# 1 Social Norms, Knowledge, and Philosophy

## A Social Norm for Testimony

We seem to have an informal rule that we should tell people something only if we know it. For example, you should tell your friend that the orchestra is playing this Saturday only if you know that it is. Some philosophers think that our having such a rule indicates something important about telling (or testimony), or about assertions more generally. This book argues that our having this rule indicates something important about what it is to know something.[1]

I think we do have a rule that you should tell people only what you know, and that in some ways it is rather like the rule that you should always dress appropriately. It is a complex social norm whose specific content – what counts as following the rule – depends on local culture and ways of doing things. What counts as appropriate dress – more specific norms about how to dress – varies not only with person, time of day, season, location, and activity, but also in different societies and across the years. Men in the United States in the 1920s were expected to wear hats outdoors in most public places, and were made to feel out of place if they didn't. Going hatless no longer counts as being inappropriately clad for adult males in the United States, but going shirtless does, on many occasions and in many places. Failure to be appropriately dressed exposes us to disapproving looks and comments but, except in a few cases, not formal legal sanctions. Most of us comply with the social norm requiring appropriate dress by complying with the more specific norms for our time, place, role, and activity (e.g., relaxing at home, recreation in public places, working in farming, education, a warehouse, a factory, medical services, tech businesses, traditional businesses) and our own cultural sub-classification (self-employed contractor, student, teacher, salesperson, engineer, lawyer, retiree, old, young, owner, management, labor). But many of us occasionally violate the

[1] Similar recent projects include Edward Craig's project in *Knowledge and the State of Nature* (1990), which will be discussed below, and Sinan Dogramaci's "reverse engineering" of epistemic evaluations (Dogramaci 2012, 2015a and 2015b), which focuses on the function of the concept of rationality, rather than the concept of knowledge. For more discussion of the purpose of epistemic evaluations, see Henderson and Greco (2015).

specific norms of dress for our society and our place in it, thereby also violating the general norm by dressing inappropriately. A few violate this norm frequently, mostly because they don't care enough about the sort of disapproval to which that exposes them to spend the time and effort to dress appropriately.

We also have a social norm, an unwritten rule, that we should testify that *p* only if we know that *p*. If someone asks a group of people "What time is it?" she expects to receive an answer from a member of the group who *knows* what time it is. If you've checked your watch or glanced at your cell phone recently, that would usually satisfy our specific norm for saying what time it is. Consulting an acceptable timepiece, or hearing from someone who has recently, is what counts as knowing what time it is, for many of us, at this cultural juncture. A thousand years ago perhaps knowing what hour of the day it was would require a glance at the sun or at the shadows it casts from familiar trees and buildings, or hearing recent testimony from someone who had seen those things. Our detailed norms for knowing the time have changed fairly recently to allow checking a cell phone in preference to looking at an old-fashioned pocket watch or other wind-up clock, and no doubt they will continue to change in similarly exciting ways in the future. Knowledge is required for answering questions such as "Where is Freddy today?" or "When does the department meeting start?" If someone answers "Freddy has a medical appointment," or "The meeting starts at four," but it then turns out that they didn't know, and should have realized that they didn't know, we think that they shouldn't have said what they did. In addition to thus thinking badly of them, we may give them a disapproving look, or in extreme cases make a (gently) disapproving remark. And of course they may anticipate our reaction, whether we have expressed it or not, and respond to it with "He told me last week that he had a medical appointment this afternoon" or "I must have seen an old meeting notice." These would be ways of extenuating or excusing non-knowledgeable testimony, by saying in effect that it was reasonable for the testifier to think she did know, that is, it was reasonable for her to think that she had complied with the relevant specific norm(s) for knowing that sort of thing.

Some philosophers doubt that we have a general norm for testimony, or for assertion. (Testimony is telling people things; the category of assertions also includes declarations that don't tell the audience anything, as in reminders of what everyone already knows, or affirmations of a creed, or re-statements of discussion points.) Others maintain that testimony or assertion doesn't require *knowledge*, exactly, but something else closely related to knowledge, such as truth, or justified belief.[2] Some hold that the norm is not, or at least is not

[2] For example, Paul Faulkner holds that we have a social norm requiring not knowledge as such, but truth telling (Faulkner 2011, 36–7).

primarily or most importantly, a social norm, but that it has some other normative character.[3] A *social norm* is an informal or unwritten rule, enforced, if at all, by approval and disapproval. I shall discuss the nature of social norms more as we proceed. But for now the examples I have been reviewing give us some reason to think that we do have a social norm for testimony and that it has some of the structure – of a general norm, and of more detailed specific norms that must be complied with, at a given place and time, to satisfy the general norm – that we noticed for the norm that requires wearing appropriate clothing. The general norm is satisfied by compliance with specific norms, and the specific norms involve detailed information most of us have about our culture, the topic, and ourselves.

## Could That Be Philosophy?

But what could a contingent fact about *testimony* (that we have an informal rule requiring us to testify only what we know), if it is indeed a fact, have to do with the larger concerns of the philosophical theory of knowledge? We prefer that people tell us what they know. We also prefer that they warn us when they don't have knowledge that *p*, if they are telling us that *p*, by prefacing what they say with something like "I think. . ." or "I'm not sure, but it seems to me that. . ." If we have this social norm, then presumably we have reasons for these preferences, reasons good enough to justify the small encouragements and discouragements needed to establish and maintain a social norm. But these small practical matters may seem a long way from philosophy.

It is natural to think that philosophy is and should be interested in a conception of knowledge that derives from important kinds of inquiry, including history, law, the various sciences, and perhaps even the intellectual projects of such thinkers as Descartes, Locke, and Kant. Should we think that any of these larger concerns will be much illuminated by reviewing our preferences for testimony, especially given that the examples of testimony with which we are apparently concerned are mostly about small practical matters, such as what time it is, or where our coworkers are? (Or which refrigerator contains the samples from yesterday's preparation, or whether the notes of the procedures as they were carried out are on this laptop, or whether philosopher X said that . . . in reply to. . .) Does such ordinary knowledge, we may wonder, count as the same kind of thing as the knowledge that scholars, historians, the courts,

[3] Timothy Williamson holds that knowledge is a constitutive norm of assertion: Being subject to evaluation according to such a norm is what constitutes a speech act an assertion (Williamson 1996, 2000). Tyler Burge holds that knowledge is a "natural norm" of belief, where natural norms are defined as adequate performance of a function given available mechanisms and resources, and are said not to depend on anyone's attitude toward such a norm (Burge 2010, 311).

scientists, and philosophers expend such lengthy and difficult efforts to obtain and to understand?

We must first learn what to count as knowledge by paying attention to uses of "know" and related terms in ordinary conversation, just as we learn how to use the other words that make up the central parts of our language. How else could children learn what knowledge is, except by finding out what their parents and other older speakers count as knowing or as not knowing? Language must begin, obviously, with matters that children can understand. But perhaps there is still some large difference between such ordinary uses of "know" and the uses of that word by historians and other scholars, lawyers, judges, jurors, witnesses in court, scientists, and medical doctors, when they are engaged in their important pursuits. Why should we think that a social norm learned in connection with small practical matters, and the low standards for being said to know that can be learned and applied by young children, have anything of interest to tell us about the concept of knowledge as it applies to our more significant adult concerns?

I shall argue that the social norm conception of knowledge helps us to understand several controversial topics in the philosophical theory of knowledge. Those topics are discussed in the following chapters: in Chapters 2–4, the nature of knowledge; in Chapters 5–6, the nature of adequate justification for belief, in Chapters 7–8 our justification for believing the testimony of others, in Chapter 9, whether we have enough, or the right kind of, control over our beliefs to make sense of norms for believing, and in Chapter 10, why we do and should prefer knowledge to mere true belief.

Each of these topics is approached through a conception of the purpose or function of classifying people as knowing, that it is to indicate and thereby encourage acceptable testimony and to discourage unacceptable testimony. But why should we think that this is the function served by the concept expressed by our word "know"?

### State-of-Nature Thought Experiments

It seems likely that we must already have a pretty good idea what function the word "know" serves, since it is an artifact that we commonly use, like a doorknob or a hammer. But the exact functions of some familiar artifacts are a little difficult to understand even for those who commonly use them: money, for example, or posts of photographs on Facebook. So it may be useful to try to make our understanding explicit.

One way to do that is to consider what else we value that would probably be missing if we didn't have the artifact in question. How would society be different if it lacked a concept of knowledge? Such thought experiments have been important in political and moral philosophy for a long time, at least

since Thomas Hobbes imagined a society without government in his *Leviathan* in 1651. It may be helpful to briefly review the outlines of that thought experiment as a model for the thought experiment of imagining a society without a concept of knowledge.[4]

Hobbes proposed to investigate the function of government by considering what human life would be like if we lacked a government. He asks us to imagine a substantial human population that has no rulers who have the power to compel obedience to their commands. Because human beings are prone to conflicts over scarce goods, such as food and water and land, and also to conflicts over social standing and mates, and because no one of us as an individual is so strong or so clever as to reliably win these conflicts without assistance from others, or so popular or persuasive as to receive such assistance reliably on a purely voluntary basis, such a society would live in constant conflict, Hobbes argued, and those who belong to it would be unable to cooperate in larger projects or accumulate goods over time. Life for people in such a society would be a "war of all against all" and their life would be "solitary, poor, nasty, brutish, and short" (*Leviathan*, chapter 13). The way to escape this miserable condition, he suggested, would be to combine with others to empower a sovereign. The group thus organized would compel each of their number to obey the commands of their sovereign, thereby reducing the frequency and severity of conflict and enabling members of the society to cooperate successfully.

There is a puzzle about how these imaginary people could cooperate to empower a sovereign if they were unable or unwilling to cooperate without one. But that problem is presumably merely an artifact of this way of considering the question of the purpose of government. There is no reason to think that any sizeable human society ever existed without some sort of leadership, however far that leadership may have fallen short of genuinely coercive powers. When we reflect on the behavior of other social mammals such as deer, wolves, or bonobos, it seems very likely that our hominid ancestors always had at least some features of an organized society, including de facto leaders for some purposes. Those features would have facilitated the development of stronger governments, as hunter-gatherer tribal societies grew larger and developed the tools needed for still larger organizations (agriculture, trade, record keeping, written contracts, etc.). We are asked to consider what human society would be like if there were no ruler(s), not because there is any real

[4] What follows is only a toy version of the argument Hobbes actually gives, which continues to exercise political philosophers and students of the history of philosophy. Although it is a toy version, it illustrates some key points about the sort of argument I want to consider for the theory of knowledge. For what little I understand of the political argument I owe a debt to the late Jean Hampton, and her book *Hobbes and the Social Contract Tradition* (Hampton 1986), on which I did some trivial work as a research assistant.

likelihood of such a state of affairs arising in human society (except perhaps briefly in "failed states"), but as a way to make clear to ourselves the main functions or purposes of government. What would a society without a government be like? If the suggestion is correct that it would lack the sorts of cooperation we ordinarily have, because human beings living without government would frequently engage in disruptive and destructive conflicts that are by comparison rare in any society that has a government, and if we can see how the law-making and -enforcing aspects of government reduce those problems to a manageable level, as in our present society, we may reasonably conclude that government functions (at least in part) to reduce conflicts and thereby to enable cooperation.

Even if there never was a sizeable human society that lacked the organized coercive power to reduce interpersonal conflict, we may still say that government evolved in order to fulfill that function and that it will continue to develop in ways that enable it to continue to fulfill that function. Similar points apply to the epistemic state of nature to be imagined below. A function of encouraging acceptable testimony, discerned in a thought experiment where one imagines a society without epistemic classifications – although perhaps no human society ever has lacked them – may still explain the past and present development of those epistemic classifications. This is a general point about the relation of functions to the evolution of ways of fulfilling those functions. We can say correctly that kidneys function to remove toxins from the blood, and illustrate this by pointing out that human beings without functioning kidneys would (and do) die from the collapse of other bodily systems caused by the accumulation of such toxins, without committing ourselves to a historical first kidney. The function helps explain the development of kidneys and many features of their operations, even if the gradualness of the development of these organs and the systems to which they belong makes it impossible to identify the first one.

Other questions about government are naturally considered by thinking about the state of nature as Hobbes describes it. If one had a free hand in designing a government, because there were no pre-existing traditions or institutions – as in the imagined state of nature – it would be natural to ask what sort of government would work best in fulfilling the function of discouraging conflict. A related question would be: When there are deaths or other incapacities among the existing rulers, how should governmental functions be transferred to others? Hobbes, under the influence of the English civil wars of the 1640s, argued for an absolute monarchy, with hereditary transmission, on the ground that attempts to share power would tend to break down as conflicts arose between roughly equal co-rulers. The "war of all against all" would also occur within the smaller ruling body, he thought. Satisfactory experience of other forms of government persuades most of us that he was

mistaken on this point. There are points about the function of improving testimony that will be analogous to these points about government. I will argue below for some claims about knowledge and justification on the ground that they make more sense than possible alternatives as producing better testimony.

Moral questions about government are also considered in connection with the state of nature. Is it morally permissible for there to be a government that compels obedience to its laws? Which persons is it permissible for a government to thus compel? What moral obligations does one have to an existing government? For the philosophical study of government these are central questions, and the thought experiment of a society gradually leaving the state of nature is an illuminating way to consider them. However, in this book I won't be considering the interesting and important moral questions about knowledge.[5]

### Empirical Philosophy?

Let's consider again how the state-of-nature thought experiment supports an account of the function of government. The problem is posed by apparently empirical claims about actual human nature: that we are quarrelsome, and too equal to one another for anyone to win reliably in these quarrels, except by organizing a government. When asked to imagine a society without government, we tend to think of how people often act in the absence of policemen, not so much about how they would act in the absence of other services provided by our own governments, such as schools, roads, old age pensions, public medical clinics, parks, and libraries. Anarchists, whether religious, libertarian, or socialist, disagree with Hobbes's conclusion that a tolerable society requires coercive government, presumably because they disagree about the quarrelsomeness of humanity or about the relative effectiveness of coercion in improving our behavior as compared to other means. These are disagreements about the causes, frequency, and consequences of contingently occurring human behaviors, and it seems they can only be resolved by experience.

So why is the state-of-nature thought experiment about government offered as philosophy? One popular, if rather extreme, conception of philosophy holds that it is, in principle and in its best practice, an *a priori* discipline. In its early

[5] See Fricker (2007). My discussions of social norms and their influence on testimony are certainly missing something important, ignoring the actualities of testimonial injustice and unjust exercises of asymmetrical epistemic power of other sorts, but one can't discuss every issue at once. My work indulges in an idealization of our social influences on one another, as if we lived in an approximately fair society and as if the influences we exert, as it were in the name of knowledge, are not corrupted by old and deep injustices. I hope the discussions below are compatible with recognizing those problems, at least as complicating influences that have to be worked into the overall theoretical picture, as well as resisted in practice when and as we can.

development it is a kind of hazy and unconvincing would-be mathematics, but its various studies sometimes evolve into branches of genuine mathematics as their concepts become clearer. That has happened to parts of logic in the past century and a half, and it may now be happening for some formal aspects of epistemology and decision theory. Shouldn't philosophy then seek to develop along these lines in its other, still mostly non-mathematical, branches, leaving empirical questions to those who have already developed special expertise in the appropriate empirical methods: psychologists, sociologists, anthropologists, economists?

I support the positive aim of developing philosophy along *a priori* lines, but I don't see any benefit in restricting philosophers' other interests. There is every reason to pursue the conceptual aspects of philosophy in an *a priori* manner, and to make those bits of philosophy as properly mathematical as the subjects will allow. But philosophy is not and never has been an entirely *a priori* discipline, despite the efforts of some mid-twentieth-century historians of philosophy to so delineate it, carving off the more narrowly conceptual bits as the only genuinely philosophical parts of Aristotle or Hobbes or Hume. Philosophers are interested in the state-of-nature speculations not because they fit comfortably into the allegedly *a priori* expertise of our discipline, but rather because they seem likely to help with the big-picture theorizing that also interests us. Philosophy has always taken as one of its projects the fitting together of various large pieces of human thought, in whatever condition they have reached, as well as we can at present. It is reasonable to hope that the more successful such efforts illuminate their subjects in a way that contributes to genuine understanding of the whole. I don't think there's any convincing reason to think that other disciplines, as currently practiced, are so much better prepared to perform this task than philosophy is that we should simply abandon it to them. Nor is there any reason to ask practitioners of other disciplines to defer to philosophers as such in evaluating the results of attempts to construct big pictures. This is evidently a multi-disciplinary project, of interest to a great variety of inquirers, and one that will benefit from skills and knowledge that are not the property primarily of any single discipline. But it is one for which philosophers do have some relevant inclinations and skills.

Another aspect of the state-of-nature thought experiment regarding government that makes it seem traditionally philosophical is that it invokes ordinary empirical knowledge that we tend to leave tacit much of the time. Anyone with playground experience knows that fights among children are more frequent and tend to last longer in the absence of adults who are willing to interfere. Adults too are more likely to fight in the absence of sufficient law enforcement, which requires security arrangements for places of public amusement and commerce. It may be partly this ordinary experience that underlies our inclination (if we have it) to agree with Hobbes that life without government would

be "a war of all against all." The state-of-nature thought experiment makes this experiential knowledge clearer to us, so that we can understand from it why we need government. Much of philosophy, as Wittgenstein famously observed, consists in assembling reminders of what we already know so as to reduce our tendency to find our own activities puzzling (Wittgenstein 1958, 127–9). The thought experiment imagining a society that lacks a concept of knowledge reminds us of what we already know through experience.

That poses a problem for the philosopher who hopes for some notice of her own assemblies of reminders. Reminders of what we already know will naturally be greeted with the comment that we already knew that. And, unless the philosopher is mistaken, we did know it. Other philosophers will claim that they presupposed what they are now being reminded of, rather than having merely overlooked it, with the implication that the reminder isn't really necessary. Still, it is often useful to make the familiar more explicit and thereby clearer, and it is surprisingly difficult to do it adequately. I think the social norm function of the concept of knowledge has been thus largely overlooked, and not merely presupposed, in philosophical discussions of knowledge. As I shall try to show, thinking about it more explicitly has useful explanatory results.

## Experimental Philosophy

That brings us to the recent development of experimental philosophy, which takes seriously the task of evaluating philosophical claims that we already know certain things. The experiments so far consist mostly of opinion polling, and the first target has been a method of philosophical investigation that goes back at least to Socrates. The reader will recall that, as Socrates' activities are represented in Plato's dialogues, and in Socrates' speech at his trial, he commonly challenged his conversational partners to give an account of some important quality such as justice or piety. If they responded, as apparently they often did, with a few ordinary examples, he would then explain that he wished to be told what those examples had in common. Their attempts to give a general informative account of what the examples had in common were then often criticized by citing other examples that satisfied the account, but didn't seem to be examples of justice (or piety, etc.), or by citing examples of the target quality that didn't satisfy the attempted account. It is the reactions to these sorts of alleged counterexamples to philosophical accounts of important qualities that have been the focus of many of the recent experimental philosophy questionnaires.

The immediate aim is to find out whether people who don't have philosophical theories to defend, or training in philosophy, agree that the cases offered by philosophers in criticism of the accounts are examples, or non-examples, of

the investigated quality X. Philosophical accounts of knowledge have been prominent in these early attempts at experimental philosophy.[6] Knowledge of a proposition *p*, some philosophers have said, is justified true belief in that proposition.[7] To put it a little more formally:

S knows that *p* if and only if:

(1) *p* is true
(2) S believes that *p*
(3) S is justified in believing that *p*.

Edmund Gettier offered two counterexamples to the sufficiency of these qualities for knowledge (Gettier 1963). His examples of justified true belief that are apparently not examples of knowledge are a bit awkward to offer to the public in opinion polls, because they depend on some bits of reasoning that may be confusing to those who haven't studied formal logic. I describe instead a somewhat more natural "Gettier" case, of a sort originally proposed by Carl Ginet and reported in Goldman (1976).

One day Harry is driving in the country and, seeing a familiar looking kind of structure a little distance from the road, he believes that it is a barn. It is a barn, so his belief is true. But, although he hasn't heard, he happens to be in fake barn country, where many barn façades have been constructed that deceive even careful observers seeing them from the highway into thinking they are observing real barns. He has been thus deceived several times already in the past thirty minutes, although he is now seeing a real barn.

After a few questions to screen out the respondents who are not paying attention or who are likely to misunderstand the story for other reasons, the experimental philosophy questionnaires ask whether Harry knows that it (the structure he is looking at) is a barn.

The philosophically encouraged answer is that he doesn't know it. The conclusion philosophers have been reaching from this case, and many related cases, is that since Harry has justified true belief that it is a barn, but he evidently doesn't know that it is a barn, justified true belief is in general not sufficient for knowledge. A large and difficult literature has grown up since the publication of Gettier's paper, consisting of attempts to correct this account of knowledge, usually by proposing a fourth necessary condition for knowledge, and sometimes by replacing the third condition too. These attempts to define

[6] See Machery (2015) for a recent and carefully conducted example of such studies in epistemology.
[7] This philosophical account of knowledge seems to be very old. In Plato's *Theaetetus*, Socrates is presented as discussing an account of knowledge as "true judgment with an account." See Burnyeat (1990), 338 ff. Gettier cited, as then recent examples of philosophers who held the view, Chisholm (1957) and Ayer (1956).