

WHAT'S THE USE?

Catherine Z. Elgin

Catherine Z. Elgin is Professor of the Philosophy of Education at Harvard Graduate School of Education. Her scholarly work focuses on issues in epistemology, the philosophy of art, and the philosophy of language. She is the author of Considered Judgment, Between the Absolute and the Arbitrary, With Reference to Reference, and co-author (with Nelson Goodman) of Reconceptions in Philosophy and Other Arts and Sciences.

ORDINARY PROBLEMS, BOTH PRACTICAL AND theoretical tend to be highly circumscribed. If we want to prove the hypotenuse-leg theorem or get from Boston to Charlottesville, a variety of factors limit our options. In the one case, we restrict our thinking to the resources, the axioms, and inference rules of Euclidean geometry; in the other, to available air routes, railroad lines, and roads, and to current schedules. We summarily exclude many conceivable alternatives. In attempting to prove the theorem, for example, induction is not an option. If millions of cases confirm the hypothesis and no exceptions have been found, we have evidence that the hypothesis is true. But no amount of empirical evidence constitutes a mathematical proof. Nor will resort to theology help. One could argue: God made the theorem true. Whatever God makes true is true. Hence the theorem is true. This is a valid deduction. If one believes the premises, one has reason to believe the conclusion, but again, not a mathematical reason. Indeed, not even every mathematical maneuver is acceptable. One easy way to prove the theorem is this: Make it an axiom, then derive it from itself, since p entails p . Clearly this won't do, even though Euclidean geometry is re-axiomatizable, and different Euclidean propositions provide equally serviceable axioms. Having good reasons isn't enough; hav-

ing good deductive reasons isn't enough; even having good deductive reasons within a Euclidean system isn't enough. We require a derivation from some standard or antecedently specified axiomatization of Euclidean geometry.

The practical problem of getting from Boston to Charlottesville is less regimented, but still subject to a variety of tacit constraints. We don't entertain such options as beaming from place to place à la Star Trek or tunnelling through the Earth. We don't consider going by way of Paris or chartering a plane. We standardly exclude options that require damming rivers, constructing highways, or creating air routes. Except perhaps for beaming, all of these are possible. We reject them out of hand because they are impractical. We automatically exclude alternatives that are inefficient, technologically unfeasible, or exorbitant. Typically, they don't even leap to mind. This suggests that the problem we set out to solve has more constraints than its statement indicates. The problem isn't how to get from Boston to Charlottesville *simpliciter*, but how to get from Boston to Charlottesville given available transportation routes and constraints on one's time and money. If we tried to be fully explicit, a good deal more would have to be said. If we switched contexts, the options would change. The Department of Transportation may have reason to seriously entertain alternatives that the ordinary traveler does not—building new roads or revising air routes, for example. Problems turn out to be like icebergs. However large they loom on the horizon, most of their magnitude lies beneath the surface, out of sight.

Everyday problems are far more complex than their ordinary statement suggests. They are circumscribed by vast networks of presuppositions, some so deeply tacit that we don't even recognize that we are making them. But far from making problems harder to solve, I suggest, this makes them easier. My point is not psychological. I'm not saying that the constraints provide cues like "It would be a good idea to check a map," or "Remember the Pythagorean Theorem." Rather, I believe, the constraints figure in the demarcation of the problem and the specification of what constitutes an adequate solution. They shape the ends in view. What is wanted is not just to prove the theorem somehow, but to prove the theorem using some standard—or even some particular—

axiomatization of Euclidean geometry, not to get to Charlottesville somehow, but to get there relatively inexpensively, directly, and efficiently using currently available modes of transportation. Typically, tacit constraints circumscribe a problem enough to allow for a solution, or at least for an understanding of what a solution requires.

If this is right, it helps explain why philosophical problems, as standardly framed, seem so intractable. In philosophical contexts, we waive presuppositions that ordinarily supply the needed constraints and fail to supply replacements to take up the slack. Rather than settling for a proof of the hypotenuse-leg theorem that takes the rules and axioms of Euclidean geometry for granted, philosophers of mathematics ask what underwrites those rules and axioms. They then endeavor to justify the axioms, validate the rules, explicate the concept of proof, determine the role of proof in mathematics, and spell out the relation of proof to truth. Instead of taking the problem of going from Boston to Charlottesville as well defined, metaphysicians back off until they confront Zeno's paradoxes: What is motion? How is it possible to move from one place to another if you have to traverse half the distance first? And so forth. The moves are familiar and, I would argue, valuable. When we uncover tacit presuppositions and subject them to critical scrutiny, we can discover whether they are well founded and whether they admit of viable alternatives. But philosophy is relentless. In Kant's words, it seeks after the unconditioned.¹ Ideally, we want to disclose, explicate, and vindicate all the considerations that underlie our convictions and actions. We seek to justify the justifiers and explicate the explicators, hoping to reach the unconditioned, where no further justification or explication can or need be given. Then and only then will conclusions be transparently correct.

This process doesn't merely demand a deeper solution to the problem at hand, it also demands a wider one. It doesn't merely ask what justifies the justifiers (for instance, how can we prove the axioms of Euclidean geometry from mathematically more basic notions), but

¹ Immanuel Kant, *Critique of Pure Reason*, trans. and ed. Paul Guyer and Allen Wood (Cambridge: Cambridge University Press, 1997) 112.

what justifies our conception of justification (in this case, what justifies this conception of proof) and then what justifies the justification of that. It requires that we keep broadening and deepening our investigation until we fully understand not only the subject at hand, but also our understanding of it.

Descartes provides the prototype of this sort of argument, but much traditional philosophy takes the same form. If I prove the theorem from Euclidean axioms, my proof is sound if I have made no mistakes, so long as the axioms of Euclidean geometry are true and the rules are truth preserving. But if I haven't established that, my proof remains vulnerable. So I need further justification. Perhaps I justify the axioms by deriving them from more fundamental mathematical truths. Then, of course, I must vindicate the more basic truths. If the series of justifiers comes to an end, there is an ultimate order of things. Theories in a domain are justified only if they are derivable in the appropriate way from the ultimate order of the things that constitute that domain. Moreover, the conception of justification must itself be justified.

Even if we are convinced, as Thomas Nagel is,² that there has got to be such an order, we lack the resources to demonstrate that we've found it, much less that we've properly derived our conclusions from it. We cannot prove that our best methods reveal the way things really are, that our results would hold up even if there were a demon as powerful as God but bent on deceiving us, or the real order afforded systematically misleading evidence to beings endowed with our perceptual and cognitive capacities. Since there is no guarantee that our best efforts succeed, or that ideal theories are true, realism leads to skepticism. "What's the use?" realists may reasonably wonder. When you play against the demon, the demon always wins.

Postmodernists seem to arrive at the same question more quickly. Seeking after the unconditioned is a mug's game, they believe, for there are no philosophical absolutes to be found.³ To think otherwise is just

² Thomas Nagel, *The Last Word* (New York: Oxford University Press, 1997).

³ Richard Rorty, *Consequences of Pragmatism* (Minneapolis: University of Minnesota Press, 1982).

to delude ourselves—to confuse features of our parochial, patriarchal, Eurocentric form of life with the order of nature. There is no God's eye view, no absolutely neutral standpoint outside theories and practices from which to evaluate them. Without such a standpoint, any pretense of justification is, they believe, at best misguided. "What's the use?" postmodernists ask. In the greater scheme of things, one stance is as good as another. More precisely, without absolutes, there is no basis outside a stance for distinguishing good from bad. The greater scheme of things is a chimera. Such reasoning may result in ironic detachment from all stances, or passionate allegiance to whatever stance one happens to occupy, or playful wafting from one adopted stance to another. But if the postmodernist position is correct, there can be no good reason to favor one frame of reference over another.

Elsewhere I have argued that both positions are epistemologically untenable.⁴ If realism is right, there is a basis whose structure is the ultimate ground of justified theories and practices. But our epistemic resources are inadequate to demonstrate conclusively that we have found it. So we don't know which, if any, of our theories and practices are in fact justified, for we don't know which ones derive from the ultimate ground. If our goal is to distinguish between justified and unjustified theories or practices, then realism, even if true, is unhelpful. This makes realism useless for action. It provides no basis for deciding to do, or to think, one thing rather than another. If postmodernist anti-realism is true, there is no absolutely independent ground for assessing theories or practices. There is no way to tell that one is better than another except that we prefer it. And preferences, though they might have causes or influences, do not admit of justification. This, too, is unhelpful. It affords no guidance as to how or why or whether to adopt, reject, or modify a theory or practice.

According to postmodernism, there is no guidance to be had, since there is no external standpoint from which a theory or practice can be

⁴ See my books *Considered Judgment* (Princeton: Princeton University Press, 1996) and *Between the Absolute and the Arbitrary* (Ithaca: Cornell University Press, 1997).

assessed. Since all criteria of evaluation are internal to the system to which they pertain, they are inevitably self-serving. This is supposed to follow directly from the failure of philosophy to discover absolutes. I disagree. The claim that there is no standpoint outside a theory or practice is equivocal. It might mean that there is no single standpoint outside all theories and practices. Or it might mean that for each individual theory or practice, there is no standpoint outside it. The former is probably true.⁵ The latter, I believe, is false. Moreover, it is false in a way that is epistemologically important. I suggest that it is possible to adopt critical stances that enable us to make epistemologically responsible assessments of our various theories and practices without falling into self-serving rhetoric or thinking that we speak with the voice of God.

Plainly, people criticize one another. Physicists sneer at sociologists and sociologists sneer at journalists. Each finds the objects of his contempt defective by his own professional lights. No one denies that work valued in one field often fails to satisfy the standards of another. But the questions inevitably arise: Who are you to say that your methods, standards, or results are better than ours? What privileges your epistemological stance? Postmodernism contends that these questions have no satisfactory answers. We can demonstrate that physics and sociology do not satisfy the same standards. That just shows that they are different disciplines. It provides no reason to think that either is epistemologically preferable to the other. The argument generalizes: not only is there no good reason to consider physics epistemically better than sociology, there is no good reason to consider chemistry epistemically better than alchemy, for there is no standpoint from which to make objective assessments. If one field successfully denigrates another, the explanation, postmodernism insists, is political. The successful denigrator has more power and uses that power to sustain its privileged position. Unless the

⁵ We could not, of course, demonstrate its truth by adopting an overarching perspective from which the non-existence of such a standpoint can be seen. But if the non-existence of such a standpoint follows from an otherwise tenable theory, and no such standpoint has been found, we have good reason to believe that none exists. Indeed, it is not clear that the concept of such a standpoint even makes sense. For any candidate, it would seem possible to adopt a yet more distant standpoint and see how the candidate standpoint relates to the things it surveys.

outputs of a theory or practice fail to satisfy the standards that they should satisfy, any sneering is unfounded. The problem is to underwrite the *should*.

If we insist that the *should* derive from *unconditional* justification, realism and postmodernism seem to exhaust our options. But we need not define the problem that way. Instead of taking “What’s the use?” as the verbal equivalent of throwing in the towel, I suggest that we take the question seriously. Once we specify what we are trying to do, we can begin to identify standards and criteria to assess our efforts. This allows for a type of conditional justification, where the conditions under which justification obtains are reasonable, revisable, and subject to evaluation.

Inquiry is a quest for information, knowledge, understanding, or the like. For simplicity of exposition, I will collapse these into the quest for information. If we can specify what we want information for, we can begin to articulate what sort of information would serve our purposes. The purposes supply constraints. An example may bring this out. Why is prediction a scientific desideratum? What use is it? To predict a future event is to infer, on the basis of antecedently available evidence, that the event will occur. We use preregistration figures and last year’s enrollments to predict how many students will take the course next term. We use the positions and trajectories of asteroids, along with the laws of motion, to predict where the asteroids will be next week. There are at least two reasons to value such inferences. One is that we want to know in advance what is going to happen. We predict the number of students who will take the course so that we can get a reasonably sized classroom and order enough books. Another use of prediction is to test a theory. If we devise a theory to accommodate a certain range of data, it is no surprise that it accommodates those very data. But merely to accommodate the antecedently available data is insufficient to show that the theory is not *ad hoc*. A good theory should extend its explanatory range beyond the data it was expressly designed to accommodate. When its predictions of previously unknown facts are borne out, the theory is confirmed.

If we want to know in advance what will happen, our inference must be complete before the event occurs. If the calculation required to pre-

dict next Thursday's weather is so complex that it takes six months to perform, the result will be obsolete by the time we get it. A relatively precise calculation that takes too long to perform does not serve normal weather-predicting ends. A faster calculation that yields a rough approximation is preferable. If we want to test a theory, a more precise, time-consuming calculation may be appropriate. If we seek to predict Thursday's weather, not in order to decide what to wear or whether to evacuate low-lying areas, but to evaluate a new theory of turbulent flow, the time lag would be tolerable. It makes no difference whether the predicted event occurs before the calculation is completed, so long as the data were available prior to the event. An accurate prediction would confirm the theory. Prediction then has two different uses. Often they are jointly realizable, if realizable at all. Sometimes they diverge. To decide whether we want a quick, rough approximation or a slow, more precise answer, we need to know what the prediction is for.

Kant distinguishes between hypothetical and categorical imperatives, arguing that if we seek a rule that commands absolutely, we need a categorical imperative.⁶ For a hypothetical imperative says only: "If you want p , you should do q ." It says nothing about whether you do or should want p , or about what you should do if you do not want p . Kant is right about the structure of hypothetical imperatives,⁷ but epistemologists should not dismiss them on that account. For investigation into hypothetical imperatives discloses connections between means and ends. If you want p , you should do q . So because you want p , other things being equal, you should do q . The *should* in this case derives from your end.

The disinterested growth of understanding turns out to be not so disinterested as epistemologists sometimes pretend. When we set out to investigate a phenomenon, we typically do not want just to understand it somehow or other. We want an understanding of a certain sort. We might hope to couch our understanding in a particular vocabulary. For

⁶ Immanuel Kant, *Foundations of the Metaphysics of Morals* (Indianapolis: Bobbs Merrill, 1969).

⁷ And he may be right about their inadequacy in ethics. On that I take no stand.

example, we might want an analysis of a painting to be in terms of colors, shapes, textures, and styles, not in terms of the chemicals found at various points on the painting's surface. We might want our understanding to reflect a particular methodology, for example, preferring controlled experiments to field work. Perhaps we want our findings to mesh with the findings in related fields. Biological investigations might restrict themselves to methods, issues, measurements, and vocabularies acceptable to chemistry or medicine. Art criticism might limit itself to vocabularies, emphases, and techniques countenanced by literary and cultural criticism. These constraints are by no means mandatory. There are plenty of other interesting and valuable things to find out about paintings and proteins. So we need the conditional: If you want to achieve a particular cognitive end, you should use (or restrict yourself to) particular means.

Investigation might disclose alternative ways to achieve that end. If so, we need to consider what, if anything, favors one over others. We might discover that although we have long thought that particular means produce a given end, or increase the probability of it, this turns out to be false. We might find that although the means bring about the end or increase its probability, the cost is exorbitant. The issue then is whether the end justifies the means. To discover that, we need to identify the end and the means and delineate the connection between the two. We need, moreover, to locate ends and means within a space of alternatives. Some means may be acceptable only when less onerous means are unavailable. Others may be unacceptable in any case.

When used as a reproach, the charge "The end justifies the means!" implicates that a course of action is morally beyond the pale. The point of the objection is that there are some means that are unacceptable, no matter what ends they would promote. This is surely true. That we could bring about the desirable result of raising test scores by executing the bottom half of the class does not justify mass execution. Nothing justifies that means. Nevertheless, it pays to consider the matter further. That there are some means that no end can justify does not entail that ends generally do not justify means. Indeed, one might ask, if the end doesn't justify the means, what does? If I want to get to Charlottesville, I have reason to board a plane that goes to Charlottesville. If I don't,

then (unless I have other, undisclosed objectives) I have no reason to get on that plane. At least part of what justifies my taking a particular plane is that it is likely to get me where I want to go.

A worry remains: means/ends reasoning applies to action. It is easy to see how my objective figures in the assessment of my getting on a plane or sacrificing a pawn. My objective might even figure in my using the proof to figure out if the theorem is true. But it wouldn't have anything to do with pure theory. We want to know whether the sequence of propositions adduced really proves the theorem. That is a different question from whether it does or even should convince me that the theorem is true. The sequence of propositions isn't an action. Does it make sense to suppose that a sequence of propositions serves some end? I think it does. In proof, the issue is not what any agent does, but what the sequence of propositions does. The goal is a proof of the hypotenuse-leg theorem. The question is: Does the sequence of propositions provide one? Granted, the goal in question is our goal. The proof is one that someone produces. But the question "Does the sequence of sentences constitute a proof?" can be asked, and the sequence can be criticized as invalid, or incomplete, or redundant, by reference to the end of proving the theorem. Similarly, we can ask of evidence whether it is evidence for a hypothesis, whether it is indicative or demonstrative, or misleading, or whatever. To ask these questions requires understanding that confirming or disconfirming the hypothesis is the end, and the evidence is a means.

Nevertheless my solution remains too thin. It disregards contextual features that bear on the justification of individual actions and statements. Many actions and statements that would be unjustified in isolation are justified in a particular context. There is no direct evidence of the existence of positrons. Considered by itself, the sentence "Positrons exist" is epistemically unjustified. But, of course, the sentence is not and ought not be considered by itself. It functions as part of a physical theory. We are justified in believing that positrons exist because symmetry principles would be violated if they did not, and symmetry principles are integral to well-substantiated physical theories. The commitment to positrons is then justified by its role in a theory that is itself justified. John Rawls distinguishes between justification *in* a practice and justifi-

cation *of* a practice.⁸ Very roughly, a practice is a constellation of rules or conventions that define the roles, positions, moves, permissions, prohibitions, and standards that give an activity or system of thought its structure. A game is a paradigm case. The rules determine what a player may do, must do, and must refrain from doing. The internally defined goal of the game determines what constitutes a good or bad move, hence what it is to play well or badly. A good move or strategy tends to increase one's prospects of winning. So whether a given move is good depends on what it takes to win. The game stipulates that. But there is nothing particularly action-oriented about practices as I've characterized them. Theories and other systems of thought qualify as practices to the extent that they are structured by suitable constellations of rules and conventions.

As postmodernists maintain, many human activities take the form of practices. Hence, the justification for many acts, statements, objectives, methods, and priorities derives from the practices to which they belong. Journalism, for example, might require at least two independent sources for each reported fact. Science might require that experimental findings be replicable. Law might require that arguments be grounded in constitutional principles and accepted precedents. If so, the justification for the inclusion of a particular sentence in a news report, a scientific journal, or a legal brief is, in the first instance, by reference to the practice in which that sentence functions. The justification for omitting a seemingly relevant sentence is that it does not conform to the rules of the practice. The reporter cannot confirm his informant's claim; the experiment resists replication; or the statement, being hearsay, is inadmissible. Within the practices in question, these constitute good reasons for omitting a sentence.

Other elements of a practice are justified in the same way. Cognitive practices supply standards of rigor and relevance, criteria for counting something as evidence or as a reason, methods for assigning weights to

⁸ John Rawls, "Two Concepts of Rules," *Collected Papers*, ed. Samuel Freeman (Cambridge, MA: Harvard University Press, 1999) 20.

evidence, and so on. The practice supplies the rationale for its components. Science limits the precision of its findings in order to achieve intersubjective agreement. If findings purported to be exact, we could not tell whether two experiments yielded the same result, rather than diverging beyond the threshold of measurement. We could never determine whether the demand for replicability had been satisfied. To achieve agreement then, science quite reasonably sacrifices exactitude. To be sure, the required level of accuracy can vary, depending on the needs and resources of the scientific activity in question. So accuracy can in principle always be increased. But there is inevitably a limit. However high we set the threshold of measurement, findings will be deemed to be accurate only to a given number of significant figures. Art criticism makes the opposite choice. It values sensitivity and sacrifices agreement to allow the finest discriminations to be made. Art criticism allows the possibility of permanently irresolvable disagreements; science does not.

Elements of a practice can be evaluated by reference to the whole. Those that enable the practice to function are justified by their contribution. Those that contribute nothing, or interfere with the conduct of the practice or the realization of its internally defined goals, are not. The commitment to positrons and the restriction on precision are thus justified in physics by the roles they play in the science.

We can also ask, independently of the practice or system they belong to, what justifies particular commitments. The answer may be: nothing. There is no good epistemological reason to believe an empirical statement that is unsupported by evidence. Nor is there any good epistemological reason to prefer imprecise statements to precise ones. So, independent of the practice they belong to, neither the commitment to positrons nor to limits on scientific precision seems acceptable. But that is irrelevant. The real question should be, not what is the justification for limiting precision or believing in positrons *simpliciter*, but what is the justification for a system that includes the commitments as elements.

Justification in the first instance, I said, may be justification in a practice. But the first instance is not the last instance. Justification in a practice consists in satisfying the standards of the practice, whatever

those standards may be. A finding or procedure is justified in alchemy if it satisfies the standards of alchemy, just as a finding or procedure is justified in chemistry if it satisfies the standards of chemistry. Manifestly this does not show that these findings or procedures are epistemologically acceptable. We need to vindicate the practices that warrant them. In some cases this involves subsuming them under other practices. Since neurophysiology is a branch of physiology, neurophysiology's standards are vindicated by the physiological standards that yield them as a special case. Often, however, this cannot be done. Even when it is done, it simply invites the same question about the subsuming practice. The findings in a branch of alchemy are vindicated by their relation to the full theory of alchemy only if the full theory of alchemy is vindicated. What determines that?

The answer, I venture, is this: In general, to ask what justifies a theory or practice is to ask what good it does, that is, to ask "What's the use?" If the sole or overriding objective of alchemy is to transmute base metals into gold, then its unremitting failure to do so shows that the theory is not justified. In that case, justification in the theory is idle. The alchemical finding that is justified in the alchemical theory remains epistemically untenable. Alchemy, being a one-problem field, is an easy case. We can readily say what it is trying to do, and easily determine that it has not succeeded. Most theories and practices have multiple aspirations and may succeed at some but not at others. Moreover, many theories and practices have unrecognized uses. This complicates the issue, but does not, I think, make any fundamental difference. To decide whether a theory or practice is justified, we may have to identify a variety of uses and assess its performance with respect to each of them. This allows for mixed verdicts. A theory or practice might be justified in some respects, but not in others.

The ends in view may be purely theoretical. The distinction between justification *in* and justification *of* rests on two requirements: 1) ends must be identifiable and characterizable independently of the theory or practice that endorses them, and 2) a theory or practice that succeeds by its own lights might in principle fail to achieve its independently characterized ends. When we ask about the ends of philosophy, pure mathematics, or physics, the issue is not, or at least need not be, what

sort of social good or practical outcome they aim to produce. It may be what kind of cognitive good they aim for, what kind of knowledge or understanding they seek to deliver, and whether and to what extent they succeed.

Once we allow for purely theoretical goods, the end might seem obvious. Every academic discipline, one might argue, is a one-problem field in much the way that alchemy is. But rather than seeking to transmute base metals into gold, academic inquiries seek to discover Truth. This answer is both too broad and too narrow to be correct. It is too broad, since every domain includes untold numbers of trivial truths that are of no interest whatsoever. It is too narrow, because every discipline has commitments that cannot be cashed out in terms of a commitment to truth. Science values simplicity. But there is no evidence that a simple theory is more likely to be true than a complicated one. Mathematics demands proof, even though it follows from Gödel's theorem that not all mathematical truths are provable. Philosophy uses idealizations, approximations, and thought experiments to advance understanding. No one would want to scrap the Myth of the Cave because it is not literally true.

Different theories, disciplines, and cognitive practices may have different ends. Mathematics seeks proof. Natural science demands empirical evidence. Some investigations seek theories that deploy a particular taxonomy, or mesh with theories in neighboring domains, or respect certain prior commitments. Biology might seek theories that accord with chemistry's accounts of the same phenomena. Thoughtful people might want psychological findings to be useful for self-reflection. Common sense might demand that fundamental physical theories be deterministic. If we adopt such objectives, various otherwise attractive theories will be found wanting. And there is room for dispute about which of these we ought to endorse and how firmly we should hold on to the various commitments.

Still, the claim that epistemic ends justify epistemic means is misleading, for it suggests a strict hierarchical order of justification that cannot be sustained. If the ends justify the means, what justifies the ends? The hierarchical model yields no satisfactory answer. One might temporize,

saying: "Further ends." Sometimes this is correct. Intermediate ends are justified because they promote final ends. But what justifies the final ends? It can't be yet further ends. There are none. So it seems that either the final ends are self-justifying, or they lack justification. If we say that they are self-justifying, we are in the embarrassing position of being unable to explain what self-justification is and which ends have it. As of yet, no one has been able to figure this out. If we say that they lack justification, we seem forced to acknowledge that ultimate epistemic ends are arbitrary. There is just no disputing matters of epistemic taste. This fails to do justice to human cognitive achievements. It seems just wrong to conclude that there is nothing to be said in favor of the theory of relativity or Euclidean geometry or the *Critique of Pure Reason* except that we like it.

We seem driven to these unpalatable alternatives only because we buy too quickly into the hierarchical model of justification. Rather than being linear, justification has, I believe, a more complex, web-like structure. Ends justify means, but means also afford some measure of justification for ends. Ordinarily, if a theory does not promote our epistemic ends, we fault the theory. So ends typically dominate means. But the dominance can be overridden. If an otherwise attractive theory fails to promote given epistemic ends and if modifications designed to realize those ends either fail or make the resulting theory manifestly less attractive, we have reason to rethink our commitment to those ends. It is, for example, antecedently reasonable to require acceptable physical theories to be deterministic. In that case, physical theories are unacceptable if they do not yield deterministic models of the phenomena they concern. But the cost of this commitment turns out to be exorbitant. We can retain the commitment only by rejecting otherwise fruitful theories and admitting blind ignorance about what is going on at the quantum level. Eventually (and perhaps regretfully), we revise our epistemic ends to bring them into accord with available and foreseeable means, and admit that irreducibly probabilistic theories are acceptable. Nor is it only when ends are unrealizable that means provoke revision. It is not unreasonable to suppose that the ultimate goal of mathematics is provable truth. That readily realizable goal accounts for the value mathematics attaches to proof and for its repudiation of empirical methods. But an aspect of mathematical practice appears unwarranted if provable

truth is its end. Mathematics values the generation of multiple, diverse proofs of the same theorem. This looks wildly redundant if all we want to know is whether a theorem is provably true.⁹ We might nevertheless endorse the end and conclude that the generation of multiple proofs is indeed redundant, and hence unjustified. Or we might reconsider the end and argue that the multiple proofs reveal something of the warp and woof of mathematics, showing how the various branches of the subject interweave. This seems to be a good thing. To make sense of it as mathematically valuable then, we need a more textured understanding of the epistemological ends of the subject than the goal of provable truth provides. That provides an incentive to redefine our mathematical ends.

The worry in all of this is circularity. If the ends justify the means and the means justify the ends, we seem left with an insular, self-sustaining, but otherwise ungrounded constellation of commitments. But there are more strands to the web of justification. To understand what is to be said for a theory or system of thought, we don't just consider a static constellation of commitments, we consider dynamics. That is, we ask how it relates to its predecessors and successors. A new theory or system of thought should be at least as reasonable as its alternatives in light of our relevant antecedent commitments. Factual commitments should reflect our prior understanding of the phenomena. Methodological and evaluative commitments should reflect our prior understanding of how to investigate such phenomena and assess our findings. New theories need not entirely agree with their predecessors. Knowledge does not always grow by accretion. But even if available theories are flawed, because they constitute our current understanding of the subject, they serve as touchstones. A new theory is preferable only if it constitutes a better understanding than the one we already had. It need not, and typically will not, entirely agree with its predecessor. But it should be reasonable in light of its predecessor. Ideally, it should explain why the old theory was right to the extent that it was and why it went wrong where it did. Minimally, it should explain why

⁹ I am grateful to James Tappenden for this point.

the previous theory did as well as it did. Norms for assessing improvement likewise evolve. We start with our current best understanding of epistemic norms and revise, correct, and amend them to bring them into accord with each other and with our assessments of the theories, beliefs, and analyses that we consider epistemically reputable.¹⁰ With the advancement of understanding come rising standards of evaluation.

Ends, once achieved, are not fixed and final resting places. They become the means to formulate, motivate, and engage in the pursuit of further ends. The mark of a good answer is that it sparks good questions. So the final aspect of the justification of a theory or system of thought is that its achievement provides a springboard for further epistemologically fruitful inquiry. Acceptable epistemic commitments then turn out to be multiply tethered. No single justificatory strand can bear the full weight. But taken together they show a statement, method, standard, or theory to be reasonable in the epistemic circumstances.

To ask, in all seriousness, “What’s the use?” is to ask how an item contributes to ends in view. Those ends provide constraints on what would constitute acceptable means, in that if one wants to achieve those ends, only a certain range of means will do. But there is nothing sacrosanct about the ends. We can investigate them in turn. Why do we want to satisfy those ends? What is to be said for and against them? What does the connection of these means with those, or other, ends reveal about their acceptability? If ends and means don’t mesh, should we revise the ends or the means? I do not claim that appeal to pragmatic factors yields the absolute grounding that Descartes sought. But pragmatic factors provide standpoints outside of a given theory from which to assess it. If we find that it does not serve its avowed ends, we may revise or reject the theory or the ends. If it does serve those ends, we may regard it as provisionally acceptable, pending further elaboration and assessment of means and ends. Justification remains provisional and precarious. The good news is that assessment is possible. The bad news is that it is always subject to review, revision, revocation. There is no last word.

¹⁰ See my book *Considered Judgment*.