

COPYRIGHT NOTICE:

Russell Hardin: How Do You Know?

is published by Princeton University Press and copyrighted, © 2009, by Princeton University Press. All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher, except for reading and browsing via the World Wide Web. Users are not permitted to mount this file on any network servers.

Follow links for Class Use and other Permissions. For more information send email to: permissions@press.princeton.edu

Ordinary Knowledge

WHILE SAMUEL JOHNSON [1709–1784] was working on his book on the lives of the English poets, James Boswell [1740–1795] volunteered his assistance in lining up a conversation with Lord Marchmont about Alexander Pope, one of the poets discussed in the volume. Johnson dismissed the offer, saying, “If it rained knowledge I’d hold out my hand; but I would not give myself the trouble to go in quest of it” (Boswell [1791] 1976, 989).¹ (Pope was, of course, the same Alexander Pope who noted that a little learning is a dangerous thing.) Johnson’s perhaps dyspeptic attitude captures a large part of the difference between ordinary and scientific knowledge, and I will use his two metaphors—raining knowledge and going in quest of it—to characterize the difference. Ordinary knowledge is almost entirely grounded in hearsay from a supposedly credible or even authoritative source, although commonly the credentials of the source are not compelling and perhaps even more commonly we can no longer remember the source or its quality. Psychologists refer to “source memory” and they commonly find that we remember a fact but have no memory of how we came by it. We will typically not double check what our newspaper or encyclopedia may once have said; we will stop our inquiry sooner rather than later.

Johnson lived in an age that we think of as having been shaped by Isaac Newton [1642–1727], the archetypical scientist, all of whose science was done well before Johnson was born.² But as a rule, Johnson no more thought scientifically than did anyone else who lived centuries earlier—or since. If he had been a poor farmer, he might have given more thought to explanations of natural phenomena than he did as a remarkably educated and urbane man. Newton’s revolutionary scientific thinking did not

¹The full exchange has the quick charm of many of Boswell’s reports. When told that Boswell ([1791] 1976, 988–89) had arranged for Marchmont to call on him the next day, Johnson retorted, “I don’t care to know about Pope.” Mrs. Thrale took Boswell’s side and chided Johnson, “I suppose . . . that as you are to write Pope’s *Life*, you would wish to know about him.” Johnson said, “Wish! Why yes, if it rained. . . .” Perhaps Johnson was merely in one of his sour moods (Boswell refers to “something morbid in his constitution”). In any case, a year later he did speak with Marchmont.

²In the morality tale for children, Newton’s understanding of gravity began when an apple rained down on his head as he sat lazily beneath a tree. He did not even put out his hand for it.

have great influence on the way ordinary people thought in his time. Indeed, it did not affect Newton's own astonishing defenses of certain religious beliefs about the date of the creation of the universe in 4004 B.C.E., an issue that alone may have occupied more years of his long life than did all of his scientific investigations. Even today, scientific thinking has far less influence on daily thinking than some might wish. I will argue that *this fact reflects rationally correct views and is commonly even a good, not bad, thing. Recognizing this fact is not to criticize ordinary knowledge or ordinary people.*

A beauty of Johnson's comment is that it shows how well he understood the nature of his knowledge. He might not have known very well how the knowledge of scientists in their specialized realms differed from ordinary knowledge, but he clearly knew that most of his bits of knowledge had a somewhat laconic quality. He happened upon them or they happened upon him, and he perhaps occasionally held out his hand to catch bits of knowledge. Most of them were not a matter of his craft but only of his experience. Most of us might not have even so much as Johnson's self-understanding that our knowledge is similarly happenstance. We are often as sure of our casual, happenstance knowledge as any scientist is of the results of some massive, careful study, whether experimental or merely observational.

Philosophical theory of knowledge is largely about a kind of public, not personal knowledge. What must interest anyone who wishes to explain behavior is the knowledge or beliefs of actual people. An *economic* theory of knowledge would address this issue. Such a theory would not focus on the objects of belief but on the ways people come to hold their beliefs, and on the overall structure of their beliefs. By an economic theory, I mean merely a theory that focuses on the costs and benefits of having and coming to have knowledge, or to correct what knowledge one has. It must fundamentally be a theory of trade-offs between gaining any kind of knowledge and doing other things, such as living well.

An economic theory of knowledge would be grounded in three quite distinct facts, all of which matter to anyone whose knowledge we wish to explain. First, knowledge has value as a resource and is therefore an economic good; hence, people will seek it. Sometimes we seek it at a very general level, as when we absorb what we can of a general education. In this case, we may have little idea of how we are ever going to use the knowledge, and we may not know in advance much about the range of the knowledge we will be acquiring. Sometimes we seek it for a very specific matter, as when we seek mortgage rates when buying a home. In this case, we know very well what we want the knowledge for, and we know reasonably well where to get it and when we have gathered enough of it.

Second, the acquisition of knowledge often entails costs, so that its

value trades off against the values of other things, such as resources, time, and consumptions. Moreover, these costs are often very high. For example, the costs of gaining enough information to judge the political candidates in an election are thought commonly to be far too high for most voters in the United States to be able to justify the expenditure, especially given that they have little to gain from voting anyway (see chapter 3 on democratic participation). Instead, they vote on the strength of relatively vague signals about issues they do not adequately comprehend.

And third, a lot of our knowledge, which we may call “happenstance knowledge,” is in various ways fortuitously available when we have occasion to use it. Some knowledge comes to us more or less as a by-product of activities undertaken for purposes other than acquiring the knowledge, so that in a meaningful sense we gain that knowledge without investing in it—we do not trade off other opportunities for the sake of that knowledge. For example, you know a language because you grew up in human society. Much of what is loosely called social capital is such by-product knowledge. If you grew up in a bilingual or multilingual community, you likely know more than one language. By-product knowledge may simply be available to us essentially without cost when we face choices. Some knowledge may even come to us as virtually a consumption good. For example, your love of gossip may lead to knowledge that is quite valuable to you. Finally, the knowledge in which you deliberately invested yesterday for making a specific choice may still be available to you today when you face some other choice to which it might be relevant.

Again, the account that follows is of subjective, not public knowledge. It is not concerned with what counts as knowledge in, say, physics, but rather with your knowledge, my knowledge, any specific individual’s knowledge. The principal interest here is not in the theory of knowledge itself but with how this account of knowledge informs analyses of various kinds of belief and behavior, such as ordinary moral choice, religious belief and practice, political participation, liberalism, extremism, popular understandings of science, and cultural commitments.

There is an important category of knowledge that is not at issue in the discussion here. Gilbert Ryle ([1946] 1971) distinguished between knowing that and knowing how. I know that the height of Mont Blanc is about 5,000 meters. I know how to ride a bicycle.³ The latter kind of knowledge is experiential. The former kind is typically not experiential for most of what ordinary people know—and we are all ordinary people in most realms. But our capacity for knowing that may turn on our knowing how. For example, it surely depends on our mastery of language, which is partly

³Indeed, if you rode one as a child but then not again for decades, you would almost instantly reacquire the motor skill of cycling. In scientific terminology, this phenomenon is called ‘savings.’

a matter of knowing how, in the strong sense that most of us could not begin to articulate our knowledge of language and how it works for us, anymore than most of us, or even the best scientist in the field, could articulate how we manage to keep a bicycle underway without falling over (Wilson [1974] 2004, esp. chap. 8). The central concern of this work is with cognitive knowledge, or knowing that.

The theory of knowledge we need here is pragmatic; it has closest affinities with recent social epistemology (see contributions to Schmitt 1994). Social epistemologists agree that much of even scientific knowledge depends on testimony, rather than on direct investigation. They wish to establish criteria according to which testimony can yield justified true beliefs. The focus of an account of pragmatic street-level knowledge and belief is on the use and subjectivity of knowledge, not on justification of any claim that it is “true” knowledge. It is subjective because it is about *your* knowledge or *my* knowledge, not about knowledge per se.

AN ECONOMIC THEORY OF KNOWLEDGE

What we need for understanding knowledge at the level of the individual, and not merely at the level of a super knower who does not suffer human limitations, is an explanatory theory, not merely a definitional or essentialist theory. I propose a pragmatic theory that can be characterized as economic. It is economic in the sense that we can explain bits of knowledge that a given person has as being substantially affected by the costs and benefits of obtaining and using those various bits. Moreover, we can explain the retention of bits of knowledge in the face of competing knowledge by the seeming costs and benefits of retaining them, balanced against the costs and benefits of revising or rejecting them. In a widely quoted remark, F. Scott Fitzgerald (1945, 2) says that “The test of a first-rate intelligence is the ability to hold two opposed ideas in the mind at the same time, and still retain the ability to function.” As is true of much of Fitzgerald, the remark is clever, but misguided. Everyone who functions at all manages to do so while holding many opposed ideas in the mind at the same time, from sunup to sundown. Surely, however, not everyone who functions at all has a first-rate intelligence.

For the understanding of human behavior—normative or merely factual—we require not a philosophically general theory of knowledge, but a street-level account, a theory of the general pattern of individuals’ available knowledge. To assess knowledge that is practical, that is actually put to use by ordinary people, our account must be subjective. That is, it must be grounded in the particular individual whose knowledge it is. To speak of knowledge is to speak of someone’s knowledge. Otherwise,

it is the knowledge that exists somewhere, even out in the ether, as, at the extreme, in the noosphere of Teilhard de Chardin. If we wish to understand knowledge at the practical level of those who have it, we must first ask why they come to know what they know. For much but not all of knowledge, this is to ask what good it does them to know it, which is in some rough sense the sum of the costs and benefits of their knowing it.

Of course, the assessment of costs and benefits depends itself on knowledge that must be subjected to the economic test. This means, of course, that the costs and benefits may be neither precise nor confidently measured. Rather, they are partly, even largely, guessed. In particular, one cannot typically know the full costs or benefits of trying out a new bit of knowledge, although one might have little choice but either to try it or not to try it, on the strength of a poorly grounded estimate of those costs and benefits. Note that, although the pragmatic or street-level account of knowledge is an economic theory, it is not the sort of economic theory that presumes full knowledge, as in the economists' rational expectations theory or much of game theory. And it is not merely about the costs of information, as it is in some economic accounts of, for example, George Stigler (1961). It is economic in the sense that the value of any bit of knowledge is how it would matter to us and our behavior, with consequences broadly defined to include the full range of costs and benefits of coming to know that bit of knowledge and of putting it to use.

The general tenor of an economic theory of knowledge is pragmatic. The central import of pragmatism in this context is roughly a combination of the following visions:

1. Wittgenstein's (1969) view that there are no foundations for much of knowledge;
2. Dewey's ([1920] 1948, 161; see also [1929] 1984) view that there is no final or best state of affairs in personal and social life, so that choices are not about reaching a maximum of some kind but about making things better; and
3. The contemporary view that much of an individual's knowledge is socially generated.

The first two of these visions are presumably clear enough without further elaboration. It might be hard to be at ease with these views, but their content is not unclear. The third vision is similarly simple enough. But the meaning that the term "socially generated" often is taken to convey is that the very fact of the matter, and not merely our knowledge of it, is socially determined. For social facts, this can be true. For the account here, all that is wanted is that much of our knowledge is grounded in social systems. That is, an individual acting alone could not have established it. The contributions of many people and even whole industries were

required to establish much of our knowledge. This vision will be spelled out further below.

It may be useful to note some caveats on what follows. First, for the problem of any practical knowledge, one must focus on the computational limits of the human agent, and on the incentives that agent may have to acquire knowledge, given that there are severe limits of capacity and cost in the acquisition of knowledge. My account here is not driven by an especially articulated theory of mental capacities, but only by a more-or-less commonsense view of limits on such capacities. Often, the main limit is time; we do not have the time to do all that we would like to do. Given that there are severe limits, we must be concerned with the ways in which bits of knowledge are acquired, and with the reasons or incentives for acquiring any bit of knowledge.

Second, on the view that there is no generally privileged position from which to judge whether someone's beliefs are true, there is no clear general distinction between beliefs and knowledge. The standard philosophers' category of justified true beliefs does not play an important role in a pragmatic theory of ordinary knowledge. Many philosophers in epistemology seem to suppose that they own the rights to usage of the terms 'knowledge' and 'belief.' Such philosophers are vastly outnumbered by ordinary people who are not impressed by claims to trademark these terms. My account is of the knowledge and beliefs of these vast numbers, and I will generally use the two terms as they are used in the vernacular.

Finally, moral and factual knowledge may not generally be distinguishable for anyone who has not been impressed with the theoretical distinction between them. The economic theory advanced here suggests that there need be no such distinction for a particular individual, because action on either kind of knowledge can entail costs and benefits that affect the content of the knowledge. Hence, for example, people seem commonly to moralize mere conventions, such as the conventions of table manners or those of tastes, that are place- or time- or culture-specific. For conventional reasons, we may refer to some bits of knowledge as beliefs, as in the case of religious beliefs. But these may be subject to the same account as any other knowledge, and may be distinguished from other knowledge in no way other than their subject matter.

Calling a theory of knowledge 'economic' implies essentially that there are choices to be made, in the sense that I must trade off other possibilities in order to invest in better knowledge. Hence, getting better knowledge does not automatically trump remaining ignorant should we choose to put our effort into other things. There is a corrupted, vernacular sense of "economic" that is not relevant to the discussion here: *An economic theory is not a theory about prices or money*, but is only about choice in the face of constraints on resources, time, and so forth. Similarly, costs

and benefits need not be tabulated in monetary terms nor assigned prices in order for us to make compelling sense of them. There are costs to you of enjoying some wonderful bit of food that are independent of the monetary price to be paid for it. Those costs include effects on your health, your wakefulness this afternoon, and your appetite for dinner with a friend later, none of which you might ever price monetarily. *Every credible theory of general choice is at least partly an economic theory in the sense used here.*

On the economic theory, we come by our knowledge in at least three quite different ways:

1. Sometimes we deliberately seek knowledge. We go in quest of it to be able to put it to use because we suppose that it will have value for us. Some of the knowledge we might deliberately seek is very specific, as is knowledge of current mortgage rates when we are buying a home. Some of it may be very general, as is the knowledge we will gain from going to university or that our children will gain from attending school, all as part of a larger plan of education for general or specific purposes. It is a distinctive aspect of modernity that education is quite general up to a fairly advanced age. In industrial society, education is typically generic (Gellner 1983, 27).
2. Sometimes we merely happen onto knowledge. It rains on us while we are engaged in some other enterprise. In this sense, much knowledge is an opportunistic by-product. In this category we may include knowledge that is a consumption good, such as the knowledge you gain from the pleasure you take in reading a newspaper or magazine. You often may gain such knowledge because it gives you pleasure to do so, and not because you think it is likely to be of any practical value to you.
3. Sometimes knowledge is imposed on us. Many people think of their schooling or their religious training as having been imposed. Indeed, the very word “training” is a peculiar term with somewhat odious implications of imposition. This is a complex issue, because the imposition cannot be simply by diktat. The Ayatollahs have massively affected the beliefs of Iranians since 1979, but to explicate the ways that that has happened is not a simple matter.

All of the knowledge you now have when you face a decision may be knowledge gained in these ways *in the past*. It is therefore, in a meaningful sense, *happenstance* with respect to the decision you now face. Most of your knowledge is a residue of the past in some sense. You might, of course, seek out some bits of additional knowledge, such as the current mortgage rates, but the bulk of your knowledge, even virtually all of it, may be happenstance in this sense for most of the decisions you ever

make. Your choices of what to know at various times in the past plus your opportunistically gained knowledge determine what you know now when you have to decide something. We cannot usually speak of *our* knowledge as though to imply we all have it. Our knowledge is simply your, my, his, or her knowledge. (Even less can we sensibly speak, as many intuitionist moral philosophers do, of *our* moral intuitions, whose Venn diagram is likely to be an empty set.)

Consider an example of a relatively trivial bit of knowledge that is, as argued below, nevertheless extremely complex. In the book *On Certainty*, Wittgenstein (1969, §170) says Mont Blanc is 4,000 meters high. It is actually nearer 5,000 meters. A standard philosophical analysis of these claims would focus on whether there are grounds for believing, for example, that the height of Mont Blanc is nearly 5,000 meters (or 15,728 feet). It would focus on the *matter* of belief or knowledge in question. The economic theory of belief focuses on the individual believer, not on the matter of belief, and on the costs and benefits to the individual of coming to have various beliefs. In such a theory we cannot speak of the justification of a particular belief *tout court*; rather, we must speak of the justification of that belief by some person in the larger framework of that individual's knowledge and experience. For this we require a theory that focuses on the individual and on the ways the individual comes to know or believe relevant things, such as how trustworthy another person, a book, or an agency is as a supposed authority on some bit of knowledge.

The economic theory has numerous implications, some of which may seem perverse on a standard philosophical account of knowledge. For example, my current knowledge is path dependent: it depends on the life I have led up to now. Indeed, even what I would count as knowledge is largely path dependent. I may, of course, come to correct part of my past knowledge, but in doing so I may not be able or may not think to correct other knowledge that is strongly connected to it. Hence, path dependence might contribute to the incoherence of beliefs. Moreover, if one has any concern with the coherence of the set of one's beliefs, path dependence might commend the rejection of new knowledge more readily than of old knowledge, for the reason that it is mentally cheaper to question a bit of new knowledge than to jettison a lot of old knowledge with the consequence of then having to restructure the broken remainder of old knowledge. A new fact might not merely challenge a particular old fact but also much that is inferentially based on the old fact. Those inferences from the old fact might be hard to track down, because they might not be mentally keyed to the fact from which they were inferred.

Hence, we may have some incentive to be conservative in our knowledge commitments. At the extreme of the very risk averse who wish to conserve limited mental and other energy, some people might be utterly

blockheaded in refusing to consider revision of any of their knowledge. One who seldom does reconsider might have little experience of the potential value of reconsideration, and might therefore become almost entirely embedded in stubborn beliefs from a dreadfully limited past experience. This is virtually a definitive condition of very traditional societies (Lévy-Bruhl [1910] 1926), if not to some degree of all societies.

David Hume ([1739–40] 2000, Appendix, §1) states the scientific creed with perhaps excessive enthusiasm: “There is nothing I wou’d more willingly lay hold of, than an opportunity of confessing my errors; and shou’d esteem such a return to truth and reason to be more honourable than the most unerring judgment. A man, who is free from mistakes, can pretend to no praises, except from the justness of his understanding; But a man, who corrects his mistakes, shews at once the justness of his understanding, and the candour and ingenuity of his temper.” (He adds, a bit less modestly, that “I have not yet been so fortunate as to discover any very considerable mistakes in the reasonings deliver’d in the preceding volumes, except on one article. . . .”)⁴

Again, on the economic theory the blurring of moral and factual knowledge can make sense. The incentive for gaining knowledge may, however, typically favor gaining practical knowledge over moral knowledge, because occasions for gaining such knowledge may be more frequent than those for gaining moral knowledge. For most of us, solving pragmatic problems is daily more common in our lives than is resolving moral issues. Hence, we may become better at factual than at moral knowledge, and may tend to bias our actions toward our own interests not because we are especially amoral but merely because we can be surer of making good self-interested choices than of making good moral or altruistic choices. On a related point, a moral theorist who is not a moral skeptic may rightly be held accountable on her own theory to higher standards of moral conduct, for the simple reason that she knows better than most people, who are not moral specialists, can be expected to know.

The economic theory also has strong implications for differences between institutional and individual capacities. Hence, pragmatic and moral considerations recommend use of institutions in some contexts just because institutions have special capacities for mastery of some kinds of knowledge. In other contexts, individuals often have advantages over institutions. To fill out the economic account of individual knowledge, let us address several topics: the social generation of knowledge, knowledge from authority, the division of labor and individual knowledge, and the internalization of norms.

⁴Stephen J. Gould (1993, 452) says virtually the same thing: “The factual correction of error may be the most sublime event in intellectual life, the ultimate sign of our necessary obedience to a larger reality and our inability to construct the world according to our desires.”

THE SOCIAL GENERATION OF KNOWLEDGE

Much of the knowledge that you or I have is knowledge that no individual alone could have established. Rather, the knowledge is generated by a social system to which many individuals and, often, institutions contribute. Again, this is not to say that the knowledge is socially constructed in the sense that we determine the underlying reality, as is sometimes argued of social knowledge and as might be true of some knowledge. As in Hume's and Smith's view of the greater productivity of the division of labor,⁵ discussed further below, knowledge grows faster if you pursue some things and I pursue others, rather than having all of us pursue everything together. But division of labor in the production of knowledge requires acceptance on faith of what others have come to know, if the advantage of division is not to be lost.

As an example of socially generated knowledge, consider again the knowledge of the height of Mont Blanc. How does one come to know such numbers as Wittgenstein's 4,000 meters—or the more widely accepted figure that is nearer 5,000 meters? Satellite measurements once suggested that K2 reaches higher than Mount Everest (Krakauer 1990, chap. 9). We have thought for about a century that Everest is the highest mountain on earth. But we might have been wrong. Suppose we eventually find that all measures corroborate the greater height of K2. What did we have to get wrong to get the particular bit of knowledge that Everest is highest wrong? These are all bits of knowledge that are not merely pieces of a system of knowledge—they are knowledge that could only be generated by a *social* system. To be confident in such pieces of knowledge is to be confident in the system that generates them. Hence, we might be able to assess these facts—but not my knowledge of them—according to the procedures followed in measuring them.

In the measurement of the height of a mountain, various tools are used, including surveyor's instruments. The users of these tools might have little understanding of the technology required to manufacture them to the fine tolerances required for their giving accurate measurements.⁶ The users must correct for the distorting effects of temperature on their tools, and also for the curvature of the Earth in measuring a mountain from which they must, of necessity, stand at a distance, and even the distorting effects of gravity from the mountain itself. They typically will not themselves have determined that curvature and might not even know how to mea-

⁵Poor Hume ([1739–40] 2000, book 3, pt. 2, sec. 2, ¶3) maladroitly speaks of “the partition of employments” and therefore loses paternity for our contemporary phrasing to Smith.

⁶This is a difficult problem. See Lehrer 1995.

sure it. Indeed, many of these things—the various technologies for manufacturing the equipment, the correction factors to control for the effects of variant temperatures on the instruments, and the correction factor for the curvature of the Earth—may all be essentially socially generated in the sense that they are the aggregate of many experiences that no individual could possibly replicate. Everyone must take at least many of these facts on faith in assuming that they have been determined correctly by a large number of contributing authorities.

Consider an example. In 1999, a German team scaled Mt. Kilimanjaro in Tanzania to measure its height. They used several global positioning system (GPS) receivers, to time and triangulate signals from satellites (whose elevations are highly calibrated). Their measurements, to within a few centimeters, at 5,892 meters, shaved 3 meters off the previously established height, as measured from atmospheric pressure (“Random Samples,” *Science* 286, 15 October 1999: 401). The 1852 measurement that proclaimed Mt. Everest the world’s tallest was carried out by the British Great Trigonometrical Survey of India with traditional surveyors’ instruments. These are dramatically different technologies all put to the same purpose. An even more massive task was the measurement of the Earth itself (Danson 2006).

KNOWLEDGE FROM AUTHORITY

We are still left with the individual-level question: How did you and I come to have whatever knowledge or beliefs *we* have of the heights of Mont Blanc, K2, and Everest? Many of us followed the second device discussed above, coming to it as a by-product of doing other things. Or rather, we didn’t follow any device at all, we merely happened onto the knowledge at some point and in some way that we might not even recall. We largely took it on authority, such as the authority of an encyclopedia or a newspaper story. And of course the encyclopedia editor or news writer most likely took it on authority from a geographer or other supposed expert.

Note one of the peculiarities of our presumptive knowledge of the height of Everest. We first have to judge a particular authority, and then we infer the truth of the authority’s claim. Most of our knowledge has this structure—it depends on reliance on some authority. Only rarely do we genuinely investigate for ourselves. But some of our authorities are now superseded in our judgments—although the facts we got from them may still linger in our minds as though they were correct. This character of our knowledge is especially important in moral judgments, because our moral beliefs may largely have been determined by the impositions of past

authorities, some of whom we would discount heavily today. But it afflicts much of our factual knowledge as well.

Ideally, each bit of our knowledge would carry a subscript to weight the quality of the source of that bit. Although much of our knowledge does not come to us associated with such weights, some of it does. Indeed, our ordinary language includes qualifications on our claims: I know, surely, more than likely, I think, the *Times* said, I've heard, maybe. In Tzotzil, a Mayan language spoken by about 200,000 Guatemalan Indians, there is a clitic, “la,” that must be inserted in every statement that reports mere hearsay as opposed to what one has personally witnessed or experienced. This hearsay clitic is in the class of “evidentials,” which qualify the commitment of the speaker to the claim being made. In English one can make similar qualifications, but it is not obligatory to do so. In Tzotzil, it is obligatory (Judith Aissen, personal communication, 27 December 2005). In the present world, in which we rely almost entirely on other sources for our knowledge, we would have to include a hearsay clitic in almost every sentence. That is a pervasive characteristic of our knowledge.

Might we ever follow the first device discussed above—deliberately seeking knowledge—in improving our knowledge? At least someone might. For example, someone working with satellite data might have tried to get the heights of all the Himalayas right. But even this person would depend inescapably on accepting knowledge from others without testing it—for example, knowledge of the solid state physics and the laser technology that lie behind the electronic measurements of height. Hence, even this expert seeker after these particular facts must have relied in large part on happenstance and personally uninvestigated knowledge. That is inherent in depending on *the system of knowledge*, as opposed merely to bits of knowledge in isolation.

Indeed, in the end it is hard even to imagine a single bit of knowledge in many areas of our seeming competence so that the problem of socially generated systemic knowledge pervades our lives. I compose these sentences at a computer whose structure is the result of millions of inputs by uncounted individuals, some of them expert on slight pieces of it and others expert on other slight pieces. The whole to which their expertise contributes is surely beyond the real comprehension of any of them or anyone else. Yet even I, who understand little of my computer's working, can presumably rely on it to radically improve my life in certain respects. My knowledge of that computer is fundamentally pragmatically determined: it works—although it sometimes fails, either because of its hardware, its software, or my errors in using these. In this case, I can try out alternative equipment to decide intelligently that some other set of machines is reasonably suitable for my task. Moreover, I can depend on social selection to have produced relatively good devices, because social selection

through market and other forces puts alternative devices to much better test than I can, and because I can expect the incentives of many of the people involved in the social selection to fit my own interests. I can also depend on social selection to have produced a lot of information on these devices, so that I may avoid testing everything or starting completely blind in my tests.

The paths to most bits of my knowledge have long been forgotten and could not plausibly be reconstructed. The overwhelming bulk has been accepted on authority of some kind, often of a questionable kind and often of a completely forgotten kind. There is a compelling sense in which it is useful to us to rely on the authority of others, because it enables us to make better choices for ourselves than we could if we had to rely only on what truths we had demonstrated in some way entirely by ourselves. It is thus deeply in our interest to rely on others. In our early years, indeed, it is a matter of life and death to rely on others and their knowledge. Throughout our lives, we would be foolish even to reject much of the knowledge we have that is probably grounded in dubious authorities, knowledge that we might even *know* to be grounded in dubious authorities.

Most of us have been stopped short on occasion by the question: How do you know that? Often the point of the question is to elicit information out of curiosity about our experience or our capacity to remember odd facts. But sometimes the query is an implicit statement of doubt, as though it were really the question: How *could* you know that?⁷ In either case, the questioner would often be satisfied by such simple answers as, I read it in the newspaper, my mother was a dentist, or I studied history. The question is then commonly not further directed at the authorities implicated in these answers, as it perhaps often should be. Nor is it so often followed by the relevant question, Are you sure that's what was claimed? For most bits of knowledge, the buck stops long before the end of the chain of authority. We typically do not seek out a first knower of any fact.

Given our radically defective reliance on authorities that are essentially unknown to us in the moment we put most knowledge to use, we take inordinately much knowledge on only a little more than blind faith. Why? Presumably because it seems to work well enough. That is the sense in which it is at least a little more than blind faith. We sweep vast quantities of putative knowledge into the maw of our minds with hardly a second glance. And we do it because we need to have enormous lots of knowledge but cannot plausibly take the time to check out the authoritative sources from which we take it to be true knowledge. I have a friend who often asks me, How do you know that? She invariably means it in the

⁷This question could be phrased with the emphasis on any of its words, or on various combinations of them, to alter its sense.

curious sense, and she seems almost as invariably to accept what I say as true, even though often I honestly can no longer say how I know whatever it is that is at stake. She, I, and no doubt you live by such knowledge, and we typically live reasonably well by it.

THE DIVISION OF LABOR AND INDIVIDUAL KNOWLEDGE

The constructive division of labor gives value to our social product. The division of labor is commonly seen from the perspective of its value to the larger society: it reduces costs and increases production. In these functions, it works in at least three ways. One, it allows you to specialize in the use of your time so that you work more efficiently with, say, a single tool rather than by switching from one tool to another as you perform the many very different tasks in producing a particular kind of object. Two, it allows you to specialize in the development of your talents to do some one thing especially well. And, finally, it allows us collectively to select from among us the one most capable of doing some difficult task: a Maria Callas can be selected to sing, a Michael Jordan to play basketball, and an Adam Smith to study the division of labor and its implications.

But the division of labor has another great value for an individual: it reduces the investment that the individual must make in understanding the world. You can generally rely on specialists who have already mastered parts of that world at least as well as you could expect to do. Those specialists might often be wrong, but they might still, on average, be less wrong than you would be if you had to rely only on your own wisdom and experience. Just as you rely on others to produce cars that work, so also you rely on others more generally to produce knowledge that works. In the social generation of knowledge, *we engage in division of labor—and this implies reliance on authority or testimony* (Coady 1992; Audi 1997; Adler 2002, chapter 5).

How does the individual detect error on the part of knowledge specialists? Primarily by having it pointed out by others. You learn that the ignition system on your car is prone to catching fire not by experience, testing, or deliberate personal investment in finding out the fact, but by reading of the problem in the newspaper. De facto, you rely on someone who has invested in finding out the facts of the matter, or in collating what many others have found out. Indeed, even if you had the experience of a fire in your car's ignition, you might still not know enough to judge the cause of or the responsibility for it. Such knowledge can come only from a larger canvass of the problems of that car's design and experience. You are spared the need to invest very substantially to learn either the fact or the nature of the problem, because you get the knowledge more or less

by chance from a credible source. That source might be very different in qualifications and in the incentives it faces than the company that made the cars, or than the individuals who own the cars; and that source might also be expected to have specialized in collecting the relevant knowledge. Competition in the production of knowledge is as valuable for productivity as competition in the production of goods is.

If we are in a traditional or small society, the prospects of our gaining insight into errors in the society's knowledge may be relatively poor. There may seldom be alternative sources of information beyond our collective social knowledge, configured more or less as it has been received from the past (Lévy-Bruhl [1910] 1926). The apparent conservatism of such societies need not, however, be a matter of psychological disposition but merely of lack of opportunity to learn on the cheap that their ways are not as effective as they might be. One of the great advantages of large, plural societies is that they offer up criticism of practices and knowledge very much on the cheap. The mass of criticism might often be disconcerting, but it is also very often enormously useful. In such a society, on many issues we may come to have no standard for what is normal, and we may be spared the normative challenges of people who might otherwise attempt to hold us to the "normal."

Appeal to authority is just a part of the division of labor for the creation and judgment of knowledge. We benefit from having specialists to assess auto safety, weather, and the truth of various matters. It is only an extension of normal reasoning to let specialists assess religious matters and moral matters of right and wrong, as will be argued in later chapters.

THE INTERNALIZATION OF NORMS

One of the most remarkable moves in ordinary understandings of the world is the internalization of norms, meaning adopting them to such an extent as to permit one to act on them without need of external sanction to motivate the action. But there are two forms of internalization of norms that are against one's interest. One is religious sanctions (exercised by the self) and the other is hierarchical social norms, such as a caste system, sexism, or racism, which some people internalize deeply enough to judge themselves inferior. Both kinds of norms have substantial social backing or enforcement, but they also seem to have internal reinforcement, which *prima facie* seems irrational. I strive charitably to show that the internalization in these cases can be arguably rational.

Many scholars think that the internalization of norms is an incoherent idea. Others think it is a real phenomenon, but one that requires complex and somewhat esoteric explanation as a product of psychoanalytic

processes that are not directly observable (Scott 1971). From the perspective of ordinary knowledge, however, it appears to be a fairly straightforward, even simple, phenomenon. First, let us start with the recognition that we can internalize lots of factual knowledge, and that this capacity or phenomenon is not unlike internalizing norms. Moreover, others can confirm that you have internalized such knowledge, so there is no insuperable difficulty in showing that the knowledge is in the black-box brain, even though we cannot directly observe the contents of that brain.

We commonly think of norms as action-guiding, as though this were a distinctive feature of them. But objective facts can also be action-guiding. I know that turning the key in the ignition will start my car. If I want to start the car, I turn the key. The more important characteristic of norms is that they are often motivating. What is of special interest to us in the internalization of norms is this fact that they are motivating.

Recall Ryle's distinction between two kinds of knowledge: knowing how and knowing that. These two may overlap. Suppose you have mastered the piano, or the technique for doing anything at all. The rest of us can know that you have internalized that knowledge from the way in which you put it to use in relevant contexts, when using it well would benefit you in some way. If one plays the piano, that might be almost entirely in the category of knowing how. But a Mozart, who can put musical theory to work, might also use a lot of knowledge-that when he plays.

Similarly, we might even say you have internalized knowledge of arithmetic, although some might insist that this is somehow hard-wired in the logic of the brain. Clearly, however, some people are very good at arithmetic, mathematics, and other logical fields and others are so poor at these things that they can hardly go beyond what they have learned, as in the rote learning of the simplest multiplication tables. The former people have come to do arithmetic with great ease and seeming natural mastery as a form of knowledge-how. The latter people, by contrast, seem to have done little more than internalize simple arithmetic through the rote learning of early schooling; they have only knowledge-that, for example that 2 times 3 is 6 and so forth.

The mastery of knowledge and techniques for doing things in the world is of course not the internalization of norms in the sense of moral norms that can guide behavior. But note that we smear the boundaries between normative beliefs that cannot be tested in the world and quasi-technological or scientific beliefs that *can* be tested. For example, many people think that their religious beliefs are matters of objective fact. We can therefore internalize such normative beliefs just as we internalize objective beliefs. They may in fact seem objective to many of us. Moreover, they may be reinforced by the reactions of other people to our behavior when those reactions either conform or fail to conform to these norms. This differs

from reinforcement of the belief that something like gravity is at work in our physical lives, but it is reinforcement nevertheless.

Moreover, we seldom attempt to test any of our beliefs, but merely accept them and sometimes act on them. Acting on a false belief that has objective consequences may de facto test the belief for us, and we might be given sufficient evidence to reject that belief. It is a peculiar characteristic of normative beliefs that they commonly do have consequences for us when we act on them, but are often held to have further consequences that cannot yet be seen. An extreme version of this is religious belief that one will be rewarded in an afterlife.

So far, we have established that internalization is a plausible and indeed likely commonplace in our lives, and that there is ample empirical evidence that many facts and techniques are internalized, evidence sufficient for others to know that you have internalized some fact or technique. Even some more or less normative beliefs may be internalized. But the hard issue to address here is whether we can internalize beliefs that are not in our interest. These include moral beliefs that lead us to act in ways that are not in our interest. They also include inferiority beliefs, such as beliefs by the proletariat that they are inferior to the higher classes, or beliefs by those in a low caste, by women, by an ethnic minority, or by many other groups that they are inferior in status to some other group.

Let us go back to arithmetic for a moment. We know that there are cultures in which arithmetic has not developed very far. Indeed, in one Amazonian culture, there is no counting beyond something roughly like a little and a lot (Everett 2005; Gelman and Gallistel 2004).⁸ In some of these cultures it is very hard to teach older people to do ordinary counting, although presumably the children could be taught readily.⁹ This means that even counting has to be taught in order for it to be internalized. Suppose you are one of the dullards who have to remind themselves of the multiplication tables to do some simple arithmetical calculations, and even have to rehearse the numbers in order to get counting right. You do this now not because you can independently believe that these moves work *but only because someone else has forced it into your head that these are the ways to do certain things*. At an age at which you are taught arithmetic, or even before that, you are also taught various moral principles and rules for behavior. In neither case—neither arithmetic nor moral norms—do you have independent proofs of your beliefs, as you might have of the relevance of your techniques for accomplishing various things, or of your knowledge of many objective facts.

⁸Some reports on the research findings say that the tribe counts only one, two. Daniel Everett (2005) claims they do not even have *these* two numbers.

⁹Hence, there may be an ethological constraint on the intellectual developments necessary for counting, just as there is for learning language (Lane 1976; Bretherton 1992, 762).

The moral principles that you are taught may come to you as facts no different in kind from other facts, such as that the moon goes through its various phases. You need not understand any of these bits of knowledge in any meaningful sense, beyond believing that they are true. For example, some of the things you believe might require that you go through various rituals in which you destroy perfectly good food, even though you are often hungry and could well use the food. Of course, you might sometimes have conflicting bits of knowledge and desires. In a state of severe hunger you might therefore cheat on the ritual you believe to be necessary for receiving more food. One might have no beliefs or norms that always trump all others or all conflicting desires.

It is not a big step now to believing, from having it drummed into one's head at an early age, that one is of an inferior caste, and that it is right that one should be badly treated as members of that caste are. Discovering that there are societies that are not organized into a hierarchy of castes might bring doubt to one's beliefs about the rightness of the caste system in which you live, but your ignorance of alternatives might otherwise characterize your life, and you might never doubt the rightness of how you are treated. If you explain all of this to your children, you might speak in moral or objective terms, you might invoke both, or you might not even draw a distinction between the two. But you and they might think that the ugly system of discrimination is fundamentally right, or is the way the world is naturally organized, just as it is organized to have monsoons and floods and droughts. Perhaps the most dramatic case of such internalization of norms of inferiority is the internalization of the so-called basic values of Stalin's Soviet Union, even to the point of accepting that it is reasonable that one should be horribly abused in the Gulag (Figs 2007).

That beliefs are internalized helps to make sense of behaviors when one has lost an earlier religious belief. Many lapsed religious believers still react to symbols and contexts that evoke earlier religious responses. If they have rejected their basic religious belief in, say, Catholic doctrine, there is no longer ground for the emotions formerly stirred by these symbols, such as the cross, or a consecrated and therefore supposedly holy wafer. Suppose you have lost your beliefs in the fundamental tenets of Catholicism. You might still not have hunted down in your mind and corrected or erased all of the minor associations of those earlier beliefs and practices. Perhaps they linger in the way memories of a lost lover or a dead friend linger in scattered places, and thus are not logically dependent on the greater beliefs in the falsity of Catholicism, the loss of the lover, or the grievous loss of the friend. Hence, you might, for example, go on a picnic where you had once gone with the lost lover or friend and find that the experience has evoked positive memories that then bring renewed sadness at the loss, even though you might well have got over the loss be-

fore the moment of the picnic. In any case, the force of the associations of the past religious beliefs makes sense now only if those beliefs and certain consequences entailed by them (if they are true) had been internalized in the mind in ways that make them available for evocation now.

STANDARD PHILOSOPHICAL THEORIES OF KNOWLEDGE

The philosophy of knowledge, or epistemology, is a highly developed inquiry. Much of it focuses on particular beliefs or types of belief and the criteria for truth or “justified true belief.” But the knowledge of people on the street bears little relation to the theories of knowledge of those in ivory towers. Ivory tower epistemology is about justification, about the truth conditions for counting some putative fact as known, rather than merely falsely believed. There are varied ways in which standard philosophical epistemology is divided into particular theories. Let us consider one of these, due to Keith Lehrer.

Lehrer (1990, 9 and 12; see also Kitcher 1994, 113) accepts the following conditionals for a theory of knowledge. If S knows that p, then it is true that p. And, if S knows that p, then S is completely justified in accepting that p. The odd quality of these conditionals is that they have no practical import. *There is no super knower who can attest to whether S meets them.* Without starting an infinite chain of queries, we can stop at one and ask, How does S know that it is true that p? Indeed, for much of life, perhaps arguably most of life, the only knower who might be concerned to attest whether S meets these conditionals is S, in which case the criteria are question-begging. Much of the theory of knowledge reads this way, as though somewhere out in the ether there were a super knower whose judgment is knowable to the epistemologist or even, illogically, the knower. But actual epistemologists are themselves merely other S’s who need have no special qualification to assess knowledge in general outside their own domains. Your local epistemologist is not a super knower, and cannot usefully judge your knowledge of most of the things that might affect what you must decide in your own life. The burden of finding competent judges of the bits of that knowledge is itself merely part of the burden of the formation of the knowledge—it is not a separate enterprise undertaken retrospectively or on the side, in order to judge the truth of the knowledge.

For a practical account, therefore, characterizations of and conditions for knowledge must finally be reduced to the perspective of S. Standard epistemology, however, is not subjective in this way but is essentially public. We can judge the knowledge of a *field* of science by the criteria of standard epistemology. It is not feasible that we judge the knowledge of an individual *person* by these criteria. Briefly consider five visions of epistemology to

assess their relevance to ordinary knowing by ordinary people: fundamentalism, coherentism, externalism, proceduralism, and communalism.¹⁰

(1) In fundamentalist epistemology, there are some things that are simply known directly. All other things are then known by their deductive fit with these directly known things. One need hardly show that this is an implausible description of the knowledge that you or I have, although one might suppose that some science could be characterized in this way. Suppose that an epistemologist could know what you know, and could therefore attempt to assess it. She might be able to show that parts of your knowledge do fit the model of deduction from fundamentals. But it would be utterly startling if she could do that generally, because your knowledge has been acquired in such a messy agglomeration of ways. The ways in which we acquire knowledge—most of it through accepting the views of manifold “authorities” or “experts” or merely sources—suggest that it is implausible that our knowledge can have happened to be deductively related to some fund of direct knowledge. We surely did not *adopt* much of our individual knowledge by personal deduction from any foundational beliefs. In any case, most of us might not have any foundational beliefs.

(2) In coherentist epistemology, each piece of possible knowledge is judged to be actual knowledge if it fits coherently with the rest of accepted knowledge. There might be a problem of circularity in determining some initial set of known facts against which to test any putative new fact. But at some point, one might settle for supposing that I just do have a coherent body of beliefs to which I can now fit new facts. However that might be, I am quite sure that my own knowledge is not coherent enough to pass a serious coherentist test.¹¹ Some of it—especially much of my professional academic knowledge—may be block-coherent. For example, my knowledge and understanding of the sociopolitical problem of collective action may all be coherent in the sense that any of the propositions I would make about it or within it fit with the other propositions I would make about it, both logically and empirically.

Much of my knowledge may seem to be unrelated to much other knowledge I have, and some distressingly large part of it may seem plausibly inconsistent with much of my other knowledge. We can readily justify living with such inconsistency, even at the base of our knowledge. Willingness to live with such inconsistency allows us to do many things other than getting our knowledge right or consistent.

¹⁰Lehrer (1990) presents a clear and careful account of the first three of these. The fourth, proceduralism, is a pragmatic theory of knowledge that fits the activities of many scientists and of many institutions. Although there are hints of it in earlier times, communalism is a relatively recent vision, except as it relates to religious and cultural knowledge.

¹¹A coherentist might rudely say that this is merely my problem—but it is inconceivable that anyone has genuinely coherent knowledge in general.

For example, consider the physicist who believes in the laws of physics and also in the basic principles of some religion, including an intervening god who manipulates outcomes. That physicist might have great difficulty supposing that these two bodies of knowledge fit together. Many people seem to suffer from superstitious reactions to many things while also formally believing that the causal connections necessary for the superstitious beliefs to work are false. (Stuart Vyse [1997] chronicles many standard superstitions.) Even the best of decision theorists might suffer from probabilistic instincts that they can formally demonstrate to be false. Again, the varieties of acquisitions of anyone's knowledge (as above, much of it from a huge number of authorities) make it implausible that it is all coherent.

Moreover, knowledge must be to some substantial extent spatially organized in the brain, and the different parts of the brain may not communicate well enough to keep all of the bits coherent. For example, consider a plausible account of why grief over the loss of a close loved one lasts as long as it sometimes does (Hardin 1988b: 181–82). My knowledge of the newly deceased person is recorded in many contexts—all those contexts in which she played an important part in my life. Her death now, however, does not instantly get recorded in all of those contexts, which therefore remain potentially live to jump to my attention with the freshness of richly lived experience that still has prospects of continuing. I have to read her death into each of these memories when they are evoked, and this is not unlike discovering her death anew on numerous occasions, over a long period of time, until finally her death informs almost all of my memories of her, and I cease encountering her alive and vibrant. If a fact of some other kind is similarly implicated in many aspects or moments of my life, I may similarly find it cropping up anew even after I have discovered it to be false. Moreover, I may find it particularly difficult (even unnatural) to purge not only the false fact but also the false inferences from it that inform my understanding of various things.

A deep problem for applying a coherentist theory of knowledge to an ordinary person is that the person is likely to have diverse blocks of knowledge, each of which might seem to be blockwise coherent, but pieces of any of which would not be coherent with other blocks. For example, your knowledge of different areas might come, respectively, from childhood teaching by long forgotten people, from your own somewhat random experience, or from instruction by masters in those areas. When you now acquire a new bit of putative knowledge, to which of these blocks are you to test it for coherence?

(3) In externalist epistemology, it is supposed that “what must be added to true belief to obtain knowledge is the appropriate connection” (Lehrer 1990, 153). That connection might be causal, or might simply be unknown so long as there is a connection. Actually, to conclude that someone

has knowledge under this epistemology requires a move different from the move the someone makes to have the knowledge. It requires something like the super knower above to conclude that the appropriate connection is there. Hence, externalist epistemology is simply irrelevant to the judgment of a subject about the subject's own knowledge. There is a version of externalist epistemology that is not so clearly afflicted with this problem, and I will discuss it separately as proceduralist epistemology.

(4) Proceduralist epistemology counts as true those facts that have been tested or established through following relevant procedures, such as experimental or other procedures of inquiry. This epistemology might well fit much of the knowledge of a well-developed, open, critical science. Historically important experiments in many sciences seem to meet routinized procedural requirements. For example, more than a century ago, the German bacteriologist Robert Koch proposed a procedural test of whether a disease is caused by a specific microbe. First, it must be possible to isolate the microbe from the diseased organism. Second, that microbe must be able to infect a healthy new host, in which it must cause the same disease. Then the microbe must be isolated from the newly infected host. For testing a human disease, for obvious reasons, the second step here usually requires an animal that is subject to the disease. In the demonstration that a microbe causes most cases of stomach ulcers, one of the proponents of the theory made himself the test host. He was massively criticized for that risky action, but his demonstration largely ended debate, and he eventually shared a Nobel Prize for demonstrating the bacterial role in causing ulcers and therefore an easy cure for them.¹²

Because AIDS has not been shown to afflict other animals, the Koch test of the causal role of HIV is said not to have been carried out. Peter Duesberg and others have therefore argued that AIDS is not an instance of retroviral infection from HIV but depends on several other factors, including lifestyle (Duesberg 1988, 1996; see, further, Horton 1996). Unfortunately, three lab workers were accidentally infected with HIV from a cloned strain (that was therefore free of AIDS infection), and they went on to develop symptoms of AIDS. Their grim misfortune completed the Koch test of the role of HIV in causing AIDS (Cohen 1994, 1647).¹³ One

¹²More than 90 percent of peptic ulcers are caused by a bacterial infection (*Helicobacter pylori*) and are therefore almost instantly curable with antibiotics. Australians Barry J. Marshall and J. Robin Warren received the Nobel Prize in physiology or medicine in 2005 after earlier enduring scorn from established medical authorities who dismissed their claim. As a major side benefit, their work may also reduce the incidence of stomach cancer, the second worst killer among cancers worldwide. Press Release from the Nobel Assembly at Karolinska Institutet, 3 October 2005. URL: <http://nobelprize.org/medicine/laureates/2005/press.html>.

¹³This is a fundamentally important finding, because by 1998 AIDS had become the fourth leading cause of death in the world (*Science* 284: 14 May 1999, p, 1101). Unlike some of the killers, it *tends* to strike people in their most active years.

might still insist that the Koch test is inadequate for establishing a causal relation here.

In the realm of social choice, we often expect institutions to apply standard procedures for reaching conclusions, as for example we require the justice system to follow standard procedures in reaching conclusions about whom to penalize for apparent crimes. The knowledge that courts “discover” is in part procedurally determined. But it is implausible to the point of being laughable to suppose that the knowledge of an ordinary individual can substantially be accounted for as the result of following some reasonable set or even sets of procedures.

(5) In addition to these four standard epistemologies, there is another that is currently developing. It is focused centrally on the social construction of knowledge, rather than on the knowledge itself. Hence, it is similar to the other epistemologies in its focus on public knowledge, although it has a strong subjective element. This epistemology can be called *communalism*. According to communalism, knowledge is what is grounded in a particular community of knowers, as in Thomas Kuhn’s ([1962] 1996) account of scientific knowledge for a particular research paradigm, in David Hull’s (1988) account of science itself as an evolutionary process, or in the contemporary literature on communitarianism, in which all of a person’s knowledge might be supposed to depend on her community (for fuller discussion, see Hardin 1995, chapter 7).

On some readings, communalism might be taken as a variant of externalism or, especially, of proceduralism. On alternative readings, it is a full-scale rejection of externalism, as in some of the writings of Richard Rorty ([1985] 1991). In the strongest positions, communalists deny the meaning of truth as a standard independent of what a particular community knows. They do not merely deny the pragmatic relevance of such a standard to a given community; rather, they assert its utter meaninglessness and inaccessibility. At the most extreme, some postmodernists evidently hold that all knowledge is merely the conventional product of communities, that there is no genuinely general truth independent of communities and ideologies (for criticisms, see contributions to Koertge 1998). In the hoax of Alan Sokal (1996a; 1996b: 62–64), the value of pi (the ratio of the circumference of a circle to its diameter) was asserted to be a social construct rather than a mathematical constant, and his deliberately silly assertion was honored with publication in *Social Text*, a prestigious postmodern journal.¹⁴

For present purposes, the problem with all of these epistemologies is that they are not applicable to the way we acquire most of our knowledge. *They do what they are intended to do*, which is to assess claims of

¹⁴See the collection of contributions to this debate: *The Sokal Hoax*, edited by the Editors of *Lingua Franca* 2000.

truth, but they do not explain our ordinary knowledge, as I wish to do. The pragmatic theory of individual knowledge has closest affinities with recent social epistemology (see, for example, contributions to Schmitt 1994 and to *Synthese* vol. 73, 1987; and Goldman [1992] 1993). Social epistemologists insist that much of even scientific knowledge depends on testimony (Coady 1992; Foley 1994). Without much if any investigation, scientists rely on the findings—the testimony—of others. Social epistemologists wish to establish criteria according to which testimony can yield justified true beliefs. Ironically, they reverse the move of Descartes, Locke, et al., who wanted to break the destructive hold of the supposedly authoritative views of Aristotle on science.

One might suppose that usefulness and how we discover our knowledge are merely prior questions, and that the usual questions of epistemology follow once we conclude that some bit of knowledge is useful. There are at least three problems with this view. First, people are typically not concerned with the central point of standard epistemology, which is justification. For the bulk of learned knowledge, people simply accept its truth; they do not attempt to establish it or test it. In ordinary affairs we might occasionally wish to check to be sure that some putative fact is actually true. For example, if I am told that some close associate or relative has violated my trust, I might strive to get the facts right. As Othello struggled with the (false) claim of Iago that Desdemona was deceiving him, he required substantial evidence before he accepted the claim. About the loyalty or trustworthiness of most people around him, the same Othello might have accepted relatively loosely grounded claims without substantially checking them, especially if they came from a source he trusted, as he mistakenly trusted Iago.

Second, as discussed further below, the criteria of standard philosophical epistemology are what could be called public criteria, rather than the personal criteria that the actual knower might apply. They are relatively esoteric criteria that the typical person does not command, and they are recent inventions that could not have governed, say, Aristotle's assessments of truth. They are public the way the contemporary laws of physics are public, and they are esoteric in the way that these laws, which a very tiny number of people know at all, are esoteric. They are decidedly not the criteria of ordinary knowledge, which cannot plausibly be said to meet them. Hence, it is of only academic interest that a philosophical epistemologist might say that John and Mary Doe's "knowledge" is not philosophically sound enough to be knowledge, even though their knowledge governs the choices they make.

And third, the criteria that a person applies to judging a bit of knowledge are not necessarily criteria for truth, but merely and genuinely criteria for usefulness. Individuals cannot typically know much about whether

their knowledge is true, but they can commonly grasp whether it is useful, whether it works in some sense. For knowledge, use is what matters to most of us most of the time, apart from when we might be doing science or philosophy. Truth might very often happen to be part of what makes some bit of knowledge useful, but it is not the same as usefulness and need not be a part of what makes for usefulness. Some bit of knowledge that would be false by any standard epistemological criteria—as arguably most of a typical person’s knowledge is—might nevertheless be very useful. Indeed, much of the knowledge we accept and act on is merely satisficing knowledge, that is, good enough.¹⁵ It is adequate for the relevant tasks and it need not be useful for us to invest the effort needed to improve the knowledge. One might say that such knowledge is “true enough” for its purpose, but that is an odd locution. In any case, it is instructively a locution misfit in the realm of standard philosophical accounts of knowledge.

Ideal knowledge is of interest only academically, not practically. When we turn to practical knowledge, we must inherently be interested in knowledge that is put to use by someone, such as you or me. But, then, what is knowledge for me might well not be knowledge for you. This is conspicuously true for knowledge of any kind, even philosophically sound knowledge, that I have acquired that you have not acquired. But the point here is even more far reaching than this. I might even know a physics that is different from your physics, and I might act on it while you act on yours. Arguably more important in our lives than different knowledge of physics is different knowledge of social causation and of social facts, both of which play far more important conscious roles in our actual choices from day to day than does any articulate knowledge of physics.

CONCLUDING REMARKS

In standard philosophical epistemology, it would commonly be thought incoherent to speak of my mistaken knowledge. Knowledge is, in some epistemologies, “justified true belief.” If I am mistaken in my belief, then I most likely lack justification for the belief. Hence, it is not knowledge. And in any case, the category of justified true belief is a category of somehow public knowledge, not personal knowledge. For most of us, most of the time, there may be no ground for claiming in general that our knowledge is philosophically justified in any such sense. There is commonly only a story to be told of how we have come to have our beliefs. There is

¹⁵This term is due to Herbert Simon (1955).

therefore little or no point, for present purposes, in distinguishing between belief and knowledge. Typically, at the street level of ordinary people, who are not philosophical epistemologists, the term “belief” is commonly used when the substance of the knowledge is a particular kind, such as religious knowledge. There is often no other systematic difference in degree of confidence in knowing those things that are labeled as knowledge and those that are labeled as belief. Indeed, people with strong religious convictions commonly claim to know the truth of the things they believe religiously far more confidently than the truth of many simple, objective things they might also claim to know. It is true that we sometimes use the term ‘belief’ to allow for doubt, as when we say, “I believe that’s the way it happened, but I might be wrong.” But this hedge applies to virtually all our knowledge.

Standard philosophy of knowledge is concerned with justification, that is, justification of any claim that some piece of putative knowledge is actually true. The ordinary person’s economic theory is economic in the sense that it is not generally about justification but about usefulness. It follows Dewey’s ([1920] 1948, 163) “pragmatic rule”: the meaning of an idea is its consequences. In essence, the theory here applies this rule to the idea of knowledge, with consequences broadly defined to include the full costs and benefits of coming to know and using knowledge.

Trudy Govier (1997, 51–76) argues that our knowledge therefore depends on trust. It might be better to say that it depends on the trustworthiness of our authoritative sources, although even this is saying too much. Very little of our knowledge seems likely to depend on anything vaguely like an ordinary trust relationship. I personally know none of the authoritative sources for much of what I would think is my knowledge in many areas. It is not so much that I take that knowledge on trust as that I have little choice but to take it. If I do not take it, I will be virtually catatonic. I am quite confident that much of what I think I know is false, but still I rely on what I know to get me through life, because I have to.

The concern of an economic theory of knowledge is, again, the problem of *a choosing agent’s knowledge* as opposed to a theorist’s or critic’s knowledge. We commonly have no recourse to a super knower when we are making life decisions, whether minor or major. Occasionally, we may turn to supposed experts for assistance, although it is we who must judge whether someone or some source is a relevant expert (see further, chapter 2 on science). Tony Coady (1994) goes so far as to ask how the lay person can distinguish between communication from an expert and communication from a spirit in the spirit world.

We ground our lives in putative but philosophically ungrounded knowledge. The task of clearing out the Augean stables in our heads would dominate our lives for decades and would get in the way of our living, not to speak of living well. Actually completing the task might well

disable us almost entirely. Wittgenstein (1969: §344) says: “My life consists in my being content to accept many things.” These are de facto *my* foundations for going on to consider other things.¹⁶

In the end, if we did attempt to clear the fouled and cluttered stables in our heads in a Herculean commitment to some epistemology, we would plausibly have so little left that we could scarcely have grounds on which to choose at all in most contexts. If we finally cleared out everything that did not fit one of the standard epistemologies, we would be very nearly catatonically incapacitated. Any of those epistemologies would arguably be pragmatically disastrous for most people. Since the principal value of knowing at all is to enable us to choose and live well, it therefore follows that we should—at least most of the time—pragmatically reject the demands of those epistemologies as deleterious to our prospects. Most of the time, we should join Samuel Johnson in not going in quest of knowledge.

Philosophers sensibly pursue the enterprise of standard epistemology, the enterprise of the justification of knowledge, but ordinary people can rightly forgo it for most of their own lives. This is, of course, merely a claim from the economic theory of knowledge, with its pragmatic focus, and is only a practical, not a theoretical claim. That economic theory of knowledge would fail the truth-tests of all the standard epistemologies, and the relevant epistemologists would dismiss it in turn. Nevertheless, although there may be good reasons for doing or knowing philosophical epistemology, those reasons are not among the practical reasons that should inform the reasoning of ordinary people when they are making decisions about their lives. For much of knowledge, if we ask “Why know?” we surely must answer “Because it is useful.” But if that is our answer, it follows that we should assess such knowledge by its utility—which is to say, by an economic theory of knowledge.

A frequently cited legal standard of rationality says that “the reasonable observer is an informed citizen who is more knowledgeable than the average passerby” (*Modrovich v. Allegheny County*, p. 407). This already sets the standard above the average, and is therefore too demanding for most people. We would live much less fulfilled lives if we were far more demanding in the standards to which we subject the flow of knowledge into our minds. This is the vague, pragmatic test to which we subject putative knowledge. Moreover, it seems clearly to be the correct test for real life. That is to say, if we are commending a practice to someone else, we should rightly commend this practice to them. We would commend rigorously following any of the standard epistemologies only if we were foolish or mean-spirited. That is, in any case, not the point of philosophical epistemology.

¹⁶Indeed, “propositions of the form of empirical propositions, and not only propositions of logic, form the foundation of all operating with thoughts” (Wittgenstein 1969: §401).