Article

Theory and Research in Education
9(3) 251–263

DOI: 10.1177/1477878511419559
tre.sagepub.com

Science, ethics and education

Catherine Elgin
Harvard University, Cambridge, USA

Abstract

An overarching epistemological goal of science is to develop a comprehensive, systematic, empirically grounded understanding of nature. Two obstacles stand in the way: (1) Nature is enormously complicated. (2) Findings are fallible: no matter how well established a conclusion is, it still might be wrong. To pursue this goal in light of the obstacles, science incorporates ethical values. These values are not mere means; their realization is integral to the sort of understanding that science embodies. The recognition of these values should be incorporated into science education.

Keywords

ethics, integrity, misconduct, science

Section 7009 of the America COMPETES Act (2009) mandates that all National Science Foundation (NSF) funded undergraduates, graduate students and postdoctoral students be given 'appropriate training and oversight in the responsible and ethical conduct of research' (National Science Foundation, 2009). No one denies that scientists should conduct their research ethically; and few would deny that neophyte scientists need to learn how to do so. Nevertheless, many principal investigators apparently consider the requirement an unfunded mandate – another demand on their already strained budgets, time and attention. Evidently they think that teaching their protégés to be morally ethically responsible scientists is something different from teaching them to be scientists. Consequently, generic handbooks and websites have been developed to instruct students how to recognize, avoid, and protect themselves from the more obvious forms of scientific misconduct. This, no doubt, is helpful. But such quick fixes treat ethical requirements as addenda rather than as integral to scientific inquiry. That is a mistake. Scientific inquiry has, I will argue, an irreducibly ethical dimension. Ignoring it skews our understanding of the scientific enterprise and its products. If I am right, then although they may not realize it, senior scientists have been imbuing their juniors with the relevant ethical values already. The point of the NSF mandate is not to ask them to take on a new task for which they are,

Corresponding author:
Catherine Elgin, Graduate School of Education, Harvard University, Cambridge, Massachusetts 02138, USA
Email: catherine_elgin@gse.harvard.edu

or may think they are, unqualified, but to do explicitly and self-consciously what they have been doing tacitly and unconsciously all along.

An overarching epistemological goal of science is to develop a comprehensive, systematic, empirically grounded understanding of nature. Two obstacles stand in the way: (1) Nature is enormously complicated. (2) Findings are fallible: no matter how well established a conclusion is, it still might be wrong. I will argue that in order to pursue the goal in light of these obstacles, science incorporates ethical values. These values, I will urge, are not mere means; their realization is integral to the sort of understanding that science embodies. I will then argue that the recognition of these values should be incorporated into science education. I do not contend that the factors I discuss exhaust the ethical obligations of scientists. Arguably scientists have an obligation to promote, or at least refrain from jeopardizing, the general welfare. My discussion has nothing to say about this. My focus is on ethical values that underwrite modern science's claim to constitute an understanding of nature.

Very roughly, the argument is this: Science requires collaboration. That collaboration requires trust. Since trust is unreasonable in the absence of trustworthiness, scientists need to be, and to consider one another, trustworthy. Since trustworthiness is a moral characteristic, morality figures in science. That being so, science education should inculcate trustworthiness and an appreciation of the ways it contributes to and figures in science.

Epistemic interdependence

Modern science accommodates itself to the complexity of nature through a division of cognitive labor. Instead of starting from scratch, scientists build on previously established findings, use methods others have devised and tested, employ instruments others have calibrated, and analyze data using mathematical and statistical techniques others have validated.[1] Research teams rather than solitary investigators frequently generate results. Such teams can do more than any one of their members can do alone, not only because teams can cover more ground, but also because team members typically have different cognitive strengths and often have different areas of expertise. Collectively a team brings to bear a broader range of talents and training than any individual scientist possesses. The sort of understanding modern science supplies is grounded in epistemic interdependence. Science as we know it is a cooperative endeavor.

Because scientific collaboration depends on a division of *cognitive* labor, it occurs in a particularly harsh epistemic climate. Collaborators need not know each other well or understand each other's contributions. The paper reporting the sequencing of the human genome had over 250 authors, working in institutions on four continents (Venter et al., 2001). It is unlikely that the authors were in a position to vouch for one another's intellectual character or monitor one another's work. Even those in the same venue were probably at best dimly aware of the contributions of some of their co-authors. Although most scientific papers have fewer authors, such lack of physical proximity and close personal ties among members of a research team is common. Moreover, all scientists build on results of, and rely on methods developed by, unknown predecessors. And, even when they work in the same laboratory, the intellectual distance between collaborators who differ in expertise can be vast. A scientist may be ignorant of the reasons behind a conclusion that a collaborator draws; and if she were told, she might not fully appreciate why they are reasons (Hardwig, 1991).

Such epistemic interdependence yields significant benefits. It extends epistemic reach: more data can be gathered, more experiments run, more matters investigated, more factors considered, more perspectives accommodated than would be possible for a single researcher working alone. It stabilizes findings: individuals are subject to bias, inattention, carelessness, and wishful thinking, as well as to limitations in background knowledge and scientific imagination. If the result of an investigation is the joint product of several minds, foibles can cancel out. What one scientist overlooks, another may notice; what one slights, another may focus on; what one never learned, another may have studied in depth. It solidifies results: because of team members' differences in expertise, the quality of evidence and range of interpretations available through teamwork is apt to be better than the quality and range available to an individual (Hardwig, 1991). Collectively, a research team can advance understanding more than any member could do alone.

But epistemic interdependence creates vulnerability. To accept a conclusion is to be prepared to use it as a basis for inference and action (Cohen, 1992: 4), and to be prepared to invite others to accept it. To the extent that one scientist does not know or understand what another did, she is in danger of accepting and basing her investigations on unwarranted findings. Should that occur, the course of her inquiry may be deflected along fruitless channels, and promising lines of investigation may be overlooked. If she co-authors a paper contaminated by a collaborator's scientific misconduct, her professional reputation and livelihood are at risk.

Epistemic interdependence thus requires trust. Scientists can reasonably depend on one another's contributions only if they can be confident that those contributions are scientifically acceptable. Perhaps personal acquaintance underwrites trust among a few investigators who work closely together. But, since the requisite epistemic interdependence extends across the scientific community, personal ties are inadequate. Trust in one's fellow scientists must be impersonal. A major function of the infrastructure of science is to secure the bases of impersonal trust.

Although the standards apply throughout scientific practice, their general, impersonal character comes to the fore with publication. In publishing her research, a scientist issues an open invitation to the scientific community to accept her results. She gives its members her assurance that they can count on her; anyone is welcome to depend on her findings and, she intimates, can do so with confidence. She invites their trust.

In so doing, she undertakes an obligation to be worthy of their trust. A critical question is what standards scientific work has to meet to vindicate the open invitation. Ideally, perhaps, we would like to demand that a research report's content be true and warranted by the research it describes. This is too much to ask. Science is fallible; it has no way to guarantee that findings are true. All that a scientist can plausibly claim is that her research satisfies current standards. To be sure, scientists make categorical claims. But those claims are surrounded by tacit caveats. Because everyone knows, and everyone knows that everyone knows, how fallible the enterprise is, that fallibility is taken for granted (Elgin, 2001). Because those standards embody the best way the discipline has devised to generate and validate understanding in its field, results that satisfy those standards are currently acceptable.

The scientific community sets standards whose satisfaction not only warrants conclusions, but shows them to be warranted. First-order requirements, such as those concerning the size and constitution of the evidence class, the standard of statistical significance, the proper sort of analysis, bear directly on warrant. Second-order requirements, by making it manifest that results hold regardless of who generated them or why, pertain to the assurance of warrant. They mandate that findings be intersubjectively observable; that experiments be reproducible; that data be unambiguous; that research reports be peer reviewed. If a phenomenon is intersubjectively observable, it makes no difference who in particular observed it; any competent scientist could have done so.[2] If an experiment is reproducible, the acceptability of its result does not turn on the character or intellectual endowments of the scientist who performed it; had any competent scientist done it, the result would have been the same. If data are unambiguous, the identity of the interpreter is immaterial; all competent scientists would interpret the data in the same way. When it makes no difference who generated a result, when any competent scientist could have done so, the personal foibles or intellectual idiosyncrasies of the investigator are irrelevant to its acceptability. This indifference makes science impersonal.

Scientific misconduct

An investigator commits scientific misconduct when she knowingly misrepresents herself as having satisfied the standards of her science. The most flagrant forms of misconduct are fabrication of data, falsification of findings, and plagiarism. Fabrication and falsification misrepresent what was done; plagiarism misrepresents by whom it was done. The obvious moral objection to all such conduct is that it is dishonest. But what, if anything, does this have to do with its being bad science?

To fabricate data is to present as evidence cases that were not observed. To falsify a finding is to report that an investigation had a result that it did not actually have. Both fabrication and falsification are lies. Thus, one might think that they subvert the goal by introducing falsehoods into science. That would obviously be bad science. This explanation assumes that a contention grounded in falsification or fabrication is false. That need not be so.

Scientists who falsify or fabricate rarely seek to introduce falsehoods into science. Because science is self-correcting, they could not reasonably expect to succeed, at least not in the long run. Rather, they seek to incorporate into science what they take to be a truth, but one for which they do not have, or do not yet have, evidence. This makes the problem trickier. If the conclusion they purport to have established is true, one might argue, our understanding of the phenomena is enriched rather than degraded by its acceptance. And if the perpetrator has good instincts about where the truth lies, his falsified or fabricated conclusion might well be true. If it is, what is the problem?

If truth is the end, and scientific research a mere means to it, falsifiers and fabricators with good instincts or good luck simply take shortcuts. They achieve the end without deploying the standard means. They skip the intermediate steps of acquiring and evaluating empirical evidence. If all that matters is truth, how a truth was arrived at should make no difference. But mere truth is not what science is after. A science is not just a compilation of discrete bits of accurate information about a topic; it is a comprehensive, systematic understanding of a range of phenomena, based on carefully acquired, scrupulously recorded, empirical evidence. The tie to evidence is crucial, for evidence underwrites epistemic standing. Conclusions with evidential support are conclusions scientists have reason to accept; conclusions – even true conclusions – without such support are conclusions they lack reason to accept.

This explains what makes the falsified ‘data’ or fabricated ‘finding’ scientifically objectionable. Even if the alleged finding is true, in the absence of adequate evidence one would be unwise to depend on it. To accept a conclusion without adequate backing is to take an epistemic risk. To put other members of the community in a position where they are unknowingly taking such a risk is to fail in one’s obligations to them. That is what proffering a falsified or fabricated finding does. It constitutes a betrayal of science because scientists have undertaken a commitment to proffer results that other members of the scientific community can, by their own collective lights, trust.

Even the best scientific research affords no guarantees. However scrupulously performed an experiment is, its result may be false. There is always some risk in accepting any finding. The standards of a scientific community constitute its consensus about what risks are worth taking. But whatever the standards, the goal of understanding nature would be subverted if the scientific community was permanently saddled with the errors of its forebears. Science therefore incorporates methods for rooting out received errors. Results that cannot be reproduced or that do not stand up under further testing and elaboration must be excised from science. But it is not enough to simply declare a previously accepted finding to be false or unwarranted. To learn as much as possible from their mistakes, scientists want to know what went wrong. Rather than simply deleting an erroneous conclusion, it is desirable to backtrack to the source of the error. This is central to the objectionability of plagiarism.

The content of a plagiarized report may be, and I suppose often is, beyond reproach. But it is not the product of the scientist who purports to have produced it. Again this is clearly dishonest. But what makes it bad science? The acceptability of a finding is independent of the identity of the scientist who generated it. Indeed, given the impersonality described earlier, if the report is beyond reproach, any competent scientist could have done the research and written it. This almost makes it look as though science should have no objection to plagiarism. Why should scientists care who did the research, so long as it satisfies their standards?

Giving credit where credit is due encourages good work. So an obvious objection to plagiarism is that failing to credit productive scientists while crediting unproductive ones threatens to drive the wrong scientists from the field and to direct research funding to lazy, inept or unproductive investigators. Plagiarism then is objectionable because it interferes with an effective incentive system by making it difficult to identify and reward those who are doing good work.

Is there any epistemological reason to credit the individuals who actually did the research and to object when others try to steal their credit? That is, does the misappropriation of other people’s work directly damage the pursuit of understanding? An affirmative answer stems from fallibilism. No conclusion is certain. However scrupulous scientists are, however well-tested their methods, and however well-honed their standards, conclusions still might be, and sometimes are, wrong. That being so, science incorporates mechanisms for detecting and correcting errors. The overwhelming majority of such errors are honest mistakes. When a mistake is discovered, one might think, the remedy is simply to excise it. Withdraw the paper and admit that there is currently no good reason to believe its conclusion. But if the research was responsibly done and satisfies current scientific standards, a question remains: What went wrong? Rather than simply deleting the error, it is desirable to track down its source. Was it due to sloppiness, bad luck, a mistaken assumption, or an unreliable method? Are the science's standards too lax or their applications uneven? To maximally profit from mistakes that have previously passed muster, it is advisable to re-examine the raw data, the interpretations, and the analyses to find out where the misstep occurred. The aim is not only to excise a particular error but correct the flawed assumptions, methods, practices, or standards that led to it. Scientists do not want to keep making the same sort of mistake.

Plagiarism stymies backtracking. The scientist who purports to have done the research did not do it. So he cannot provide information about what was done and what went wrong. Fallibilism underwrites a demand for corrigibility. Plagiarism unduly limits corrigibility; it undermines the self-correcting character of science.

The dishonesty of falsification, fabrication, and plagiarism is not an accidental feature of practices whose scientific objectionability lies elsewhere. Their dishonesty thwarts science's goal. The practices in question are bad science because they are dishonest. An obligation to be truthful then lies at the heart of science.

Epistemic responsibility

Scientific misconduct frustrates the quest for understanding because it undermines collaboration or corrigibility. This suggests that members of a scientific community bear additional obligations to one another. For outright dishonesty is not the only threat to collaboration and corrigibility.

Science is subject to epistemic requirements that constrain experimental design, data gathering, analysis, and assessment. Hypotheses should be formulated sharply enough to admit of empirical testing. Experiments should be designed to yield data capable of confirming or disconfirming the hypothesis under investigation. Irrelevant and obviously misleading evidence should be discredited. Data should be painstakingly gathered and carefully recorded in a timely fashion. They should be displayed in such a way that their significance and their limitations are manifest. Analyses should be rigorous, relevant, complete, and apt. Assessments should be judicious. Research reports should be grounded in a comprehensive, fair-minded, current understanding of the topic and the methods for investigating it. They should acknowledge the assumptions made, the limitations of the approach being used and the potential constraints on validity. These requirements are familiar; their satisfaction yields conclusions there is reason to accept.

There is nothing overtly ethical about these requirements or the benefits of satisfying them. They seem to be purely epistemological requirements. The ethical element emerges from the open invitation implicit in publication. Results that justifiably inspire confidence must be epistemically responsible. Not only intellectual honesty, but also competence and conscientiousness are mandatory. In publishing a result, a scientist implicates not only that it was honestly arrived at, but also that it satisfies the epistemic standards of the profession. Because scientists depend on one another, the epistemic requirements are ethical requirements; if scientists fail to satisfy the epistemic requirements, they fail in their ethical obligations to the community.

Carelessness, fickleness, obtuseness, and bias are not just epistemological failings; in the context of science, they are also ethical failings. Results infected by such faults are honest mistakes. Confirmation bias, for example, leads a scientist to overlook evidence that tells against a hypothesis she favors. Someone subject to confirmation bias is not being dishonest; she genuinely believes that her evidence supports her conclusion. But inasmuch as she can, by being more rigorous, avoid confirmation bias, her mistake is a culpable error. The same holds for carelessly taking or recording data, overstating one's results or understating their limitations, and so on. By inviting her colleagues to accept a conclusion that by their own lights they ought not accept, a scientist does them wrong.

In agreeing to work together, people engender obligations in one another – obligations whose content is determined by the objective they jointly pursue. Ceteris paribus, they owe it to one another to discharge those obligations. In science, the goal is understanding. So the obligations are in the first instance epistemic; they pertain to generating and sustaining the understanding that science seeks. But because scientific inquiry is a joint epistemic venture, the epistemic obligations are also ethical obligations. Scientists owe it to one another to satisfy the epistemic requirements, and to make it manifest that they do so.

This pattern is not peculiar to science. Members of other collectives, such as athletic teams and orchestras, are under similar obligations to one another. In the first instance, the obligations of a quarterback are athletic; those of a bassoonist aesthetic. But the quarterback owes it to his teammates to discharge his athletic obligations well; he is something worse than a poor athlete if he carelessly or negligently lets down his side. And the bassoonist owes it to his fellow musicians to perform well; he is something worse than a poor musician if he neglects to learn his part or is cavalier about playing in tune. The commitments that team members make to one another are akin to implicit promises; because the common knowledge of those commitments creates a basis for legitimate expectations, one wrongs one's fellows if one does not abide by them.

Arguably, in science there is a difference in scope. The members of a research team, like the members of a football team or an orchestra, consist of a limited number of identifiable individuals who have well-defined obligations to other people they may know reasonably well and work reasonably closely with. But, as we have seen, in publishing findings, a scientist effectively invites the entire scientific community to join her team; since she provides assurance that her findings satisfy their epistemic standards, she is under an obligation to proffer findings that they all can be confident of. Indeed, to the extent that the rest of the population depends on the findings of science, in publishing her results, a scientist offers her assurance that everyone can have confidence in them.

Incentives

Because it is common knowledge that only work that at least ostensibly satisfies the scientific community's requirements will be accorded scientific standing, practitioners have a strong incentive to do only work that at least ostensibly satisfies the requirements. Typically the most straightforward way to at least ostensibly satisfy the requirements is to actually satisfy them. So scientists have an incentive to produce work that satisfies the requirements. This hardly handles every case. But it may go some way toward explaining why scientific misconduct is relatively rare.

Another incentive is this: scientists recognize that the requirements are not arbitrary. They are instrumentally valuable because their satisfaction is conducive to realizing the goal of understanding some aspect of nature. Inasmuch as scientists want to realize the goal, they have reason to employ the requisite means. As Bernard Williams notes, Watson and Crick wanted to discover the structure of DNA; they did not want merely to be thought to have discovered it. In order to realize their ambition, they had to satisfy the standards of biology. No doubt they were personally ambitious. But 'their goal [was] fame, above all fame and prestige in the scientific community, and that [would] come from the recognition that they have done good science' (Williams, 2002: 142).

A yet deeper incentive lies in the recognition that satisfying the requirements is not merely conducive of the goal, it is at least partially constitutive of the goal. Being comprehensive, impersonal, evidentially grounded, and reflectively endorsed are integral to the sort of understanding science seeks. Such an understanding is subject to, and holds up under, critical scrutiny, where the grounds for criticism and bases for assessment have been tempered through a succession of increasingly refined, increasingly rigorous processes of self-correction. Moreover, since that understanding is open to extension, correction, and refinement, a reflectively endorsed finding is not considered fixed or final, but it is an acceptable springboard for further inquiry. Scientists have an incentive to satisfy the requirements then because their satisfaction is built into scientific understanding.

Even so, the mutual trust that sustains the scientific community may seem fragile, and relying on it may seem naive. Even if scientists, qua scientists, have reason to uphold the standards, scientists are people who play other roles as well. They may be entrepreneurs with a financial stake in the outcome of their experiments; they may be parents who need to support their children; they may be egotists whose self-worth depends on the respect of their peers. One might wonder whether the incentives to cut corners are greater than those to behave with integrity.

If talented, amoral scientists emerged full grown from the head of Zeus, born with all the factual knowledge and skill needed to ply their trade, the incentives I identified might be too weak. Normally, scientists do not attempt to replicate one another's results; and when they do, an irreplicable result is typically construed as an honest mistake. So, although there is an enormous disutility in being found guilty of scientific misconduct, if his research reports are plausible, a culpable scientist is unlikely to get caught. An adroit, amoral scientist might recognize this and, being devoid of scruples, consider the risk worth taking.

Even if such a scientist grants that in general the best means to achieving the goal is to satisfy the requirements, this would not obviously provide him with an incentive to do so, if he thinks well enough of his own intuitions. He might rationalize falsifying or fabricating on the ground that, although in most cases the requirements should be satisfied, his particular case constitutes an exception. He has a good enough feel for the subject that he can just tell how the experiments would come out. I conceded earlier that if truth is the end and the scientific method a mere means, a scientist with good instincts (or luck) could achieve the end without deploying the standard means. An amoral scientist may believe that the only disincentive to cutting corners is the possibility that his conclusion will be found to be false. But, he might think, it probably is true, and if it is found to be false, it is apt to be consigned to the realm of honest mistakes. Again, it might seem, the risk he takes is slight.

Science education

Luckily scientists do not emerge full grown from the head of Zeus. People become scientists by studying science. And the ethical values I discussed are instilled in the course of their education. Even if a brilliant, amoral senior scientist could contrive a plausible, fraudulent result, doing so is not easy. For a beginning student, the best way – very likely the only way – to ostensibly satisfy the requirements is to actually satisfy them. The novice who fakes, fabricates or plagiarizes is apt to get caught; so is a novice who is careless, inattentive, or lax. Learning to do scientific research is learning to do honest, truthful, careful research. It is not learning to do research, with honesty, truthfulness, and conscientiousness tacked on as afterthoughts. So, even if he starts out as amoral as his Zeus-born senior counterpart, the novice has a strong, purely self-interested incentive to conform his behavior to the requirements. He is simply unable to cut corners and produce plausible results. What he learns in learning to be a scientist is to satisfy the requirements. He does not learn that it is normally a good idea to conduct the experiments one purports to have conducted. He learns that scientific research consists in conducting those experiments.

As his education progresses he develops the habit of conforming to the requirements; doing so becomes second nature to him. In this respect, his ethical education as a scientist is Aristotelian. He develops scientific integrity by behaving with scientific integrity, whatever his initial motives for doing so may be. And as he learns more about the subject, he is apt to gain respect for the requirements. He comes to recognize that their satisfaction is conducive to the goal, and eventually, that it is integral to the goal. That is, he learns what scientific understanding is. The trustworthiness required of scientists is not excessively fragile nor is the expectation that most scientists will be trustworthy naive, because the grounds for trustworthiness are so deeply embedded in science that they cannot help but be inculcated in scientific education.

This suggests that the education that future scientists standardly receive already fosters scientific integrity. It invites the idea that nothing more need be done. Is the NSF requirement redundant? Sadly, that is too optimistic. Some scientists are guilty of misconduct. Some students have no clear sense of where the boundaries are drawn. So the question arises: How can science education foster integrity?

When something has become second nature, the agent automatically and unthinkingly behaves in accord with it. Most scientists probably do not even consider flouting the requirements. This is both good and bad. It is good in that they simply do the right thing; it is bad in that they may fail ever to consider why it is the right thing to do. If they do not know, for example, why data must be recorded in a timely fashion, or what exactly is problematic about conflicts of interest, they may be ill equipped to withstand temptation, to adjudicate competing claims when all the desiderata cannot be simultaneously satisfied, or to recognize their obligations when they detect misconduct. By highlighting aspects of a practice that had been tacit, scientists put themselves in a position to articulate the relevant factors, discover their basis and their strength, subject them to scrutiny, and consider whether they are worthy of reflective endorsement. And they put themselves in a position to deepen science education by sensitizing students to often unrecognized aspects of scientific practice that are critical to its success. So, in addition to forestalling explicit scientific misconduct, the NSF requirement may prompt a reorientation to the practice of scientific inquiry that can augment the understanding of the discipline and its results.

Moreover, the NSF requirement pertains only to students training to be scientists. If the ethical values that underwrite trustworthiness are integral elements of science, an appreciation of their role should be conveyed in all science education, not just in the advanced education of future scientists. For all science education aims at inculcating an understanding of science. To reconfigure science education so as to highlight the ethical factors is not to skew science education; it is not a diversion designed to perform some independent educational good. Rather, to fail to appreciate the roles that these values play is to misunderstand the scientific enterprise. It is to fail to grasp why scientific findings are worthy of intellectual respect.

How might K–12 science education and basic college science courses convey an appreciation of such values? At any level, science education can incorporate the division of epistemic labor and highlight the benefits it produces. Science teachers frequently lament the lack of resources. Students must share microscopes, balances, test tubes or graphing calculators because there simply are not enough to go around. This dearth can be turned into an asset. If activities are so designed that students must genuinely work together rather than simply take turns with the equipment, they can learn something about the interdependence of scientific teams, and the need to trust and be trusted by other members of their teams.

Consider an exercise, taken from the middle school textbook *Introductory Physical Science* (hereafter *IPS*): 'If you mass the same object several times, will you find the same mass each time? Mass two objects, such as a penny and a rubber stopper, alternately several times. To avoid being influenced by your previous results, have your lab partner read and record the results' (Haber-Schaim et al., 1999: 21). This simple activity can teach two important lessons beyond the avowed objective of learning how to determine the sensitivity of a balance. It can teach the epistemic value of working together. Through their division of labor, lab partners minimize the likelihood that later readings are influenced by earlier ones. It also can underscore the need for trust. Lab partners need to trust one another in order to work effectively together. One student must place the objects properly on the balance; the other must carefully read and scrupulously record the results. Since the activity requires that the two not monitor each other's actions (else later readings might be influenced by earlier ones), each must trust that the other is doing her job.

Regarding another experiment, *IPS* says, 'To have more confidence in your conclusion, you would have to repeat the experiment several times to be sure that you have not spilled any salt or made an error in reading a scale. Many repetitions would use up much time and would be boring. So, instead of asking you to repeat this experiment, we shall bring together the results of all the experiments done by the whole class' (Haber-Schaim et al., 1999: 28). Again, there is more going on than *IPS* acknowledges. Bringing together the results of the whole class controls for systematic as well as sporadic error. If a single team did the same experiment multiple times but systematically misread the scale, their results would be mutually consistent but wrong. But if their results were compared with the results of others, the error would be manifest. Taking the results of the whole class into account thus affords greater reason for confidence than each team repeatedly doing the experiment for itself. This is so, however, only if the students are, and consider one another to be, trustworthy in their experimental work. Here trust and trustworthiness must extend to a wider community than in the previous case. Students' knowledge that the results will be compared with the results of others affords an incentive to be trustworthy in performing the experiment. Falsified or cavalier results are highly unlikely to align with the results of other teams. The practice of coalescing the results of the class encourages intellectually responsible work, not only because individual students do not want to look like fools or frauds before their classmates, but also because students recognize that only if the individual experiments were responsibly conducted does the collective result merit confidence. They gain an appreciation of the importance of repeatable results and intersubjective agreement in science, and of the reasons why trustworthiness must underlie that agreement.

Another way to foster appreciation of the role of trustworthiness in science is to assign projects where students have to build on one another's findings. If Team B must take Team A's results as inputs, and the success of Team B's investigation turns on the adequacy of Team A's results, and conversely, then each team needs to trust and be trusted by the other. Since building on the findings of others is crucial to the advance of science, such activities would augment students' understanding of the nature of scientific progress.

Science is fallible. Even the most strongly supported findings may turn out to be false. Moreover, science does not leap from success to success. The course of an investigation is typically littered with failed attempts. This suggests that science teachers might do well to take failures more seriously. If an experiment does not work out as anticipated, it can be worthwhile to figure out why. Rather than simply giving a low grade to the student who performed the experiment or telling her to do it over correctly, it may be fruitful to ask the student to explain what she did.[3] The class can then troubleshoot the experiment to find out what went wrong. Learning how to benefit from serious, intellectually responsible efforts that fail is a crucial lesson in science. Like practicing scientists, students can learn from their scientific mistakes, but only if they have the requisite information. This underscores the scientific objectionability of plagiarism and the importance of being aware of one's assumptions, of accurately describing one's methods and of carefully recording one's data. The student who plagiarized her lab report or failed to scrupulously record her data does not have the resources to account for her experimental failure.

It is not my purpose to tell science teachers how to revise their courses. The point of these suggestions is to show how the recognition of trustworthiness and its several components – honesty, conscientiousness, intellectual responsibility, etc. – might be integrated into and augment basic science education. Such integration is straightforward for the sort of inquiry-based science education that *IPS* exemplifies; more radical revisions might be necessary for courses that focus more on imparting facts than in conveying insight into how science is done. But any science course should answer the question, 'Why should we accept these claims?' If trustworthiness is critical to scientific consensus, students should be made aware of it.

Conclusion

Throughout this article, I have emphasized the obligations that members of the scientific community bear to one another, and emphasized the ways these obligations relate to the epistemic goals of science. Because scientists need to trust one another, they have devised institutional arrangements to secure the bases of trust. In publishing a paper, a scientist issues an open invitation to the community to accept her findings and use them as a springboard for further investigation. She gives her colleagues her assurance that the standards they collectively endorse have been satisfied. Extending this to the science classroom, the idea would be that students give one another their assurance that the jointly recognized standards of classroom science (which, of course, vary with grade level) have been satisfied.

This may suggest that science is insular – that the relevant assurances are conveyed only to colleagues. But the fruits of science nourish more than further research; they feed into the beliefs and actions of non-scientists as well as scientists. We take medicines, drive cars, pay bills on-line, and eat genetically engineered vegetables. In doing so, we rely on deliverances of science. As members of panels and juries, we hear scientists testify as expert witnesses (Brewer, 1998). The open invitation the scientist issues thus is wide open; it is issued to the world at large, not just to the scientific community. In publishing a finding, speaking at a conference, or testifying in court, a scientist conveys her assurance that anyone – layman or expert – can rely on her scientific claims. Although the standards to be satisfied are set by the scientific community, the moral obligation to behave with integrity is utterly general. Thus a major function of science education should be to enable educated laymen to understand why, and to what extent, the assurances proffered by scientists are credible.

The question arises whether, or how far, my argument generalizes. Do the same sorts of considerations hold for other disciplines? Do they, or their close relatives, hold for collaborative activities that do not take understanding as their goal? It seems obvious that the answer to these questions is 'yes'. The specific contours of trustworthiness may vary with the practice in question, for the nature of responsible behavior varies from one practice to the next. But the need to ground collaborative activities in trust is pervasive. This suggests that, throughout their education, students should be made aware of the ethical requirements that figure in the subjects they are studying. And they should be given the opportunity to develop the virtues that are integral to the activities they participate in. This applies not merely to academic subjects but to other activities, such as music and athletics as well. To highlight the ethical requirements at the heart of a practice, be it football or history, mathematics or marching band, is not to deflect attention from the 'real' subject, but to deepen understanding of that subject. All genuine education involves moral education.

Acknowledgements

I am grateful to Jonathan Adler, Randall Curren, Judith Jarvis Thomson, Harvey Siegel, the audience at Vassar College and an anonymous referee for comments on an earlier draft and to the Spencer Foundation for research support.

Notes

1. Kuhn (1970) denies that science is cumulative. In scientific revolutions, much that has been accepted is called into question. This does not undermine my contention that scientists build on one another's findings. For even a revolutionary scientist draws on evidence, methods, and instruments that others have generated.
2. Competence is field-based. Anyone with the appropriate knowledge and training in a specific branch of science could have done the work. There is no claim that a competent geneticist could have performed or interpreted an experiment in quantum electrodynamics, or even that the bench scientist could to the statistical analysis provided by another member of his team.
3. Alternatively, to avoid stigmatizing particular students, it might be advisable to assign an activity where the obvious approach won't work, and then ask the students collectively to figure out what went wrong.

References

Brewer S (1998) Scientific expert testimony and intellectual due process. *Yale Law Journal* 107: 1535–682.

Cohen LJ (1992) *An Essay on Knowledge and Belief*. Oxford: Clarendon Press.

Elgin C (2001) Word giving, word taking. In: A Byrne, R Stalnaker and R Wedgwood (eds) *Fact and Value*, pp. 96–116. Cambridge, MA: MIT Press.

Haber-Schaim U, Cutting R, Kirskey HG and Pratt HA (1999) *Introductory Physical Science*, 7th edition. Belmont, MA: Science Curriculum Inc.

Hardwig J (1991) The role of trust in knowledge. *Journal of Philosophy* 88: 693–708.

Kuhn T (1970) *The Structure of Scientific Revolution*. Chicago, IL: University of Chicago Press.

National Science Foundation (2009) Responsible conduct of research. Available at: http://nsf.gov/bfa/dias/policy/rcr.jsp (accessed 10 August 2011).

Venter JC et al. (2001) The sequence of the human genome. *Science* 291(5507): 1304–51.

Williams B (2002) *Truth and Truthfulness*. Princeton, NJ: Princeton University Press.

Biographical note

Catherine Z. Elgin is Professor of the Philosophy of Education at Harvard Graduate School of Education. She is the author of *Considered Judgment* (1996), *Between the Absolute and the Arbitrary* (1997), *With Reference to Reference* (1983), and co-author with Nelson Goodman of *Reconceptions in Philosophy and Other Arts and Sciences* (1988). *Address*: 404 Larsen Hall, Graduate School of Education, Harvard University, Cambridge, Massachusetts 02138, USA. [email: catherine_elgin@gse.harvard.edu]