

Changing Core Values

Catherine Z. Elgin

Old fogies tend to see declining standards everywhere. ‘Look at the clothes teenagers wear!’ ‘Listen to what passes for music!’ ‘Fruits used to be more flavorful. Now they taste like cardboard!’ Back when we were young, we insist, things were not just different but better. Standards change. Things are not what they once were. To some extent, no doubt this is true. But questions remain. How do standards and values change? How much do they change? Do the changes make things better or worse, or merely make them different? One common way to answer such questions relies on a distinction between intrinsic and extrinsic values.

An extrinsic value is one that promotes a valuable end. So an extrinsic value is conditional. It is a value only on the condition that the end it promotes is valuable. Fuel economy in a car is extrinsically valuable because gasoline is expensive, and fossil fuels are limited. That makes it good to conserve fuel. Were gasoline cheap and inexhaustible, fuel economy would not be a value. Fuel economy thus has value merely because something else is a value. Change in extrinsic values seems easily explained. If an end loses its value, the means that promote it also lose value. If an end increases in value, the means that promote it do too. If a better way of promoting or producing the end is found, the old means diminishes in value.

Even for extrinsic values, this analysis is too crude. For not every means of

promoting a valuable end is valuable. Saving five lives would be good, but not by killing a healthy person and transplanting his organs into five ailing patients. The end, although valuable, does not justify the means. This shows that the value of the means does not wholly derive from the value of the end it promotes. No values are entirely extrinsic. Means also have a measure of independent value or disvalue.

Nor can all values be extrinsic. An extrinsic value is a value because it promotes some other valuable end. But what makes that end valuable? Perhaps it promotes a further valuable end. For example, vitamin C is valuable because it prevents scurvy. Preventing scurvy is valuable because scurvy is a debilitating disease. Avoiding debilitating diseases is valuable because doing so promotes health. Health is valuable because it promotes human flourishing. The sequence can continue for a while, but not forever. Eventually reasons give out. When asked why human flourishing is valuable, we may have nothing more informative to say than 'It just is'. Recognition of this leads some philosophers to conclude that some values must be *intrinsic*. They are valuable in themselves and not for any further goods that they promote. We need not, and perhaps cannot intelligibly ask, 'What is human flourishing good for? What good does it promote?' Similarly for other intrinsic values. There is nothing more to be said than that they are valuable. If values are intrinsic, their value is basic and underived.

On this account, at the end of the sequence of conditional values lie values that are unconditional. They neither have nor need further justification. We know that their value is intrinsic because we recognize that they neither have nor need further justification. If values were permanently established, incontrovertible, and indisputable, this answer might be plausible. But even seemingly intrinsic values change. And we

need to account for their changing. If they really are basic, it is hard to see how they could change. One explanation is epistemological. Intrinsic values do not change, but our judgments about them can be mistaken. So when we revise our views about whether something is intrinsically valuable, we judge that our previous views were mistaken. We now believe that the values we had taken to be intrinsic are not intrinsic, and perhaps are not really values. This raises a problem: What reasons could we have to think that we were mistaken? If there is nothing by reference to which an intrinsic value is justified, what could supply evidence that we were wrong? There is no further court of appeal that could overturn our initial assessment. We might still change our minds. But we would have no reason to do so. There would be no basis for judging our later views better or worse than our earlier ones.

Richard Rorty denies that any values are truly intrinsic. But, like those who believe in intrinsic values, he thinks that some values function in such a way that they neither have nor need justification. They are, as Wittgenstein says, hinges on which our practices turn.¹ Rorty recognizes that they are not static. He says: ‘In the process of playing vocabularies and cultures off against one another, we produce new and better ways of talking and acting -- not better by reference to any previously known standard, but just better in the sense that they come to seem clearly better than their predecessors.’² Rorty can explain why such values change, but he maintains that aside from their seeming so, there is no standard for judging that the new ones are better or worse than the old. According to Rorty, seeming better, or at least seeming clearly better, suffices for being better. The difficulty, of course, is that every time styles change, the new fashions seem clearly better than their predecessors. We do not think that changes in style provide

a good model for understanding shifts in core values.

I believe that the difficulties arise because the intrinsic/extrinsic distinction relies on a linear model of justification. Either some further value must justify a given value, or nothing does. If the chain of conditional (extrinsic) values terminates, the endpoint seems arbitrary. p is valuable, and there is nothing to be said about why. If the chain does not terminate, the regress is vicious. p is valuable if q is; q is valuable if r is, and so on. But that one thing is valuable if another is does not establish that anything is valuable. For the 'if clauses' may fail to be satisfied. The conditions they specify may fail to be met. If values are arbitrary, so are changes in value.

When we focus on fashion, changes in value seem trivial and transient. Whether they are arbitrary makes no difference. The pendulum swings back and forth. For a time miniskirts and bell bottom trousers are in style. Then they are out of style. Then they are back in. When they are in style, they seem better, when they are out, they seem worse. Although there may be some sociological interest in the fashion pendulum's arc, it is not on the whole very important.

Not all changes in value are so superficial. The change from feudalism to a more democratic system of political rights and responsibilities does not seem to be just a change of style. Core values change. And some of the changes seem to be definite improvements. If we are to understand ourselves, it is important to know the causes, consequences, and normative status of such changes. The linear model is unhelpful, I believe, because core values are deeply and intricately interwoven into institutions, practices, and ways of life. Little of their value derives from their being either conditionally or unconditionally good in themselves. Their value largely derives from

their roles in valuable practices, institutions or ways of life. This does not make them extrinsic. For they do not merely promote valuable practices, institutions, or ways of life. They are integral parts of them. That is, they are part of what makes such practices, institutions, or ways of life valuable. To understand the values and reasons they change then, we need to assess the larger wholes they figure in.

Core values are not necessarily the values that we most loudly champion. Indeed, they may be nearly invisible, being so engrained in the fabric of our lives that we take them for granted and do not recognize them as separate strands. For example, a tribe might be committed to treading lightly on the Earth, to living at one with nature rather than to setting itself up in opposition to nature. This commitment might be integral to the way they farm, the way they cook, the way they build, the way they bury their dead. They might live this commitment, recognize and acknowledge the import of its several manifestations, but never acknowledge the core value that links the various manifestations together. Still, a change in that value would reverberate throughout the entire culture.

Moreover, the values we loudly endorse and publicly appeal to may be relatively peripheral. We may claim to be committed to democracy. But if we regularly support undemocratic regimes and undermine unfriendly democratic ones, our commitment to democracy does not run very deep. We may claim to value education. But if we do not adequately fund education, properly respect teachers, reward educational achievements, and so forth, our actions belie our words. Whether or not they are acknowledged, core values, as I understand them, are the values that deeply and extensively inform our lives.

Core values differ in scope. Broad-scope values lie at the heart of traditions,

cultures, or ways of life. They are apt to extend across generations. Narrow-scope values are central to more restrictive practices. We speak of the core values of education, medicine or the law. These are the deepest values of those professions. Respecting them is an integral aspect of good professional conduct. But they may fail to be core values, or indeed be values at all, outside of the professions that endorse them. At the other extreme are more global values, which transcend the limits of particular cultures. Any decent society must make adequate provisions for the care of its young. So an obligation to care for children is probably a universal core value. How it is realized varies from one society to the next. But the basic value seems to be, and should be shared. A value's status as a core value thus stems not from the intensity with which it is held, the number of people who hold it, or its obviousness, but from its centrality to a complex, norm-laden social structure.

Because they are so central and so intricately interwoven into our practices, core values are deeply implicated in our actions and assessments. We characterize actions in terms of them – as, for example, honorable or wasteful, and we judge the acceptability of things in terms of them. But if core values are the touchstones by reference to which we judge other things, how can we assess them? To justify a value is to account for its being a value. If we insist that core values must be justified in isolation, our prospects look bleak, as the discussion of intrinsic and extrinsic values showed. But because they are entwined in our practices, institutions and ways of life, we have additional resources. At least part of the justification for a value lies in the justification for the practice, institution, or way of life it belongs to.

An example may be helpful here. Consider the question whether a physician

ought to prescribe an antibiotic to treat a child's ear infection.³ There are several reasonable perspectives on the question. A perspective that focuses on the good of the individual patient highlights the benefits and costs (if any) to the child. The drug would be effective. It would cure the infection, spare the child several days of pain, and prevent scarring of the eardrum. It is not wildly expensive and would not do her any lasting harm. Considering only the good of the individual patient, the physician ought to prescribe the drug. A public health perspective highlights different factors. Use of antibiotics leads to the evolution of antibiotic resistant bacteria which pose a long term public health risk. Future generations will suffer because of over use of the drugs by the current generation. For the greatest good of the greatest number, antibiotics should be used sparingly, so that they continue to be effective against very serious diseases. Children's ear infections are not serious enough to merit the drug. So it ought not, according to a public health perspective, be prescribed. An economic perspective highlights the price of the drug, and the way that medical costs have increased dramatically, largely because of the price of medications. Again prescribing the medication seems unwarranted. None of these perspectives is irrelevant or unreasonable. But they pull in different directions. Nonetheless, as things stand, at least in the west, the physician's duty seems clear. A core value of western medicine is that the physician must do what is best for his patient. In light of that value, the physician should set aside the public health and economic considerations and look to the good of the individual patient. If the medication is good for her, he should prescribe it. This is the answer, so long as the current core values of western medicine are accepted.

But the tensions among these perspectives may give us reason to question the

adequacy of the core values of western medicine. This is not done by juxtaposing the separate core values and attempting to decide which of them is best. For to determine the value of a value – the value of doing what is best for the individual patient or the value of doing what is best for everyone even remotely affected by an action – we need to know how these values function. Rather than attempting to evaluate their merits in isolation, we should consider how they figure in and what they contribute to the practices, institutions or ways of life they belong to. We need to assess then not just the particular value, but the practice of medicine at whose core it lies. Given what we want medicine to do – heal the sick and the injured, alleviate pain and suffering, bring about health and well being -- is the current system optimal? We need to ask not, within western medicine as it is now configured, whether the physician should put the interests of the individual patient first, but whether western medicine should be configured as it is. What values does it realize, what values does it sacrifice? Is there a better way of organizing our practices and institutions?

If we restrict our attention to what is required within medicine, these questions cannot be raised. Within western medicine, the physician's duty is clear. But the questions are reasonable, legitimate, and perhaps even pressing if we step outside. We need to ask what goods medicine seeks to achieve, how well it achieves them, at what cost, and who should pay the cost. The vindication of a core value then rests on the vindication of the practice that the core value figures in. This may involve asking what needs to be sacrificed to achieve the goods that medicine seeks to achieve, and whether a preferable overall bundle of goods could be achieved by trading off medical goods for goods of other kinds. To judge the continuing value of core values, we need to assess not

just the values themselves, but also the larger wholes they figure in, the goods that those wholes provide, and the costs and benefits of realizing those goals, as opposed to others that we might try to achieve instead.

Answers to such questions are not fixed and final. Circumstances change. Something may be a worthwhile core value in one context, but not in another. Another example shows this. Until the 1960s the Hmong people of northern Laos did not have an alphabet. Because they could keep no written records, a good memory was a tremendous asset. Having a good memory was, and should have been, a core cognitive value for the Hmong. But with the introduction of an alphabet, this changed. Once they acquired the capacity to record information in writing, a highly developed memory became less important. There is less need to remember exactly occurred, when a written record of the events is available. So what was once a valuable core cognitive value – having a good memory – has rightly evolved into the less central value that it is in other cultures that have written languages. The Hmong were not mistaken in holding a good memory as a core value when they had no alphabet. It genuinely was an asset, not just to the individual but to the community that had to rely on the memories of its members to retain access to information about things distant in time and space. But when circumstances changed, a good memory became less valuable. For a new and more reliable mechanism for preserving information replaced it. The value of the value changed because of changes in surrounding circumstances.

If this picture is basically right, then the justification of core values largely derives from the justification of the larger structures they belong to. This raises two related questions. What justifies the larger structure? And how must a factor figure in the larger

structure to be justified when it is? John Rawls's discussion of justification *of* a practice and justification *in* a practice helps here.⁴ A practice, in the sense that interests us here, is a network of mutually supportive factors. Each separate factor is more reasonable in light of the others than it would be alone. So a value is justified in a practice if it is more reasonable in light of the other elements of the practice than it would be without them. It is then justified *in* the practice it belongs to. It is not, in Wittgenstein's terms, an idle wheel, unconnected to the rest of the mechanism. This alone is not enough to wholly justify it, however, for the practice itself may be unjustified. That torture of political enemies is justified in a despotic regime does not make such torture justified, because the despotic regime is unjustified. So for a consideration to be justified, the system of factual and evaluative commitments it belongs to needs to be justified as well.

The justification of such a system derives from its relation to our antecedent commitments. If the system is at least as reasonable as any alternative in light of what we were antecedently committed to, it is justified.⁵ These antecedent commitments concern substantive beliefs, goals, values, and methods. Given what we are trying to do, what we believe about the domain and about our resources and options, the methods at our disposal, the constraints on our resources, and the values we want to uphold, is a system that incorporates this value as good as any available alternative? If so, it is justified. If not, it is not.

Justification, on this model, is a matter of reflective equilibrium.⁶ The considerations within the system are in equilibrium, being mutually supportive. The system as a whole is one we can on reflection accept, being reasonable in light of what we were already inclined to endorse.

The reflective equilibrium standard does not require that each individual component of the system answer to or be justified by an antecedent commitment. Our antecedent commitments are apt to be incomplete. So we integrate into our developing system considerations about which we had no specific prior commitments, because the resulting structure is stronger for them. Even if for example, physics has no direct evidence of the existence of positrons, it might incorporate a commitment to positrons because of considerations of symmetry. There is plenty of evidence for the existence of electrons, and good reason to endorse symmetry principles. And symmetry demands that negatively charged particles be balanced by positively charged counterparts. So even with no direct evidence of positrons, physics has good reason to incorporate them into a tenable physical theory. We may also repudiate some antecedent commitments. Perhaps we find that what we thought were general principles only work for a restricted range of cases. For example, we may discover that a grammatical principle that we thought was universal applies only to subject-verb-object languages. Then the system that emerges, rather than incorporating the principle we once held, should recognize that it only works within a limited sphere. Perhaps we recognize that a principle we once held was due to prejudice. Now that we understand the principle and the reasons we held it, we see that it is unsound. In all these cases test the developing system by reference to our prior commitments. But we need not incorporate all of our prior commitments into the new theory. For the most part, however, the new account should explain why the prior commitments seemed reasonable, when they did. Sometimes the explanation is that they were in fact reasonable. In other cases it is that they were understandable mistakes.

Equilibrium is a balance, and the balance can be upset. So maintaining reflective

equilibrium is a continuous process. A barrage of factors from evolving bacteria to developing alphabets, to technological advances, to innovative methods and new ideas can upset the balance. What was reasonable in one set of circumstances may become unreasonable as conditions change. Another medical example illustrates this point. In the middle of the 20th century, medicine developed the capacity to transplant kidneys. Often the kidneys are taken from cadavers, but live donors can be used. (And there is evidence that kidney donations from live donors are better for the recipient.) There is no denying that removing a healthy kidney from a healthy person is doing him medical harm. The surgery itself harms the organism. The loss of a kidney puts him at risk. And so on. Still, the kidney donation would save someone's life, the donor is willing, and the risks are minimal. So, one wants to say, on balance it is medically appropriate for such donations to take place. But removing a healthy kidney from a healthy person violates a core value of western medicine that goes back to Hippocrates: First do no harm. Advances in medicine then seem to put pressure on the core value.

When a consideration undermines reflective equilibrium, a variety of adjustments might be made to restore balance. We could hold fast to the core value and rule that organ transplants from live donors are impermissible because they violate the harm principle. Or we could simply introduce an exception for organ transplants. Or we could revise our conception of harm. We might decide that even if the donor's body is harmed, *he* is not truly harmed because he voluntarily donates his organ. His life, we might argue, is improved by being permitted to help save the life of someone he loves. Or we could revise the core value and replace it with something like: do no harm unless the benefits far outweigh the costs and the person harmed is willing to be harmed. No doubt a variety

of other revisions might also be made. How should we choose among them?

There is no rule for deciding such matters. In general we seek the best system overall. The choice among revisions is guided by our assessments about which systems will best match our antecedent convictions. If we think that donating a kidney to a dying brother should be permissible, we will not opt for the outright rejection of the procedure. If we think that introducing an exception for organ transplants is ad hoc, we will be reluctant to choose that option. We might be tempted by a revision of our conception of harm if we think that voluntarily doing something one considers worth doing is not harming oneself.

Rather than attempting to decide just how this particular case should be handled, I want to highlight the fact that the apparent value of such organ donations puts pressure on the system that would seem to forbid them, and that holistic considerations influence what revisions we are willing to accept. W. V. Quine's principle of minimal mutilation is at work.⁷ In revising a system of commitments, we should preserve as much as possible of the system we previously endorsed. We would not accept a revision that permitted killing one to save five, even though that too would be a way of accommodating the good of organ donations from live donor.

The principle of minimal mutilation may seem unwarranted. If any of a variety of revisions would restore equilibrium, why should we favor the least drastic? One reason is practical. Change in received opinions and practices increases the likelihood of mistakes. The more similar the new system is to old one, the fewer mistakes we are apt to make. Moreover, revisions often have unforeseen and unintended consequences. Some of these may be deleterious. If we replaced the 'do no harm' principle with one that said

‘do no harm unless the good more than compensates’, we would find ourselves unable to object to killing one to save five. A third reason is that if the previous system seemed acceptable until the new consideration arose, that constitutes reason to think that it is on the whole fairly good. It is likely that a modest revision can preserve most of what is good about it.

The principle of minimal mutilation seems to insulate core values against revision. If we can always make peripheral revisions, and we should always make the least disruptive revisions possible, then it might seem, the core of any system is untouchable. Some more modest revision, one is apt to think, will always be possible. This is not quite true. To a certain extent core values are insulated from revision. They are not easily disrupted. But they are not utterly immune to revision. For it is not always less disruptive to change the core than to make multiple changes in the periphery. For example, Kepler’s laws of planetary motion constituted a drastic revision in astronomy, but turned out to be less disruptive than the multiple, increasingly implausible, ad hoc revisions needed to keep the Ptolemaic model in accord with observations. Revisions within the Ptolemaic system were probably initially more tenable the more drastic changes advocated by Copernicus and Kepler, but eventually it became clear that revisions of core astronomical commitments yielded a better overall account.

Nor is it the case that we can always restore equilibrium while preserving the current core. The principle that every event has a determinate cause has long been central to physics, as well as to common sense. But there seems to be no way for quantum mechanics to preserve it. Events at the quantum level are evidently irreducibly probabilistic. We could of course retain our commitment to physical determinism by

rejecting quantum mechanics. But that would require repudiating the scientific evidence that supports quantum mechanics. And that would require repudiating the methods that generated that evidence and the standards that supported it. So the mutilation of our understanding of science that would be required to retain our commitment to determinism would be far from minimal.

I said that reflective equilibrium is assessed against the antecedent commitments we hold. An acceptable system must be reasonable in light of them. So the question arises: who are *we*? That is, whose commitments form the backdrop against which a system is assessed? The answer varies. At the outset, everyone's views get a hearing. But as we learn more about a subject, we refine our views about what it takes to understand that subject, and about who has the requisite expertise. Given that this is what physics is, that is what is required to understand physics, and those are the people who satisfy the requirement. Thus their relevant commitments form the backdrop against which changes in physical theory ought to be assessed. The opinions of non-experts carry far less weight. In other areas, expertise is not so much a matter of academic training and credentials. It is a matter of moral sensitivity, social awareness, and wide experience with the community. This is the sort of expertise that the elders of a tribe are expected to possess, and the reason why their views carry weight. Given their wisdom and experience, their views are a reasonable backdrop against which to assess revisions in core social values. In general, we have opinions about whose opinions should be listened to on a given subject. And we assess revisions in light of the relevant expert opinions. Since these opinions are only the backdrop, it is possible for the revised theory to diverge from what the experts currently believe. To say that the opinions of the physics

professors or the tribal elders deserve consideration is not to say that they will prevail. As scientific and social revolutions demonstrate, a radical revision may prove best on balance. Then we revise and update our list of experts on the basis of how well their opinions are borne out in the revisions we find acceptable.

The reflective equilibrium model then shows that although core values are resistant to change, they are capable of changing. Moreover, it explains what makes such changes improvements. The system that results from any revision must consist of mutually supportive elements. And, to be an improvement, it must be more reasonable than any available alternative in light of our antecedent commitments. The standard of reasonableness itself is among the antecedent commitments we appeal to. If a revision in core values belongs to a system that satisfies this standard it is an improvement over its predecessors.

I have been talking quite generally about changing core values. I want to end by considering briefly how my account of such change bears on education. Here too we need to distinguish between good in a system and good of a system. What is good in a given educational system is easily recognized. We have curricula, rubrics, benchmarks, and so on. We say that the student is expected to have mastered *a*, *b*, and *c*, and provide exams or other assessment measures to determine how well he has done. It is not hard to identify the students who do well in the system. But it is much harder to determine whether someone who has done well in an educational system is well educated. We have to ask how good the educational system is. This requires determining what the goal of an educational system is, what it seeks to do. To answer this requires figuring out how education contributes to or figures in a good life. So we cannot tell whether someone is

well educated without at least implicitly answering Aristotle's question: What is the good for man?

The answer may vary with circumstances. When lives were shorter and children were expected to follow in their parents' footsteps, they needed to acquire the capacity to play an antecedently defined social role. Now lives are long, opportunities abound, and people are apt to have multiple careers in a lifetime, and to live long after retirement. So versatility, adaptability and a capacity to learn how to learn are more central values. Critical reasoning and reflection are too. As alternatives become available, we need the resources to assess them critically and responsibly. With changes in circumstances, core values and other central commitments can be unsettled. With the introduction of an alphabet, the Hmong had to revise their education, and their views about what was educationally important. Students had to learn to read and write, to value reading and writing, and to set less store on memorizing. With advances in technology, we may have to do something similar. At least we should seriously investigate whether, for example, the sorts of mathematical skills or the sorts of research skills that are traditionally taught are the ones our children now need. We should ask about the roles of art and literature and music as well. Do we provide children with the aesthetic education they need in order for their lives to be enriched by the arts?

Human beings need to learn a lot more than other animals do. So education lies at the heart of a good human life. A good system of educational commitments must be in reflective equilibrium, and must be part of the wider reflective equilibrium that constitutes human flourishing. Such an equilibrium is dynamic. It involves continuous adjustments to changing circumstances. We should not expect to find a fixed and final

network of educational values. Rather we should hope to discover ways to monitor and assess our current commitments with an eye to improvement. Core values, in education as elsewhere, have a considerable inertia. They are more resistant to change than more peripheral values. So they are not, and should not be, given up lightly. But if they cease to be effective, they are subject to revision. If others are more effective, the received values are subject to rejection. A good education should equip people with the resources to recognize the need for revision and with the capacity to make revisions restore reflective equilibrium in their practices, institutions and ways of life.

¹ Ludwig Wittgenstein, *On Certainty*, New York: Harper Torchbooks, 1969, #341.

² Richard Rorty, 'Introduction', *Consequences of Pragmatism*, Minneapolis: University of Minnesota Press, 1982, p. xxxvii.

³ I am indebted to Sam Elgin for useful discussions about this example.

⁴ John Rawls, 'Two Concepts of Rules,' *Philosophical Review*, 64 (1955), pp. 3-32.

⁵ See Catherine Z. Elgin, *Considered Judgment*, Princeton: Princeton University Press, 1996.

⁶Cf., John Rawls, *A Theory of Justice*, Cambridge, MA: Harvard University Press, 1971; Nelson Goodman, *Fact, Fiction, and Forecast*, Cambridge, MA: Harvard University Press, 1984, p. 64.

⁷ W.V. Quine, *The Philosophy of Logic*, Englewood Cliffs, NJ., Prentice Hall, 1970, p. 7.