Facta Philosophica, 1, 1999, 39-54.

Epistemology's Ends, Pedagogy's Prospects

Catherine Z. Elgin

Plato's *Meno* ends on a disheartening note. Virtue cannot be taught, Socrates concludes, because there are no teachers of virtue. And there are no teachers of virtue because no one -- not even those who are virtuous themselves -- knows what virtue is.[i] The background assumption is this: since teaching consists in conveying knowledge, you cannot teach what you do not know.[ii] Let us call this Plato's Teaching Assumption (PTA for short). At first glance, PTA seems plausible. I cannot hope to teach you the atomic number of gold if I have no clue what it is. Even if I happen to guess the correct answer and impart my opinion to you, we would hardly dignify my accomplishment by calling it teaching. `Teaching', in the sense that interests us here, is a success term, and mere inculcation of opinions (even, Plato insists, true opinions) does not qualify as the right sort of success. But the implications of accepting Plato's teaching assumption are bleak, for Socrates' conclusion generalizes far beyond the ethical realm. If one cannot teach what one does not know, it is not just virtue that cannot be taught. Neither can (much of) anything else. The requirements on knowing and teaching are too high.

Plato maintains that knowledge differs from (mere) right opinion through having a tether -- something to secure it or hold it fast.[iii] This seems right. Although epistemologists disagree vociferously about the nature and strength of the requisite tether, they generally agree that some sort of tether is needed and that it must be strong enough to confer a right to be sure. For without such a right, one does not know. Lucky guesses turn out to be correct. But because we have no right to be sure of them, they do not qualify as knowledge. Epistemologists differ over whether knowledge is contextual or acontextual, whether it rests on justification or reliable mechanisms, whether an internalist or an externalist stance is appropriate. The common denominator is that knowledge requires tethered true belief. So if you can teach only what you know, you can teach only what you have tethered true beliefs about. And if teaching consists in conveying knowledge, then teaching is limited to conveying tethered true beliefs to your students. These conditions turn out to place exceedingly restrictive constraints on the ends and means of education.

Even if we manage to evade global skepticism, we must concede that we don't actually *know* much of what we and our colleagues purport to teach. I won't embarrass you by asking how much philosophy you actually know. (Are your philosophical views true? Are they adequately justified or reliably produced? Are they so much as mutually consistent?) Even the `mature sciences' rarely yield knowledge, strictly so-called. Anomalies, discrepancies, and outstanding problems challenge the adequacy of our most strongly supported theories. They provide substantive, local reasons to suspect that those theories are not quite true as they stand. But if something is not quite true as it stands, it is strictly speaking false. Although epistemologists disagree about whether knowledge requires that remote possibilities of error be discredited, they generally agree that specific, local challenges must be met. Anomalies are not just dark reminders that Descartes' demon has yet to be decisively defeated. They pointedly call into question the adequacy of the particular theories they confront. So the strategy of bracketing skeptical worries, even if sound, will not avail us here. As long as a theory lacks the resources to meet or deflect the challenges anomalies pose, it is insecurely tethered. For there is a clear and present danger that it is false. That being so, it is not a repository of knowledge. Nor are the models it generates. Since they involve idealizations, approximations, and simplifying assumptions, they neither are nor purport to be true (or wholly accurate) representations of the phenomena they concern. If PTA is correct, no more than virtue can philosophy or science be taught.

At the cutting edge of inquiry, where cognitive advances take place, matters are controversial, truth is elusive, and any tether is bound to be fairly loose. The latest findings in a field do not immediately merit the status of knowledge. They have to stand the test of time and become enmeshed in acceptable, confirmed theories before they are adequately grounded. But if the latest findings do not qualify as knowledge, then according to PTA they cannot be taught. This means that in graduate seminars that focus on recent work in a field, teaching does not occur. Odd though that seems, it may be right. Advanced seminars at their best are collaborative exchanges, not conduits of antecedently established knowledge. Maybe `teaching' is the wrong word for what goes on there. As we back away from the cutting edge, we retreat to what seem to be more solidly grounded conclusions. So perhaps it is in less advanced courses that teaching occurs. In that case, we teach students what is known, thereby equipping them to investigate what is unknown. This accords with our words. We lead seminars, we say, but we teach introductory ethics, astronomy, metallurgy or whatever.

There are at least two problems with this proposal. One has to do with systematicity, the other with accessibility. The worry about systematicity is that there is no effective way of isolating antecedently established results from what is going on at the cutting edge. New discoveries can unsettle findings we consider firmly established and shift the grounds we take to establish them. It's not just the permanent possibility of scientific revolution that causes difficulties. We might feel fairly safe in considering steel a metal and fairly safe in believing that future inquiry is unlikely to lead scientists to conclude otherwise. So we might think the fact that steel is a metal is a bit of knowledge capable of being taught. But even if we are sanguine about the fact that steel is a metal, we may be (and probably should be) more circumspect in our assessment of our grounds. Further investigation might easily result in the refinement of the criteria for classifying something as a metal. Even if the newly sanctioned criteria didn't require us to revise our classification of steel, they might constitute a revision in the grounds for the classification. In that case, our previous claim to knowledge is undercut. If scientists used to think that something is a metal because it has a particular lattice structure *L* and later conclude that it is not *L* but related structure *L** that makes something a metal, then their previous reason for counting steel as a metal was incorrect. They believed the right thing for the wrong reason. In that case, their earlier conviction that steel is a metal was not adequately tethered. So they didn't really know it. Current investigations are designed to elaborate, extend, refine and/or challenge accepted theories. A theory has to be well established indeed before it is reasonable to hold that further advances will not unsettle the grounds for the judgments that it sanctions.

An argument from superficiality might seem to rebut this objection. "Look," one is tempted to say, "I admit that my claim that steel is a metal doesn't take us very far. It is a pretty shallow claim, backed by pretty shallow reasons. As such, it neither has nor needs a lot of deep theoretical support. My reasons for believing that steel is a metal don't rely on details about lattices, so nothing scientists discover about lattices is going to discredit my claim." This defense has the disconcerting consequence that metallurgists don't know that steel is a metal, but pastry chefs and philosophers, owing to their relative ignorance of metallurgy, do. The consequence is particularly distressing in view of the fact that just about everything pastry chefs and philosophers purport to know about metals derives ultimately from the findings of metallurgists.[iv] Nevertheless, it is worth investigating whether putatively scientific knowledge can be sustained at a suitably superficial level.

To even begin to make the case, we need to distinguish between two questions:

1) What makes steel a metal?

and

2) What justifies the belief that steel is a metal?

If contemporary metallurgy is on the right track, commitment to lattices is unavoidable in answering the first question. So scientific discoveries about lattices could easily discredit whatever answer we give to (1). Even so, it may be possible to answer the second question without adverting to lattices. Perhaps our reason for believing that steel is a metal consists in the recognition that magnets attract steel plus the seemingly well founded belief that magnets attract only metals. Nothing in the shift from *L* to *L** affects the connection we take to hold between magnets and metals. So our justification for believing that steel is a metal survives revisions in views about lattice structure.

This seems right. But it doesn't go far enough to vindicate PTA. For it does not show that our grounds are in general impervious to the revisions that relevant theories may be expected to undergo. Even if the justification we rely on is independent of changes in views about lattices, it looks exceedingly vulnerable to changes in views about magnetism. If, for instance, physics finds that some ceramics are magnetic, the justification is undermined. Considerations at the forefront of science intertwine with and support commonsensical beliefs. That being so, we should expect the grounds for our everyday beliefs to shift as science advances.

Interpretation too is subject to revision in light of new findings. What we understand when we understand the statement `Steel is magnetic' depends on and derives from the theory or system of thought it belongs to. For it is that theory that spells out the presuppositions, truth conditions, implications and implicatures of the statement. As the theory is extended, revised, and/or deepened, the interpretation of the statement evolves as well. There is, as Quine and Davidson insist, no sharp distinction between matters of meaning and matters of fact. As data accumulate, investigators often modify their assignments of extensions to terms. The discovery that magnets attract some ceramics convinces scientists that the extension of `magnetic' is broader and more variegated than their forebears thought. This leads to a further shift. Previously, `Only metals are magnetic' was, for good reason, considered an exceptionless generalization that could serve as a criterion restricting the application of the term `magnetic'. Now it is at best a stereotype. The empirical discovery that some ceramics are magnetic thus affects what we mean, how we tell, and what class we claim steel belongs to when we say that steel is magnetic. This is so whether we are privy to the latest findings or simply defer to experts to back our claims. What it means to say that steel is magnetic changes as inquiry proceeds.

Holism pulls against knowledge. If, as Quine says, statements face the tribunal of experience as a corporate body,[v] we can't know individual facts. To know the fact that steel is magnetic, we need to know a good deal about metallurgy and magnetism. We need to know what it means to claim that steel is magnetic, what such a claim commits us to, what sort of evidence supports that claim, and what makes that evidence adequate. To the extent that the relevant theories are vulnerable and to the extent that our grasp of the theories or their grounds is tenuous, so is our claim to know the fact. And according to PTA, if a theory or our grasp of it is vulnerable, so is our competence to teach that fact.

The other worry concerns accessibility. To convey knowledge, we need to impart both content and grounds. But the more elementary the course, the less equipped students are to understand the complexities of the subject. Perhaps there are adequately tethered truths about magnetism, truths whose grounds are sufficiently stable that we have no reason to expect future science to unsettle them. Perhaps the instructor knows those truths. Still, according to PTA, to teach them -- that is, to instill knowledge of them -- requires conveying both the truths and the tethers. And to impart the tethers is to convey to the students, in a way that they can grasp, both the grounds for believing them and the reasons for considering those grounds adequate. This may seem unproblematic. We're not, after all, trying to convey the intricacies of electromagnetism to a fourth grade science class. The truths imparted in elementary courses tend to be more superficial and less nuanced than the ones that more advanced students and professionals grapple with. Hence, one might think, they and their grounds are more easily taught. But the complexities that emerge at higher levels are integral to the content and grounds for the superficial generalizations we seek to impart. If a particular alignment of atoms is what makes something magnetic, then to know what is being claimed in saying that a material is magnetic requires appreciating the significance of that alignment. If teaching is conveying knowledge, we cannot teach magnetism to students who lack the resources to understand what that alignment is and why it matters.

Maybe the worry about accessibility is misguided. Granted, the instructor can't convey to novices the full content and grounds for the facts she imparts. But, one might argue, if those facts are secured by an adequate theory, and the instructor knows as much, then in imparting the facts to her students, she teaches them. This is not wholly implausible. We purport to know a variety of more or less free floating facts -- the atomic number of gold, the main product of Bolivia, the causes of the Franco-Prussian War, and so on. Often these isolated bits of information are products of educational encounters. Why shouldn't we say that we were taught such facts, we learned them, so now we know them? But if a parrot were trained to recite on demand the causes of the Franco-Prussian War or the atomic numbers of the elements, we wouldn't say that it knew them, for it wouldn't understand its own words. Even if *we* can provide its utterances with content and grounds, the parrot cannot. So it does not know. No more should we claim that a student who memorizes such matters by rote knows them. For he, like the parrot, knows not whereof he speaks. To understand an assertion requires an appreciation of what its acceptance would commit one to, and what would count as reason to accept it. Neither the parrot nor the rote memorizer has such an appreciation. Inculcating true sentences without providing the recipients with a grasp of the contents and grounds is not teaching.

Surely, one might object, students -- even students who learn things by rote -- are not in the position of the parrot. Typically they have some understanding of the contents of their claims and often have some understanding of the grounds.

This is clearly right. The mere fact that the sentences they memorize belong to languages they speak virtually insures that they have some grasp of what they are saying. But this is not enough to satisfy PTA. For often the understanding of the layman involves to no more than a stereotype. The student who can augment her recitation of atomic numbers with the claim that the elements are the fundamental units of matter would fall far short of having knowledge. (The elements aren't *really* fundamental. The difference between matter and energy is not so clear. . . .) The rote memorizer can give a gloss on her claims. But it typically is insufficiently discriminating and often is inaccurate as well. It's not nothing. But it falls short of the requirements for knowledge.

Teaching looks to be well nigh impossible. PTA insists that we can teach only what we know. Given the stringent demands on knowledge and the systematic interdependence of seemingly established and tentative findings, we don't know much. Moreover, since `teaching' is a success term, and success is achieved only if students learn, we can teach only what our students are capable of learning. If teaching is a matter of conveying knowledge, then unless the students can grasp the entire theory, or a suitably extended, isolable fragment that provides a statement of fact with content and grounds, they cannot learn, so we cannot teach them, that fact.

Rather than abandoning hope of teaching, I suggest that we reject PTA. Even if we concede (as we should) that inculcating one's lucky guesses is not the same as teaching, and that competence with the subject matter is a requirement on teaching, it does not follow that teaching consists in imparting knowledge, or that you can teach only what you know. Rather, I suggest, teaching consists in advancing understanding. How does this help?

Although one cannot know that *p* unless `*p*' is true, it is possible for understanding to be couched in *p,* even if *p* is false. This is why idealizations, approximations, and sketches are epistemically effective. The ideal gas law represents gas molecules as perfectly elastic spheres. They are not. So the law is not true. Hence we do not know that $pV = nRT$. Nonetheless, the equation provides an excellent approximation. To think of gases as though they conform to the ideal gas law is to understand a good deal about their behavior. To be sure, truth is an important epistemic desideratum. Other things being equal (or perhaps even nearly equal), we want our theories to be true. But other things are not always equal, or even nearly so. A gerrymandered, computationally intractable theory, all of whose sentences are true, is not in general epistemically preferable to an elegant theory that, by smoothing curves, invoking idealizations, and employing approximations, reveals significant regularities and generates good predictions. Truth does not, and should not, always trump.

That understanding does not require truth has several advantages. The first is sheer utility. The truth about a subject is sometimes too detailed or convoluted to be cognitively useful. The smooth curve that the data points approximate can be more revealing than the jagged line they actually mark out. If so, the understanding achieved via the approximation is preferable to the knowledge that mirrors the truth.

A second advantage derives from the fact that current knowledge is inevitably limited. Scientists always know less than they wish they did. That's what makes science a mode of inquiry rather than a repository of antecedently established fact. But this does not mean that they are utterly clueless about their subjects. Although a claim qualifies as an approximation because of its relation to a truth, we can recognize it as an approximation without knowing the truth it approximates. We are justified in considering it an approximation if it accords reasonably well with the relevant measurements we are in a position to make and the other relevant evidence we are in a position to adduce. A more accurate approximation fits the data better. Oncologists might have a model that explains how a cancer spreads. If it does not explain why some patients respond to treatment and others do not, they cannot claim that the model yields knowledge of the course of the disease. The fact that the disease does not always display the trajectory that the model predicts shows that the model is not entirely correct. The fact that it generally displays something close to that trajectory shows that the model is pretty good. So the oncologists can legitimately claim to understand the cancer better than they did before they developed the model, even though they cannot claim that it provides the knowledge they seek.

A third advantage is pedagogical. The truth about a subject may be too complex, abstruse, or counterintuitive for novices initially to grasp. What it means, why it matters, and what supplies evidence that space is non-Euclidean may be too complicated for beginning science students to understand. Providing them with a Euclidean model gives them some understanding of spatial relations, a network of cognitive commitments to refine. It may turn out that the best way to teach the geometry of space is to introduce the seemingly natural Euclidean model, show why it is inadequate, and then introduce modifications which yield a different, less intuitive geometry that evades the inadequacies. Understanding admits of degrees. A rough approximation exhibits some understanding of it subject matter; a close approximation may provide greater understanding. PTA needs to assume that education is a matter of moving from easily learned truths to more difficult truths, for knowledge requires truth. But much education proceeds by a series of approximations. We begin with a crude outline, and elaborate, extend, and emend it as we go. That space is Euclidean is not strictly true. But it is an excellent first approximation. Teaching spatial relations as though they were Euclidean provides an effective introduction to the subject.

Defenders of the epistemological centrality of truth might reply as follows: Granted that epistemically acceptable theories use idealizations, approximations, and models that are not strictly true. But they don't use any old idealizations, approximations, or models that happen to be lying around. To explain which ones they do, and should use, we still need to appeal to truth. Strictly speaking,

(a) Objects near the Earth fall toward the Earth at a rate of 32 ft./sec.^{2}

is false, since it neglects air resistance and the gravitational effects of bodies other than the Earth. But it is a good approximation, because

(b) Objects near the Earth fall toward the Earth at a rate of approximately 32 ft./sec.^{2}

is true. Even when we're dealing with approximations then, truth must be invoked to explain what distinguishes the epistemically acceptable from the epistemically unacceptable.

This is defense unpersuasive. That `approximately *p*' is true whenever `*p'* is approximately true is uncontroversial. It is also trivial. That `approximately *p'* is true contributes nothing to `*p*''s being an epistemically good approximation -- one that in current, or in ideal epistemic circumstances we have reason to endorse. Every exact truth is surrounded by a cloud of true approximations. Most of them are cognitively inert. We have no reason to believe or care that they are true. Good approximations are readily converted to truths. But bad ones are too. So if we want to distinguish between epistemicallly good and bad approximations (or epistemically better and worse approximations), we cannot look to truth. For it is not the case that whenever `approximately *p*' is true, `*p*' is an epistemically good approximation. Nor is it the case that the epistemic value of an approximation depends on its proximity to the exact truth. We need to determine what makes for an epistemically good approximation, and how good an approximation needs to compensate for its inexactitude. Once we've settled that, we can, if we like, import the `approximately' into the content of the claim and obtain a truth. But the epistemological significance of the approximation is not enhanced by this semantic slight of hand. It is the object-level approximation that does the epistemic work. The meta-level truth claim gets a free ride.

Unqualified truth, adequacy, and accessibility are epistemic desiderata that can trade off against each other. Some epistemically accessible truths are inadequate, being too imprecise or irrelevant to our cognitive goals. The knowledge that gases are often invisible, although useful in other contexts, is unlikely to do much to advance our understanding of fluid dynamics. The knowledge that gases expand when heated, although relevant, doesn't take us very far. We want to know how and why and how much gases expand when heated. But when we turn to these questions, our epistemic access to unqualified truth begins to wane. Still, more or less adequate models are available and provide some measure of insight into the phenomena we seek to understand. Even if we don't know, we're not entirely ignorant about what is going on. Moreover, those models may turn out to be epistemically preferable to the truth. Even if we could ascertain the exact sequence of molecular interactions in a volume of gas and the thermodynamic consequences of each interaction, the disorderly truth we would learn would probably do less to advance our understanding of the phenomena than the streamlined approximation that the ideal gas law supplies. For the approximation highlights patterns and regularities that would be lost in the welter of details that the motley truth disclosed. This is not to say that we should never descend to the level of the details. It is merely to emphasize that something is lost in the descent. The ideal gas law is not merely a temporary expedient or a pedagogical tool to be given up when the truth is in hand.

The quest for understanding is no mere matter of seeking to establish the truth of hitherto doubtful propositions. Tappenden makes this point vividly in connection with mathematics. Mathematicians, he notes, often value new proofs of an old theorem. If they were merely bent on establishing the fact that the theorem is true (or even provably true), their enthusiasm would be unwarranted. It is not. New proofs are valuable when they show the theorem in a different light, disclosing mathematically significant properties and relations that had not previously been recognized or exploiting resources that had not previously been used. We understand the truth in question better when we recognize that it admits of this sort of algebraic proof and that sort of geometric one, for we see what it contributes to and derives from different branches of mathematics.[vi] Similarly in other branches of inquiry. Insight is often gained by adducing new lines of evidence for previously established truths, and by embedding such truths in new contexts, thereby disclosing features and relations that had previously been obscure.

Categorization is crucial. Any object can be described in a multitude of ways, each description saying something true. But not all truths are created equal. Some are irrelevant, pedestrian, uninformative. Others engage interests, activate intellects, resolve dilemmas, and/or open avenues of inquiry. Whether we call it judgment, taste, or sensitivity, a crucial aspect of understanding is the capacity to distinguish between fruitful and fruitless characterizations. Both may yield tethered true beliefs. But it is the fruitful ones that advance understanding.

To understand a subject involves grasping problems, methods, standards, and terminology as well as facts. If the objective of teaching is the advancement of understanding, then the scope of teaching is far wider than PTA assumes. We've got to master a lot more than the established facts to understand a subject. We've got to convey a lot more than established facts to teach a subject. To teach science, for example, requires conveying an understanding of and respect for the scientific method. It requires conveying an appreciation of the role of anomalies and outstanding problems, the significance of evidence, the power of the idealizations, and the importance of the requirement that results be replicable. Merely to impart a list of facts that scientists have discovered (that $e=mc^2$, that vitamin C prevents scurvy, that hydrogen is lighter than oxygen, etc.,) would not be to teach science. To teach philosophy requires enabling students to understand and assess the significance of the arguments that constitute a philosophical position and to contrive arguments of their own. Merely to impart a list of positions philosophers have held (Thales believed that everything is water, Descartes believed that mind and body are distinct, Quine believes that whatever is is physical, etc.) or a list of the philosophical propositions the professor believes to be true (whatever is is actual, there is no necessary connection between matters of fact, etc.) would not be to teach philosophy. To teach a subject -- philosophy or physics or geometry -- is to teach how it various commitments interweave to provide an understanding of the items in the domain.

The question remains: what constitutes understanding? Truth, I said, is not required. Nevertheless, there must be some standard that distinguishes understanding from mere opinion. If we say (as we should) that there are no absolutely secure propositions on which to build our theories, and no failsafe rules of reasoning, how do we decide what belongs in a good theory or system of thought? In *Considered Judgment* I argue that we understand a subject when our relevant commitments constitute a system of thought in reflective equilibrium.[vii] Understanding advances when a system in reflective equilibrium is extended, elaborated, or supplanted by a better system. Let me say a bit about what reflective equilibrium is and what sort of justification it provides.

Whether or not they are justified, we accept some sentences, stances, standards, and methods without reservation. Being our current best guesses about the matter at hand and the appropriate ways of dealing with it, these function as our working hypotheses. We do not contend that they are surely right or to be held true come what may. But because they are our best guesses, they have some claim on our epistemic allegiance. We need a reason to give them up. I call such considerations initially tenable commitments.

To be sure, reasons for rejection are often all too readily available. Our working hypotheses may be mutually incompatible, jointly untenable, or otherwise at odds with each other. Our methods may yield inconsistent answers or provide no answers to questions we consider both relevant and significant for the subject at hand. Our standards of acceptability may endorse too many, or two few, or intuitively wrong answers. And so on. For any number of reasons, available resources may be inadequate to achieve our cognitive and practical objectives. To arrive at an acceptable theory or system of thought, we typically need to revise, extend, and correct the judgments, methods, stances, and strategies we started with. Through a process of delicate adjustments, we seek to devise a system of mutually supportive, independently supported commitments. Such a system, I maintain, is in reflective equilibrium. To achieve reflective equilibrium we may have to draw new evaluative and descriptive distinctions or erase previously drawn lines, reorder priorities or impose new ones, reconceive the relevant facts and values or recognize new ones as relevant. Although there is a presumption in favor of initially tenable commitments, it is defeasible. Nothing is immune to revision. A system is accurate if it reflects (closely enough) the commitments we began with; it is adequate if it realizes our cognitive and practical ends. Reflecting closely enough does not require and is not insured by exact replication of the commitments we began with. We realize that those commitments are incomplete and suspect that they are flawed; we recognize that our initial conception of our objectives is vague and may be muddled. We do not expect our working hypotheses to be precisely right. Nonetheless, since we have nothing better to go on, they function as guides that direct the course of research.

Reflective equilibrium is not mere coherence. A system of thought is coherent if its components constitute a mutually supportive network. Each element in a coherent system would be reasonable if we already accepted the others. But apart from the support they lend each other, there may be little reason to accept any of them. And their mutual support may depend on selective blindness to contravening considerations. When a system is in reflective equilibrium, independently motivated, initially tenable commitments underwrite coherence. Such a system is justified then because it is reasonable in light of our prior commitments.

Still, one might worry that reflective equilibrium is too easy to achieve. If, as

Quine says, we can hold any claim true so long as we make compensatory adjustments elsewhere in the system,[viii] what is to prevent us from retaining our most benighted superstitions by sacrificing the reasonable ideas that are in tension with it? There are at least two lines of defense against this eventuality. The first is that among the commitments that tether our systems are second order commitments. These characterize the epistemic requirements, standards and practices that we take our theories to be subject to. That factual statements require empirical evidence, that observation is more reliable than intuition, that controlled experiments are epistemically preferable to anecdotal evidence, that reading the entrails of birds affords no epistemic access to future events -- commitments like these constrain the revisions we are in good conscience prepared to countenance. Like any other commitments, they are subject to revision. But because of their broad scope, their revision would reverberate widely. A multitude of seemingly well grounded convictions would be called into question if, for example, we revised our views about the reliability of wishful thinking. So the payoff must be considerable for such revisions to yield a tenable system of thought.

The second is that understanding, as I construe it, is social. A system of thought is in reflective equilibrium only if it is reasonable in light of what *we* are antecedently committed to. This might seem to make matters worse. Granted, it protects against idiosyncratic commitments, but it looks as though it validates widespread misconceptions. Clearly popular opinion cannot be the hallmark of tenability, for many unpopular ideas are right and many popular ideas wrong. I agree. But it is possible to incorporate a social dimension into one's epistemology without making popularity the measure of warrant. Among our initially tenable commitments are commitments about epistemic authority -- about whose views on various matters deserve to be taken seriously. We recognize that people have different areas of expertise, so we don't give every opinion equal weight. (Some we don't even give the time of day). But if we acknowledge someone as an expert in an area, her opinions in that area rightly carry some weight with us. They are initially tenable. This does not make expert popular opinion is the standard of epistemic adequacy. Like the rest of us, the experts may harbor opinions that do not mesh into tenable systems of thought (even when their adherents think they do). So we ought not accept their deliverances blindly. But because they are initially tenable, the opinions of the experts belong to the backdrop against which we are to test the adequacy of the theories we construct. Expertise, it should be noted, is not always a matter of credentials. Someone can be an expert due to talent, experience, attentiveness, or merely being in the right place at the right time. An eyewitness's report is initially tenable simply because he observed the event we are interested in and we have no reason do doubt that he is a competent observer of such events. About some matters, expertise is widely distributed across the community. About others, the field of experts is relatively narrow. With the advancement of understanding, we refine our criteria for expertise. So over time, as our understanding of a subject and the methodologies for investigating it evolve, the standards by which we assess the experts and their commitments rise as well.

The factors I've mentioned yield grounds for cautious optimism, but afford no guarantees. As far as I can see, there are no guarantees. We can stipulate, if we like that you don't know that *p* unless *p* is true, and that a theory is epistemically unacceptable unless it is true. But since we can't entirely eliminate the possibility of epistemic misfortune, we are left with the consequence that we never know whether *p* is true, or whether the theory that incorporates it is acceptable. The commitments that belong to systems in reflective equilibrium are reasonable in the epistemic circumstances. In some epistemic circumstances -- where evidence is sparse or methods are primitive, for example -- false opinions are reasonable. I don't think this is a flaw in the theory.

A system of thought is in reflective equilibrium just in case its components are reasonable in light of one another, and the system as a whole is as reasonable as any available alternative in light of our relevant antecedent commitments. Such a system is one that on reflection we can endorse. It is tethered, not to epistemological absolutes, but to our prior understanding of the matter at hand. It does not purport to deliver irrevocable truths or permanently tenable epistemic commitments. New evidence and further refinements can upset the balance. But the commitments that constitute such a system are warranted in the epistemic circumstances. Insofar as they hang together to constitute a creditable system, they provide an understanding of the subject.

Understanding, as I construe it, is holistic. It is a matter of how commitments intertwine to form a mutually supportive, independently supported system of thought. It is advanced by bootstrapping. We start with what we think we know and build from there. This makes education continuous with what goes on at the cutting edge of inquiry. Physicists take the scientific community's consensus about electromagnetism as their working hypotheses. Fourth graders start with what they take themselves to know about magnets and metals and whatever else seems relevant. Both groups build from what they already accept, extending, revising, reconceiving as necessary to advance their understanding of the phenomena. Methods, standards, categories and stances are as important as facts. The understanding that a scientist or a fourth grader obtains from her inquiries is inseparably linked to the methods she uses, the standards she takes her investigations to be subject to, the assumptions she takes to be the uncontroversial background to her work, and the conceptual resources she has to work with. So something like E. D. Hirsch's list of facts every fourth grader should know is slightly silly. At least, knowledge of those facts would not make a child an educated fourth grader. What makes for a good fourth grade education is not the set of facts the fourth grader knows, but the level of understanding she has achieved and the resources she can deploy to advance that understanding. Facts are part of the story, but so are fictions, methods, standards, and categories. A major part of understanding is recognizing what problems remain to be solved.

Holism undermines PTA because the content of a claim derives from and depends on the system of thought it belongs to. What it means to say that iron is magnetic turns on what such a claim commits us to. Since, according to PTA, the children who haven't mastered electromagnetic theory don't know what their words commit them to, they don't know what they are saying. My account is not vulnerable to this difficulty because it contends that understanding a claim, like understanding the facts it pertains to, is a matter of degree. Since the children's system of thought is a sparse and crude fragment of the physicist's, it is reasonable to think that the physicist has a deeper, more sophisticated, more accurate conception of a magnet than the children have. She consequently draws on a richer network of presuppositions and background assumptions and her use of the term commits her to a more complex constellation of implications and implicatures. But it does not follow that the children's words are vacuous. They too draw on a network of commitments in reflective equilibrium. That network supplies them with an understanding of their words and their objects. Their network is sparser than the scientist's. So the children's conception of a magnet is comparatively impoverished. This is as it should be. But there is enough overlap between the two conceptions that we can (sometimes with a dollop of the principle of charity) recognize them as conceptions of the same thing. This overlap affords a basis for communication and a platform for teaching. Scientists and science teachers, having a greater understanding of the subject, can raise questions and introduce considerations that push the children to broaden and deepen their understanding.

If we look back at *The Meno*, we see an example of this. Socrates insists (what no one has ever believed) that he is not teaching the slave geometry.[ix] True, he is not imparting geometric truths to the boy. But he is asking leading questions that guide the slave toward a better understanding of the relation between the length and area of a square. Socrates began the exercise with a better understanding of geometry than the slave had. But there is no reason to think that he either understood or needed to understand the truth, the whole truth, and nothing but the truth to teach effectively.

Should we say than that you can teach only what you understand? Maybe. But if we do, we should recognize that the principle is more a terminological stipulation than an insight about education. If I understand enough about a matter to reliably direct your efforts to advance your understanding, as Socrates directed Meno's slave, we call what I do teaching and what you do learning. If we're pretty much on a par, grappling with the material, puzzling it out together, we call what we do collaborative investigation. It is a difference in degree, not in kind. And often it may be unclear which description is appropriate.[x]

Catherine Z. Elgin

Notes

i Plato, *Meno*, tr. G. M. A Grube, (Indianapolis: Hackett, 1976), 89d-96d.
ii2. *Meno,* 87c.
iii Plato, *Meno,* 98a.
iv For a discussion of how contemporary theories of knowledge often yield such untoward consequences, see my paper, `The Epistemic Efficacy of Stupidity,' in Nelson Goodman and Catherine Z. Elgin, *Reconceptions in Philosophy and Other Arts and Sciences,* Indianapolis: Hackett, 1988, pp. 135-152.
v W. V. Quine, `Two Dogmas of Empiricism,' *From a Logical Point of View*, (New York: Harper Torchbooks, 1961), p. 41.
viJamie Tappenden, `Geometry and Generality in Frege's Philosophy of Arithmetic,' *Synthese*, 102 (1995), p. 339.
viiCatherine Z. Elgin, *Considered Judgment,* Princeton: Princeton University Press, 1997.
viiiQuine, *op. cit.,* p. 43.
ix Plato, *Meno*, 82e.
x I am grateful to Dennis Phillips and Harvey Siegel for valuable comments on an earlier draft of this paper.