

Nevertheless I remain convinced that pain is really bad, and not just something we hate, and that pleasure is really good, and not just something we like. That is just how they glaringly seem to me, however hard I try to imagine the contrary, and I suspect the same is true for most people. That doesn’t mean our visceral responses are infallible, any more than our prereflective visual perceptions are. They are merely the starting points for the exploration of a domain that may require extensive practical and moral reasoning to understand. On the Darwinian account, this must be regarded as an illusion – perhaps an illusion of objectivity that is itself the product of natural selection because of its

contribution to reproductive fitness. However, in my case the scientific credentials of Darwinism are not enough to dislodge the immediate conviction that it is not an illusion.

But what is the realist alternative? Describing it is tricky, since it is obvious that biologically hard-wired pleasure and pain play a vital role in the fitness of conscious creatures even if their objective value doesn't. The realist position must be that these experiences which have desire and aversion as part of their essence also have positive and negative value in themselves, and that this is evident to us on reflection, even though it is not a necessary part of the evolutionary explanation of why they are associated with certain bodily episodes, such as sex, eating, or injury. They are adaptive, but they are something more than that. While they are not the only things that have objective value, these experiences are among the most conspicuous phenomena by which value enters the universe, and the clearest examples through which we become acquainted with real value.

In the realist interpretation, pleasure and pain have a double nature. In virtue of the attraction and aversion that is essential to them, they play a vital role in survival and fitness, and their association with specific biological functions and malfunctions can be explained by natural selection. But for beings like ourselves, capable of practical reason, they are also objects of reflective consciousness, beginning with the judgment that pleasure and pain are good and bad in themselves and leading on, along with other values, to more systematic and elaborate recognition of reasons for action and principles governing their combination and interaction, and ultimately to moral principles.

If the faculties that generate these judgments could also be given a Darwinian explanation, then the realist interpretation would be refuted, for we would then have no reason to regard them as discoveries of what is true independent of our judgments. The

objective goodness of pleasure and badness of pain, and the objective truth or falsity of all our more complex value judgments, would be completely irrelevant to the understanding of these faculties or what we are doing when we exercise them. Therefore on the realist view, although the association of pleasure and pain with sex and injury can be explained by natural selection, their objective value, our capacity to recognize it, and all that follows from this cannot be.

Street concludes that realism cannot be right; I conclude that something is missing from Darwinism, and from the standard biological conception of ourselves. Street's negative conclusion comes with a positive alternative, a form of constructivism.⁶ Unfortunately my negative conclusion leaves the positive alternative quite mysterious, and unless it is possible to say more, comparison between the competing views of how to respond to the shared hypothetical premise will be difficult.

5. As with the placement of consciousness and cognition in the natural order, the problem of value and the natural order has both a constitutive and a historical aspect. The constitutive question is about our nature: What kind of beings are we, if realism is true and we do indeed recognize and respond to values and practical reasons that are not just the products of our own responses? The historical question is about our origins: What must the universe and the evolutionary process be like, to have generated such beings? Both these questions seem to require some alternative to physicalist naturalism and its Darwinian application in biology, but what are the possibilities?

⁶ See "Constructivism about Reasons," *Oxford Studies in Metaethics* vol. 3, ed. Russ Shafer-Landau (Oxford: Clarendon Press, 2008).

The most conspicuous constitutive consequence of realism would be that human beings are able not only to detect but to be motivated by value. In the case of basic experiential values such as the goodness of pleasure and the badness of pain, an instinctive motivation is built into the experience itself: The desire that it continue is part of pleasure and the desire that it stop is part of pain. But when we think of pleasure or pain that is not actually present – either our own experience at other times or the experience of others – this is not the case. Yet we can be motivated by the recognition that pain is bad, and that there is reason to do what will prevent it, whether for ourselves or for others. Such considerations can get us to resist the immediate, built-in motivation of present pleasure or pain, giving it only its objective value. And of course the same is true of much more complex values such as honesty and dishonesty, justice and injustice, loyalty and betrayal.

What is this motivational capacity? As I have said, it should not be construed in terms of an extra metaphysical component of the world, which exercises a causal influence on us. The features of the world that confer value and provide reasons are ordinary facts about the experiences of people, their relations to one another, and the implications for people's and other creatures' lives of different possible courses of action. A reason for action is a fact, such as the fact that aspirin will cure your headache, and its being a reason is just its counting in favor of your taking aspirin or my giving you some. My capacity to respond to real values is the capacity, for example, to be motivated to give you aspirin for that reason – because I know that it will cure your headache and I recognize that that counts in favor of it, because headaches are bad. We do this kind of thing all the time. What does it mean?

I believe it involves a conscious control of action that cannot be analyzed as physical causation with an epiphenomenal conscious accompaniment, and that it includes some form of free will – though it is, as always, very obscure what sense to give to that notion. I respond consciously to value when I decide to give you aspirin because I learn that you have a headache and I know that aspirin will make it better. Of course I want your headache to go away, but that too is the result of my recognition that headaches are bad. The explanation of my action refers to these facts about headaches and aspirin in their status as reasons – as counting for and against doing certain things. It is through being recognized as reasons by a value-sensitive agent that they affect behavior.

My convictions about the traditional problem of free will are incompatibilist, but I am not sure that issue has to be settled for present purposes. If there is a way that conscious motivation by reasons for action can be reconciled with causal determinism of action, either physical or psychological, then value realism may be compatible with determinism. For the moment, I wish only to insist that it is not compatible with a Darwinian conception of how the sources of our motives are determined.

If we leave the issue of determinism aside, the distinctive conception of human beings that is implied by value realism is that they can be motivated by their apprehension of values and reasons, whose existence is a basic type of truth, and that the explanation of action by such motives is a basic form of explanation, not reducible to something of another form, either psychological or physical. I give you aspirin because I know it will relieve your headache, thereby manifesting my recognition that this fact counts in favor of the action. And this is not merely a superficial description of something else, which is the real, underlying explanation.

Human action, in other words, is explained not only by physiology, or by desires, but by judgments. We are the subjects of judgment-sensitive attitudes, in Scanlon's phrase,⁷ and those judgments have a subject matter beyond themselves. We exist in a world of values and respond to them through normative judgments that guide our actions. This, like our more general cognitive capacities, is a higher development of our nature as conscious creatures. Perhaps it includes the capacity to respond to aesthetic value as well – construed realistically as a judgment-independent domain which our experiences and judgments reveal to us.

If we ask again in this case about the choice between reductive and emergent explanations, it seems very unlikely that the exercise of consciousness in evaluative judgments and practical reasoning will lend itself to a reductive analysis, even if a reductive monism explains the existence of consciousness itself, in complex organisms. As with cognition in general, the response to value seems only to make sense as a function of the unified subject of consciousness, and not as a combination of the responses of its parts. In that respect it is different from those experiences of pleasure and pain that provide some of the most important raw material of value: They might be explained reductively through some form of psychophysical monism. But practical reasoning and its influence on action involve the unified conscious subject, and that suggests an emergent answer to the constitutive question. On the other hand, the problem of how to reconcile the unity of the subject, such as it is, with a reductive account of the mental afflicts psychophysical monism at every level. Perhaps at this extravagantly speculative stage the part-whole problem should not rule out reductivism for practical reason any more decisively than it does for taste sensations.

⁷ *What We Owe to Each Other*, p. 20.

The most important metaphysical aspect of a realist view of practical reason is that consciousness is not epiphenomenal and passive, but that it plays an active role in the world. This is part of the ordinary view we have of ourselves, but it adds to the mere psychophysical irreducibility of experiential consciousness a further respect in which the physicalist form of naturalism fails to provide a complete account of the nature of human beings.⁸ Whether practical reason is emergent or reducible to activity at the microlevel through some form of psychophysical monism, value realism requires consciousness to be active, and rules out epiphenomenalism in human action. (A subjectivist view of value may also be incompatible with epiphenomenalism, but that is far from clear: Subjectivism might be consistent with the possibility that our conscious value judgments are all side-effects of the physiological causes of action.)

Setting this issue aside, we come back to the essential contrast between a subjective and a realist conception of value. Subjectivism interprets our value judgments as an outgrowth and elaboration, made possible by our linguistic and rational capacities, of our natural motivational dispositions, nothing more. Realism interprets them as the result of a process of discovery, starting from initial appearances of value that are comparable to perceptual beliefs and moving (we hope) toward a better understanding of how we should live. If realism is true, practical reason in this sense is one of our cognitive faculties.

6. I have not suggested an answer to the constitutive question from the standpoint of realism, only indicated the conditions such an answer would have to meet. But let me

⁸ For a persuasive account of the relation between consciousness and free will, see “A Plain Person’s Free Will” by David Hodgson, *Journal of Consciousness Studies* (2005) 12:3-19.

now try to say something equally inconclusive about the historical question. It is clear that unlike realism, subjectivism lends itself, at least speculatively, to a Darwinian account of how creatures capable of having values might have arisen. At least this question does not add anything to the difficulties for such an account already presented by the problems of consciousness and of reason in general. Value begins, on this view, from our desires and inclinations, which are natural facts of animal and human psychology, and higher-level value judgments are motivational elaborations from this base, generated by experience, reflection, and culture.

On a realist view, by contrast, the historical question is much more obscure, partly because the result is obscure. We want to know what has to be added to the standard Darwinian picture to account for the appearance through evolution of creatures like us, who can control their actions in response to reasons. While this capacity must be consistent with the influence of natural selection in that it is not inimical to reproductive fitness, and therefore not liable to be extinguished by natural selection, its appearance on the evolutionary menu would have to be explained by something else.

Value realism must make sense of the fact that the biological evolutionary process and the physical and chemical history that preceded it have given rise to conscious creatures, to the value that fills their lives and experiences, and ultimately to self-conscious beings capable of judgment-sensitive attitudes who can respond to and be rationally motivated by their awareness of those values. The story includes huge quantities of pain as well as pleasure, so it does not lend itself to a simple teleological interpretation. Nevertheless, the development of value and moral understanding, like the development of knowledge and reason and the development of consciousness that

underlies both of those higher-order functions, forms part of what a general conception of the cosmos must explain. As I have said, the process seems to be one of the universe gradually waking up.

What is the actual history of value in the world, so far as we are aware of it? In the broadest sense, it seems to coincide with the history of life. First, with the appearance of life even in its earliest forms, there come into existence entities that *have* a good, and for which things can go well or badly. Even a bacillus has a good in this sense, in virtue of its proper functioning, whereas a rock does not. Eventually in the course of evolutionary history there appear conscious beings, whose experiential lives can go well or badly in ways that are familiar to us. Later some descendants of these beings, capable of reflection and self-consciousness, come to recognize what happens to them as good or bad, and to recognize reasons for pursuing or avoiding those things. They learn to think about how these reasons combine to determine what they should do. And finally they develop the capacity to think about reasons they may have that do not depend only on what is good or bad for themselves. This begins with the lives of other beings like themselves, but the question can be extended to good and bad wherever it is found, in the lives of other conscious creatures, and perhaps even in forms of life devoid of conscious experience. Reasons for action apply only to beings with reason, but the concept of value has a much wider range of application. Only beings capable of practical reason can recognize value, but once they recognize it, they find it in the lives of creatures without practical reason. In the broadest sense it is probably coextensive with life, though how much of this value we humans have any reason to care about is a question I will leave open. It seems too simple to hold that only the value in conscious life generates reasons;

cf. Scanlon's example of cutting down a great old tree just to try out one's new chain saw.⁹

Two things must be explained to answer the historical question about value: first, the appearance of value in the myriad forms it takes in the variety of lives capable of having a good; and second, the appearance of reasons for action and of those beings capable of recognizing them and acting on them. The first would be explained by whatever explains the existence of life, including the place in animal life of consciousness in its many forms. The second requires something more. Does it require more than an explanation of the appearance of reason in a general sense? I believe so, because what has to be explained is the appearance of the capacity to recognize and respond to reasons for action, and not just a general cognitive capacity. Practical reason is a development of the motivational system and of the will, not merely a development of the system for forming beliefs.

If value is tied to life, its content will depend on particular forms of life, and the most salient reasons it gives us will depend, even in a realist conception, on our own form of life. This is how a realist account can accommodate one of the things that make subjectivism seem most plausible, namely the fact that what we find self-evidently valuable is overwhelmingly contingent on the biological specifics of our form of life. Human good and bad depend in the first instance on our natural appetites, emotions, capacities, and interpersonal bonds. If we were more like bees or lions what seems good to us would be very different, a point that Street emphasizes.¹⁰

⁹ *What We Owe to Each Other*, p. 68.

¹⁰ "Imagine, for instance, that we had evolved more along the lines of lions, so that males in relatively frequent circumstances had a strong unreflective evaluative tendency to experience the killing of offspring that were not his own as 'demanded' by the circumstances, and so that females, in turn, experienced no

It would be a mistake to try to find a common denominator such as pleasure and pain to accommodate in a realist conception the diverse values that are generated by all the actual, not to mention imaginable, forms of life. Instead, value must be seen as pluralistic: The domain of real value, if there is such a thing, is as rich and complex as the variety of forms of life, or at least of conscious life. Just as most of these lives are only dimly accessible to our understanding from the inside, so the value they generate, positive and negative, is largely beyond our full appreciation. It is also unclear how far the reasons generated by those values can reach – as is true even of the values we recognize in forms of human life other than our own. And who knows what unimaginable forms of life and their associated value exist elsewhere in the universe, unrelated to us by common descent? But since value realism can accommodate agent-relative reasons for action, the recognition of what is objectively valuable in the life of one creature does not automatically settle the question of what reasons it implies for the actions of another.

However, our direct access to value comes from human life, the life of one highly specific type of organism in the specific culture it has created. The human world, or any individual human life, is potentially, and often actually, the scene of incredible riches – beauty, love, pleasure, knowledge, and the sheer joy of existing and living in the world. It is also potentially, and often actually, the scene of horrible misery, but on both sides the value, however specific it may be to our form of life, seems inescapably real. Our susceptibility to many of these goods and evils plays a vital role in our survival and reproductive fitness – sexual pleasure, physical pain, the pangs and satisfactions of

strong unreflective tendency to ‘hold it against’ a male when he killed her offspring in such circumstances, on the contrary becoming receptive to his advances soon afterwards.” (Street, “A Darwinist Dilemma,” p. 120)

hunger and thirst – but they are also good and bad in themselves, and we are able to recognize and weigh these values. Initially we recognize them in our own lives, but it cannot stop there.

In looking for a historical explanation, a realist must suppose that the strongly motivating aspects of life and consciousness appear already freighted with value, even though they find their place in the world through their roles in the lives of the organisms that are their subjects. Sex, food, and drink are wonderful, in addition to being adaptive. Value enters the world with life, and the capacity to recognize and be influenced by value in its larger extension appears with higher forms of life. Therefore the historical explanation of life must include an explanation of value, just as it must include an explanation of consciousness.

If we recall the three potential types of historical explanation – reductive, teleological, and intentional – it is hard to see how a reductive explanation would be possible. Even if there were a partly reductive answer to the constitutive question about the existence of value – if the value of an experience of pleasure were constituted, for example, by the combined value of its protomental parts – that doesn't lead anywhere with regard to the historical question. Even with the help of some form of psychophysical monism, it seems impossible to imagine a reductive historical explanation of the origin of life, the development of conscious life, and the appearance of practical reason that would make it anything other than a complete accident that what we care about has objective value.

By contrast, once we recognize that an explanation of the appearance and development of life must at the same time be an explanation of the appearance and

development of value, a teleological explanation comes to seem more eligible. This would mean that what explains the appearance of life is precisely the fact that life is a necessary condition of value, and ultimately of its recognition.

I will again set aside the hypothesis of an intentional explanation. According to the hypothesis of natural teleology, the natural world has a tendency to give rise to beings of the kind that have a good – beings for which things can be good or bad. These are all the actual and possible forms of life. They have appeared through the historical process of evolution, but the explanation for the existence of that process and of the possibilities on which natural selection operates is in part that they bring value into the world, in a great variety of forms.

Since the emergence of value is the emergence of both good and evil, it is not a candidate for a purely benign teleological explanation: a tendency toward the good. In fact, no teleological principle tending toward the production of a single outcome seems suitable. Rather, it would have to be a tendency toward the proliferation of complex forms and the generation of multiple variations in the range of possible complex systems.¹¹

If we were not inclined to recognize objective reasons for action, and were motivated exclusively by our desires, we would have no reason to believe in the existence of value in a realist sense. There would be nothing to explain, beyond the system of subjective motives, and their capacity to be guided by the information delivered by perception, memory, and theoretical reason. But if we take our impressions of objective value to be substantially correct, rather than completely illusory, then we must regard the

¹¹ Such a possibility is described by C. D. Broad in *The Mind and Its Place in Nature* (London: Routledge, 1925), pp. 81-94.

appearance and evolution of life as something more than a history of the development of self-reproducing organisms, as it is in the Darwinian version.

We recognize that evolution has given rise to multiple organisms that *have* a good, so that things can go well or badly for them, and that in some of those organisms there has appeared the additional capacity to aim consciously at their own good, and ultimately at what is good in itself. From a realist perspective this cannot be merely an accidental side-effect of natural selection, and a teleological explanation satisfies this condition. On a teleological account, the existence of value is not an accident because that is part of the explanation of why there is such a thing as life, with all its possibilities of development and variation. In brief, value is not just an accidental side-effect of life; rather, there is life because life is a necessary condition of value.

This is a revision of the Darwinian picture rather than an outright denial of it. A teleological hypothesis must acknowledge that the details of that historical development are explained largely through natural selection among the available possibilities on the basis of reproductive fitness in changing environments. But even though natural selection determines the details of the forms of life and consciousness that exist, and the relations among them, the existence of the genetic material and the possible forms it makes available for selection have to be explained in some other way. The teleological hypothesis is that these things may be determined not merely by value-free chemistry and physics but also by something else, namely a cosmic predisposition to the formation of life, consciousness, and the value that is inseparable from them.

In the present intellectual climate such a possibility is unlikely to be taken seriously, but I would repeat my earlier observation that no credible account, even a

purely speculative one, seems to be available of how a system as staggeringly functionally complex and information-rich as a self-reproducing cell, controlled by DNA, RNA, or some predecessor, could have arisen by chemical evolution alone from a dead environment. This is not only the opinion of defenders of intelligent design. Although scientists continue to seek a purely chemical hypothesis about the origin of life, it is not unusual to find even card-carrying scientific naturalists like Francis Crick who say that it seems like a miracle.¹²

Some form of natural teleology, a type of explanation whose intelligibility I briefly defended in the last chapter, would be an alternative to a miracle – either in the sense of a wildly improbable fluke or in the sense of a divine intervention in the natural order. The tendency for life to form may be a basic feature of the natural order. And once there are beings who can respond to value, the rather different teleology of intentional action becomes part of the historical picture, resulting in the creation of new value. The universe has become not only conscious and aware of itself, but capable in some respects of choosing its path into the future – though all three, the consciousness, the knowledge, and the choice, are dispersed over a vast crowd of beings, acting both individually and collectively.

These teleological speculations are offered without conviction, merely as possibilities. What I am convinced of is the negative claim that, in order to take our questions and judgments about values and reasons seriously in a realistic sense, we must reject the idea that they result from the operation of faculties that have been formed from scratch by chance plus natural selection, or that are incidental side-effects of natural selection, or are products of genetic drift. When we ask ourselves, for example, whether

¹² Francis Crick, *Life Itself: Its Origin and Nature* (New York: Simon and Schuster, 1981), p. 88.

revenge is a true justification, or just a natural motive, or what kind of weight we should give to the interests of strangers, or of other species, we should think of ourselves as calling on a capacity of judgment that allows us to transcend the imperatives of biology.

I believe this is also true of our use of theoretical reason to determine the real character of the world presented to us by sense perception. (In this I part company with Street.) As I have said, the judgment that our senses are reliable because their reliability contributes to fitness is legitimate, but the judgment that our reason is reliable because its reliability contributes to fitness is incoherent. That judgment cannot itself depend on this kind of empirical confirmation without generating a regress: To make the judgment is necessarily to take it as having authority in its own right. I don't think any other mental stance is available in the theoretical case.

But in the case of value and practical reason, I believe it is coherent to be subjectivist – to regard all impressions of objective value or objective reasons for action as illusory, and to think of the processes of practical deliberation and moral reasoning as sophisticated ways of deciding what one really wants. The theoretical use of reason, standing on its own, provides a frame within which it is in principle possible to think of one's motivational system as a product of natural causes merely, and never as a response to what is really good and bad. Although I find it impossible to take up this position, I do not think it is unintelligible. The question is one of relative plausibility. And I have to acknowledge that for someone not disposed to accept value realism, the radical consequences I have drawn from it will only increase its implausibility.