

Chapter 1

Modes of Knowledge

SYNOPSIS

- *The traditional idea that knowledge is true justified belief requires both clarification and qualification.*
- *Propositional knowledge is subject to various distinctions, among which that between explicit and inferential knowledge and that between occurrent and dispositional knowledge are particularly prominent.*
- *When “knowledge” is construed as inferentially accessible knowledge, it becomes possible to construct what can plausibly be seen as a “logic” of knowledge.*

IS KNOWLEDGE TRUE JUSTIFIED BELIEF?

It is something of an oversimplification to say the knowledge involves belief. For one thing, believing is sometimes contrasted with knowing as a somewhat weaker cousin. (“I don’t just *believe* that, I *know* it.”) And there are other contrast locutions as well. (“I know we won the lottery, but still can’t quite get myself to believe it.”) Still in one of the prime senses of belief—that of acceptance, of commitment to the idea that something or other is so—the knowledge of matters of fact does require an acceptance that is either actual and overt or at least a matter of tacit implicit commitment to accept.

Various epistemologists have sought to characterize knowledge as *true justified belief*.¹ In his widely discussed 1963 article, Edmund Gettier followed up on

suggestions of Bertrand Russell by offering two sorts of counterexamples against this view of knowledge as consisting of beliefs that are both true and justified.

Counterexample 1

Let it be that:

1. X believes p
2. p is true
3. X has justification for believing p , for example, because it follows logically from something—say q —that he also believes, although in fact
4. q is false.

Here X clearly has justification for believing p , since by hypothesis this follows logically from something that he believes. Accordingly, p is a true, justified belief. Nevertheless, we would certainly not want to say that X *knows* that p , seeing that his (only) grounds for believing it are false.

To concretize this schematic situation let it be that:

1. X believes that Smith is in London (which is false since Smith is actually in Manchester).
2. Smith's being in London entails that Smith is in England (which conclusion is indeed true since Smith is in Manchester).
3. X believes that Smith is in England (because he believes him to be in London).

That Smith is in England is accordingly a belief of X 's that is both true and for which X has justification. Nevertheless we would clearly not want to say that X *knows* that Smith is in England, since his (only) reason for accepting this is something quite false.

The lesson that emerges here is that knowledge is not simply a matter of having a true belief that is *somehow* justified, but rather that knowledge calls for having a true belief that is *appropriately* justified. For the problem that the counterexample clearly indicates is that in this case the grounds that lead the individuals to adopt the belief just do not suffice to assure that which is believed. Its derivation from a false belief is emphatically *not* an appropriate justification for a belief.

Counterexample 2

Let it be that:

1. X believes p -or- q .
2. q is true (and consequently p -or- q is also true).

3. X disbelieves q .
4. X believes p -or- q , but does so (only) because he believes p .
5. p is false.

Here p -or- q is true. And X has justification for believing p -or- q since it follows from p which he believes. And since p -or- q is true—albeit in virtue of q 's being true (when X actually disbelieves)—it follows that p -or- q is a true, justified belief of X 's. Nevertheless, in the circumstances we would certainly not say that X knows that p -or- q , seeing that his sole grounds for believing it is once more something that is false.

The difficulty here is that X holds the belief p -or- q which is justified for X because it follows from X 's (false) belief that p , but is true just because q is true (which X altogether rejects).

To concretize this situation let it be that:

1. X believes that Jefferson succeeded America's first president, George Washington, as president
2. X accordingly believes that Jefferson or Adams was the second American president, although he thinks that Adams was the third president.
3. Since Adams was in fact the second American president, X 's belief that Jefferson or Adams was the second American president is indeed true

So X 's (2)-belief is indeed both true and justified. Nevertheless we would certainly not say that X *knows* this since his grounds for holding this belief are simply false. However, maintaining that knowledge is constituted by true and *appropriately* justified belief would once again resolve the problem, seeing that a belief held on the basis of falsehoods can clearly not count as appropriate. The difficulty is that the grounds on which the belief is held will in various cases prove insufficient to establish the belief's truth. And this blocks a merely *conjunctive* conception that knowledge is a matter of belief that is both true and justifiably adopted. But this problem is precluded by the *adjectival* conception of knowledge as belief that is at once correct and appropriately seen to be so, so that truth and justification are not separable but blended and conflated.

What is critical for knowledge attribution is that the believer's grounds for the particular belief at issue endorsed by the attributor as well.² When one says "A truth is known when it is a justified belief" one is not (or should not) take the line that it is believed *and* (somehow) justified, but rather that it is *justifiedly believed* in that the belief's rationale is flawless. The basic idea is that there can be no problem in crediting x with knowledge of p if:

X believes p on grounds sufficient to guarantee its truth and realizes this to be the case.

And so, the crucial point is that when knowledge is characterized as being true justified belief one has to construe justification in a complex, two-sided way because that belief must be accepted by its believer

- on grounds that *he* deems adequate

and moreover

- these grounds must be that *we* (the attributers of the belief) endorse by way of deeming them adequate as well.

The “subjective” justification of the attributee must be satisfied by the “objective” justification of the attributor if an attribution of knowledge is to be viable.

This line of consideration brings to light the inadequacy of cognitive conclusion: the view that knowledge is belief that is caused in the proper way. Because here “proper way” means proper as *we* (the attributers) see it, which may fail to be how the holder of that true belief sees it. (Those belief-engendering causes may fail to correspond to his actual cognitive grounds or *reasons* for holding the belief.)

And much the same is true of reliabilism: the view that knowledge is belief produced by a reliable process. For this reliability again holds for that attributer’s view of the matter and its issuing from that reliable process may not in fact be the believer’s own ground for holding the belief. (Causes can only provide *reasons* when their confirmatory operation is *correctly* recognized.)

So much for what is at issue with someone’s actually knowing a fact. But of course here, as elsewhere, there is a distinction between (1) something actually *being* so, and (2) having *adequate grounds for claiming* that it is so. And the former (actually being) always goes beyond the latter (having adequate grounds). We can have good reason for seeing our belief grounds as flawless even when this is actually not the case. In cognitive matters as elsewhere we must reckon with the prospect of unpleasant surprises. The prospect of error is pervasive in human affairs—cognition included.

Consider the issue from another angle. It is part of the *truth conditions* for the claim that something is an apple—a necessary condition for its being so—that was grown on an apple tree, that it contains seeds, and that it not turn into a frog if immersed in a bowl of water for 100 days. And yet many is the time we call something an apple without checking up in these things. The *use condition* for establishment to call something an apple are vastly more lenient. If it looks like an apple, feels like an apple should, and smells like an apple, then that is quite good enough.

And the same sort of thing also holds for knowledge (or for certainty). The truth conditions here are very demanding. But the use conditions that author-

ize responsible employment of the term in normal discourse are a great deal more relaxed.

Actually, to have knowledge is one sort of thing, something that goes well beyond is required for its being the case that you or I have adequate grounds for claiming that it is the case. With claims appropriate, assurance of all sorts stops well short of *guaranteeing* actual truth. And this is so with our subjectively justified knowledge claims as well.

It is important in this regard to note that one of the basic ground rules of (defeasible) presumption in the cognitive domain inheres in the rule of thumb that *people's conventionally justified beliefs are true*—where “conventional justification” is constituted by the usual sorts of ground (“taking oneself to see it to be so,” etc.) on which people standardly base their knowledge claims. And so whenever this groundrule comes into play the preceding formula that true justified beliefs constitute knowledge becomes redundant and knowledge comes to be viewed simply as appropriately justified belief where appropriateness may stop short of full-fledged theoretical adequacy. However what is at issue here is accordingly not a truth of general principle: it is no more than a *principle of practical procedure* which—as we fully recognize—may well fail to work out in particular cases.

Use conditions are geared to the world's operational realities. They bear not on what must invariably be in some necessitarian manner, but on what is usually and normally the case. And here—in the realm of the general rule, the ordinary course of things—it is perfectly acceptable to say that “knowledge is true justified belief.” For ordinarily subjective and objective warrant stand in alignment. Those cases where knowledge fails to accompany true and (subjectively) justified belief all represent unusual (abnormal, nonstandard) situations.

The *practical* justification of the principle at issue is nevertheless substantial. For in our own case, at any rate, we have no choice but to presume our conscientiously held beliefs to be true. The injunction “Tell me what *is* true in the matter independently of what you genuinely believe to be true” is one that we cannot but regard as absurd. And the privilege we claim for ourselves here is one that we are pragmatically well advised to extend to others as well. For unless we are prepared to presume that their beliefs too generally represent the truth of the matter one cannot derive information from them. An important principle of practical procedure is at issue. In failing to extend the credit of credence to others we would deny ourselves the prospect of extending our knowledge by drawing on theirs.³

MODES OF (PROPOSITIONAL) KNOWLEDGE

When one speaks of “belief” in matters of knowledge one has in view what a person stands committed to accepting by way of propositional (“it is the case

that”) contentions. Knowledge is not an activity—mental or otherwise. Nor is knowing a psychological process. What is at issue here is not a matter of a mode of action or activity. The question “What are you doing?” cannot be answered by saying “I am knowing that Paris is the capital of France.” This is something you may be thinking about or wondering about, but it is not something you can be knowing. The verb “to know” admits of no present continuous: one cannot be engaged in knowing. We can ascribe knowledge without knowing what goes on in people’s heads (let alone in their brains). To know something is a matter not of process but of product.

To realize that Columbus knew that wood can float I do not have to probe into the complex of thought processes he conducted in Italian or Spanish or whatever—let alone into his brain processes. Knowledge in the sense a issue here is a matter of a relation being held to obtain between a person and a proposition, and will in general have the format Kxp (“ x knows that p ”). Like owning or owing, knowing is a state: a condition into which one has entered. It is not an action one does or an activity in which one engages. (Those philosophers who do or have spoken of “the act of knowing” talk gibberish.) Instead what is at issue is a state or condition, namely the state or condition of a person who stands in a certain sort of relationship—a *cognitive* relationship—to a fact.

What is this relationship like? Here it is both useful and important to distinguish between explicit, dispositional, and inferential knowledge.

Explicit knowledge is a matter of what we can adduce on demand, so to speak. It takes two principal forms:

- *Occurrent knowledge.* This is a matter of actively paying heed or attention to accepted information. A person can say: “I am (at this very moment) considering or attending to or otherwise taking note of the fact that hydrogen is the lightest element.” The present evidence of our senses—“I am looking at the cat on the mat”—is also an example of this sort of thing where what is claimed as knowledge is geared to circumvent cognitive activity.
- *Dispositional knowledge.* This is a matter of what people would say or think if the occasion arose—of what, for example, they would say if asked. Even when X is reading Hamlet or, for that matter, sleeping, we would say that an individual knows (in the presently relevant dispositional manner) that Tokyo is the capital of Japan. Here those items of knowledge can automatically be rendered occurrent by suitable stimuli.

Inferential knowledge, by contrast, is (potentially) something more remote. It is something deeply latent and tacit, not a matter of what one *can* produce on demand, but of what one *would* produce if only one were clever enough about

exploiting one's occurrent and dispositional knowledge. We understand unproblematically what that means, namely—he must think, reflect, try to remember, exert some intellectual effort and patience. Inferential knowledge is a matter of exploiting one's noninferential (occurrent or dispositional) knowledge. It is the sort of extracted knowledge that heeds materials to work out—grist to its mill. Since it is transformative rather than generative, there can be no inferential outputs in the absence of noninferential inputs. (This, of course, is not to say that it cannot be new—or even surprising in bringing previously unrecognizable relationships to light.)

What an individual himself is in a position to infer from known facts are facts he also knows

if Kxp and $Kx(p \vdash q)$, then Kxq .

Here Kxp abbreviates “ x knows that p .” This sort of inferential knowledge is a matter of what the individual does—or by rights should—derive from what is known. It is this feature of knowledge—its *accessibility* that provides for the inferential construal of the idea that is at work in these deliberations. For our concern in epistemology is less with the impersonal question of what people do accept than with the normative question of what, in the circumstances, it is both appropriate and practicable for them to accept. Thus in attributing knowledge we look not only to the information that people have in a more or less explicit way but also to what they are bound to be able to infer from this.

This points toward the still more liberal conception of what might be called *available knowledge*, which turns on not what *the individual themselves* can derive from their knowledge but on what *we* can derive from it, so that something is “known” in this sense whenever $(\exists q)(Kxq \ \& \ q \vdash p)$. We arrive here at a disjunctive specification of fundamentally recursive nature:

X knows that p iff (i) X knows p occurrently, or (ii) X knows p dispositionally, or (iii) p can be inferentially derived by using only those facts that X knows.

The “logical omniscience” of the consequences of what one knows is a characterizing feature of this particular (and overly generous) conception of knowledge. K^* is thus radically different from K .

Since one of our main concerns will be with the limits of knowledge, it will be of some interest to see if there are things not known even in this particularly generous sense of the term that is at issue with *available knowledge*. In general, however, it will be the more standard and received sense of *accessible knowledge* that will concern us here.

OTHER BASIC PRINCIPLES

Accessible knowledge involves its holder in acceptance (endorsement, subscription, credence, belief, etc.). To accept a contention is to espouse and endorse it, to give it credence, to view it as an established fact, to take it to be able to serve as a (true) premiss in one's thinking and as a suitable basis for one's actions.⁴ And a person cannot be said to *know* that something is the case when this individual is not prepared to "accept" it in this sort of way. And accordingly, the claim "*x* knows that *p*" is only tenable when *x* holds *p* to be the case. It is senseless to say "*x* knows that *p*," but he does not really believe it or "*x* knows that *p*, but does not stand committed to accepting it."⁵

Accordingly, one cannot be said to know something if this is not true.⁶ Let "*Kxp*" abbreviate "*x* knows that *p*." It then transpires that we have:

- *The Veracity Principle*

If *Kxp*, then *p*.

This relation between "*x* knows that *p*" and "*p* is true" is a necessary link that obtains *ex vi terminorum*. Knowledge must be veracious: The truth of *p* is a presupposition of its knowability: if *p* were not true, we would (*ex hypothesi*) have no alternative (as a matter of the "logic" of the conceptual situation) to withdraw the claim that somebody *know* *p*.

Some writers see the linkage between knowledge and truth as a merely contingent one.⁷ But such a view inflicts violence on the concept of knowledge as it actually operates in our discourse. The locution "*x* knows that *p*, but it is not true that *p*" is senseless. One would have to say "*x* *thinks* he knows that *p*, but . . ." When even *the mere possibility* of the falsity of something that one accepts comes to light, the knowledge claim must be withdrawn; it cannot be asserted flatly, but must be qualified in some such qualified way as "While I don't actually *know* that *p*, I am virtually certain that it is so."

Knowledge veracity straightaway assures knowledge consistency. If someone knows something, then no one knows anything inconsistent with it.

- *Knowledge-Coherence Principle*

If *Kxp* and *Kyp*, then *p* compat *q*.

Proof. If *Kxp*, then *p*, and if *Kxq*, then *q*. Hence we have *p* & *q*, and so, since their conjunction is true, *p* and *q* cannot be incompatible.

Again, certain other "perfectly obvious" deductions from what is known must be assumed to be at the disposal of every (rational) knower. We thus have the principles that separably known items are cognitively conjunctive:

- *The Conjunctivity Principle:*

Someone who knows both p and q separately, thereby also knows their conjunction:

If Kxp and Kxq , then $Kx(p \ \& \ q)$ —and conversely.

A person who is unable to exploit his information by “putting two and two together” does not really have *knowledge* in the way that is at issue with specifically *inferential* knowledge. It seems only natural to suppose that any rational knower could put two and two together in this way.

To these principles of *inferential* knowledge we can also adjoin the following:

- *Knowledge-Reflexivity Principle*

Actually to know that someone knows something requires knowing this fact oneself.

If $KxKyp$, then Kxp ; and indeed: If $Kx(\exists y)Kyp$, then Kxp .

Letting i be oneself so that Kip comes to “I know that p ” we may note that the Knowledge-Reflexivity Principle entails

If $KiKyp$, then Kip ; and more generally: if $Kx(\exists y)Kyp$, then Kxp .

To claim to know that someone else knows something (i.e., some specific fact) is to assert this item for oneself. Put differently, one can only know with respect to the propositional knowledge of others that which one knows oneself.

However, it lies in the nature of things that there are—or can be—facts that X can know about Y but Y cannot. Thus X can know that Y has only opinions but no knowledge, but Y cannot.

Knowledge entails justification (warrant, grounding, evidence, or the like). One can maintain that someone knows something only if one is prepared to maintain that he has an adequate rational basis for accepting it. It is senseless to say things like “ x knows that p , but has no adequate basis for its acceptance,” or again “ x knows that p , but has no sufficient grounding for it.” To say that “ x knows that p ” is to say (*inter alia*) that x has *conclusive* warrant for claiming p , and, moreover, that x accepts p on the basis of the conclusive warrant he has for it (rather than on some other, evidentially insufficient basis). It is senseless to say things like “ x knows that p , but there is some room for doubt” or “ x knows that p , but his grounds in the matter leave something to be desired.”⁸

G. E. Moore pointed out long ago that it is anomalous to say “ p but I don’t believe it” or indeed even “ p but I don’t know it.”⁹ That this circumstance of what might be called “Moore’s Thesis” be so is readily shown. For in maintaining a

claim one standardly purports to know it. Accordingly, one would not—should not—say $p \ \& \ \sim Kip$ unless one were prepared to subscribe to

$$Ki(p \ \& \ \sim Kip),$$

where i = oneself. But in view of Veracity and Conjunctivity this straightaway yields

$$Kip \ \& \ \sim Kip,$$

which is self-contradictory. (Observe, however, that “I suspect p but do not fully believe it” [or “really know it”] is in a different boat.)

It is also important to recognize that which is known must be compatible with whatever else is actually known. No part of knowledge can constitute decisive counterevidence against some other part. The whole “body of (genuine) knowledge” must be self-consistent. Accordingly we shall have

$$\text{If } Kxp \text{ and } Kxq, \text{ then compat } (p, q)$$

and consequently:

$$\text{If } Kxp \ \& \ Kxq, \text{ then } \sim(p \ \vdash \ \sim q).^{10}$$

Given this circumstance, logic alone suffices to assure the implication:

$$\text{If } Kxp \text{ and } p \ \vdash \ q, \text{ then } \sim Kx\sim q.$$

Accordingly, one decisive way of defeating a claim to knowledge is by establishing that its denial follows from something one knows.¹¹ But of course not knowing something to be false (i.e., $\sim Kx\sim q$) is very different from—and much weaker than—knowing this item to be true (Kxq). And so—as was just noted in the preceding dismissal of “logical omniscience”—the just-indicated principle must *not* be strengthened to the objectionable

$$\text{If } Kxp \text{ and } p \ \vdash \ q, \text{ then } Kxq,$$

which has already been rejected above.

A further, particularly interesting facet of knowledge-discourse relates to *the automatic self-assumption of particularized knowledge attributions*. It makes no sense to say “You know that p , but I don’t” or “ x knows that p , but not I.” In *conceding* an item of knowledge, one automatically *claims* it for oneself as well. To

be sure, this holds only for that-knowledge, and not how-to-knowledge (even how to do something “purely intellectual”—like “answering a certain question correctly”). It makes perfectly good sense to say that someone else knows how to do something one cannot do oneself. Again, abstract (i.e., unidentified) knowledge attributions will not be self-assumptive. One can quite appropriately say “ x knows everything (or ‘something interesting’) about automobile engines, though I certainly do not.” But particularized and identified claims to factual knowledge are different in this regard. One cannot say “ x knows that automobile engines use gasoline as fuel, but I myself do not know this.” To be sure, we certainly do *not* have the omniscience thesis:

If Kxp , then Kip (with $i = I$ myself).

One’s being entitled to claim Ksp follows from one’s being entitled to claim Kxp , but the content of the former claim does not follow from the content of the latter.

Moreover, the ground rules of language use being what they are, *mere assertion is in itself inherently knowledge-claiming*. One cannot say “ p but I don’t know that p .” To be sure, one can introduce various qualifications like “I accept p although I don’t actually know it to be true,” But to affirm a thesis flatly (without qualification) is *eo ipso* to claim knowledge of it. (Again, what is at issue is certainly *not* captured by the—clearly unacceptable—thesis: If p , then Kip .)

Such “logical principles of epistemology” will of course hinge crucially on the exact construction that is to be placed on the conception of “knowledge.” In particular, if it were not for the inferential availability character of this concept, the situation would be very different—and radically impoverished—in this regard.

One further point. The statement “possibly somebody knows p ” $\diamond(\exists x)Kxp$ says something quite different from (and weaker than) “somebody possibly knows p ” $(\exists x)\diamond Kxp$. It is accordingly necessary to distinguish

1. Only if p is true will it be possible that somebody knows it:

If $\diamond(\exists x)Kxp$, then p

from

2. Only if p is true will there be someone who possibly knows it:

If $(\exists x)\diamond Kxp$, then p .

Because the modality in (1)’s antecedent precedes that qualified the statement carries us beyond the bounds of the actual world into the realm of merely

possible existence. But with (2)'s antecedent we remain within it: we discuss only what is possible for the membership of *this* world. Hence (2) is a plausible thesis, while (1) is not.

Be this as it may, the cardinal point is that of clarifying the nature of knowledge. And in this regard *putative* knowledge is a matter of someone's staking a claim to truth for which that individual (subjectively) deems himself to have adequate grounds, while *actual* knowledge by contrast is a matter of someone's correctly staking a claim to truth when that individual has (objectively) adequate grounds.

The various principles specified here all represent more or less straightforward facts about how the concept at issue in talk about actual knowledge actually functions. Any philosophical theory of knowledge must—to the peril of its own adequacy—be prepared to accommodate them.