

---

April 2006

## On Public Knowledge

Joanne Lau

Follow this and additional works at: <https://scholarworks.wmich.edu/hilltopreview>



Part of the Philosophy Commons

---

### Recommended Citation

Lau, Joanne (2006) "On Public Knowledge," *The Hilltop Review*: Vol. 2 : Iss. 1 , Article 8.  
Available at: <https://scholarworks.wmich.edu/hilltopreview/vol2/iss1/8>

This Article is brought to you for free and open access by the Graduate College at ScholarWorks at WMU. It has been accepted for inclusion in The Hilltop Review by an authorized editor of ScholarWorks at WMU. For more information, please contact [wmu-scholarworks@wmich.edu](mailto:wmu-scholarworks@wmich.edu).

## On Public Knowledge

Joanne Lau  
Department of Philosophy  
Western Michigan University

In “Elusive Knowledge”, Lewis begins by saying that we know a lot of trite information about the ordinary world<sup>1</sup>. I know what carrots are. I know how to drive a car. I know that when someone greets me and holds out their hand, they expect me to shake it and so on. However, when we talk in philosophical terms about knowledge, we have no reason to say that we *really* know anything. There seem to be two contexts of knowing something. Regarding the everyday, general type of knowledge, I know a lot. But when I put on my philosopher’s hat, it looks like I can’t say I really know very much at all. As a result, it is argued that either we know lots of things, or we must submit to skepticism<sup>2</sup>.

Lewis does not think this is the case, and I agree. I do not think that we need to be skeptics; we can claim that we know a lot of things. In a further blow to skepticism, the things that we *do* know are usually not isolated. So my argument is that, unlike the skeptic, not only do I know a lot of things, I know a lot about things about the world that I have not checked for myself.

### *The argument against Cartesian skepticism*

In the traditional skeptics’ view, Descartes begins with the individual thinker and no knowledge, and from that view we must first establish that I exist, and that God exists, and that material objects exist, etc. in order to find out about our world<sup>3</sup>. However, Coady argues that people do not actually operate that way. We do not begin with nothing and then establish facts which entail other facts to ascertain whether or not we know something<sup>4</sup>. If we accept deductive closure, then to know  $p$  entails knowing all of the entailments of  $p$ <sup>5</sup>. Thus, if I look at a tree, then it follows that I know what a tree is, that I am not hallucinating, that I am not a brain in a vat on Alpha Centauri, and so on. But I have not personally checked that my eyes are working properly,

---

<sup>1</sup> Lewis, D., *Papers in Metaphysics and Epistemology*, Cambridge University Press 1999. pp 418-445.

<sup>2</sup> This is the philosophical position where, roughly put, we cannot say we really know anything because we cannot attain the required level of certainty for any information we may claim to know.

<sup>3</sup> Descartes’ *Meditations on First Philosophy*, edited by Cottingham, J., Cambridge University Press 1996.

<sup>4</sup> Coady, C. A. J., *Testimony*, Oxford University Press 1992, pp 144-147.

<sup>5</sup> This is a highly contentious point. Deductive closure is where knowledge of  $p$  and that  $p$  entails  $q$ . So if I know  $p$  and that  $p$  entails  $q$ , then  $q$ . Nozick and Dretske argued against it, but it is not a major issue here and for the purposes of the paper I will simply accept deductive closure: Nozick, R., *Philosophical Explanations*, Harvard University Press 1981.

or that I am not in a vat, or that I am not on Alpha Centauri. Further, if I had to find out all these peripheral things (such as proving that my eyes are working), there would be other facts following from the peripheral facts that I would need to also prove (such as the fact that I actually have eyes) and so on. But, luckily, we do not actually operate in that fashion. To know  $p$ , I do not have to establish the truth of other propositions  $q$ ,  $r$  and  $s$  which are entailed by  $p$ <sup>6</sup>. So while I have not utterly defeated the Cartesian skeptic, I have at least pointed out that the skeptical view is neither practical nor applicable to ordinary people in ordinary situations.

Similarly, Searle argues that we are equipped with a background which tells us, in a given situation, what information is superfluous, and what we should be paying attention to<sup>7</sup>. This background is made up of our experiences and interactions with the external world. But in that external world are also other *people*, so receiving information from people must be considered in the same manner as receiving information about trees. When receiving information from another person, the skeptic argues that you must ascertain whether or not that person is reliable, because you can rely only on testimony from established sources. So, according to the skeptic, you must figure out that there is an external world, and that certain things in the world are people, and that they express words, and that some are reliable. Then, based on the words you have heard and your past interactions with this person, you can determine whether or not to believe what they have said. But again, we do not actually function in this way. Instead, given the necessity for social interactions, we generally trust that a person is telling the truth and accept what he or she is saying.

I will discuss the concept of shared knowledge later in this paper. In the meantime, I will examine *what* it is that we are sharing, before discussing *how* it is that the information is shared in a social setting.

### ***The Context and Content of Knowledge***

Lewis claims that when we say that we know something, we mean that we know it *within a certain context*<sup>8</sup>. The background lets us know what the appropriate context is for a given situation – so that you and I will understand each other and can exchange information<sup>9</sup>. Suppose I am at a bar and observe that “every glass is empty”. “Every” is a universal quantifier, but surely I do not mean that “every glass in the entire world is empty” at the time of my utterance; that would be absurd. Instead, what I mean is that I am narrowing my utterance to meaning “every glass within the context of this bar is empty”. In everyday language, there is an implicit contextual limitation in what we say.

We must distinguish between the context of knowledge and its *content*. I might know that water boils at 212 degrees Fahrenheit, because this is a type of regularity that I have observed whilst making tea at home. My knowledge about boiling water is limited to the context of

---

<sup>6</sup> Coady, op cit.

<sup>7</sup> Searle, J., *The Construction of Social Reality*, Free Press 1995, pp 129-132.

<sup>8</sup> Lewis, “Elusive Knowledge”, pp 425-426.

<sup>9</sup> Searle, op cit.

ordinary tea making. But were I to try and boil water on the top of Mount Everest, I might find out that, there, water does *not* boil at 212°F. So after my expedition to Mount Everest, the content my knowledge about the boiling point of water has changed: in all *ordinary* circumstances, and at *normal* pressure, water boils at 212°F. But I also know that there are exceptional circumstances at which water does not boil at 212°F.

When a fact's context and the content are put together, we encounter situations where the content of what we assert remains the same while the context changes. I might say that I know the table is flat, but there is nothing that is *completely* flat<sup>10</sup>. So what I mean is that "relatively speaking, and ignoring the microscopic bumps here and there, I know that the table is flat." However, flatness changes with context. A crease on a cricket pitch may be considered "flat" for the purposes of playing cricket, but that surface is not "flat" for someone who is an ice-skater or plays curling, because, for those sports, the surface of the ground requires a higher degree of flatness than it does for cricket. So while the content of our facts (that is, "this is/is not a flat surface") remains the same, the context changes depending on other factors (such as its material, intended use, etc.).

### ***The context of knowledge made public***

When I say to you that "the table is flat", I am not only expressing my belief that the table is flat, but I am also assuming that we share enough background knowledge so that *you* have certain basic knowledge about what a table is, and you know what flatness is, and so on. Similarly, you would also be assuming that when I am talking about a table being flat, what I am talking about conforms to what you know about tables and what you know about flatness.

There is a lot of knowledge that is shared in the public domain, or by at least two individuals. Searle's argument for collective intentionality is relevant, and I will adopt it as follows<sup>11</sup>. Given our social environment, there is a large set of knowledge that is shared, or at least tacitly assumed for us to be able to interact with each other and function socially<sup>12</sup>. We make these assumptions about shared knowledge regularly. The rules of multi-player games, road laws, etiquette, etc. all rely on the presence of a common knowledge among its participants. Wittgenstein's argument for "language-games" also reflects this view in that language exists fundamentally for the purposes of human activity and as such language is possible only when people share a "form of life", such as those discussed in this section<sup>13</sup>.

So there is a large pool of knowledge which we assume is shared in order for everybody to be able to interact and work together in a society (such as everyone driving on the same side

---

<sup>10</sup> This example is attributed to Lewis, op cit., pp 425-426. See also Unger, P., *Ignorance*, Clarendon Press 1975, especially chapter II.

<sup>11</sup> Searle, op cit., pp 23-26.

<sup>12</sup> For example, we assume that two musicians playing a duet would know what each other was doing in order for them to play together in that they share knowledge about how to play their instruments, how to read music, etc.

<sup>13</sup> Wittgenstein, L., *Philosophical Investigations*, 1953, II.

of the street, etc.). In fact, a lot of the things we do are based on *convention* – that is, we roughly have a convention when:

- 1) everyone conforms to a regularity *R*;
- 2) everyone expects everyone else to conform to *R*;
- 3) everyone prefers to conform to *R* on condition that the others do, since situation *S* is a coordination problem and uniform conformity to *R* is coordination equilibrium in  $S^{14}$ .

Many of our conventions are tacit – nobody has discussed the social norms of what we do but what we *actually* do is founded on a set of mutual expectations. I expect that, in a certain situation, you expect me to act in a certain way, and so I act in that way because that is what I expect you to expect from me<sup>15</sup>. In other words, in a case of independent decision between rational people, if I am figuring out what you will do, I am trying to replicate your practical reasoning. In order to do this, I need to figure out what *you* expect *me* to do, and this means that you would be replicating *my* reasoning. So, to replicate your reasoning, I need to replicate *your attempt* to replicate my reasoning. In this way, a mutual expectation arises, since we each expect the other to behave in a certain way<sup>16</sup>. For example, suppose you and I have arranged to meet at a particular time in a foreign city, but we have not specified a particular location in that city. In order for us to meet, I must figure out where *you* are expecting that I will expect to meet, and go there. Perhaps there is a large bridge in the middle of the city; if I think that you would think that I would consider a noticeable bridge to be a good place to meet, and if *you* also think that I think that you would choose such a place, then we will be able to meet at the bridge<sup>17</sup>.

### ***The causes of shared knowledge***

Further, a social environment such as ours is constantly being saturated by information. Labels on food, street signs, the news on television, covers of books etc. are bearers of information which we can obtain as knowledge. For example, as I walk up the street, there is information that I have come know, because I can see and read it. So I know that today, gas is at a certain price. I know that there is a special book offer with today's newspaper. I know that you can park here only if you have a particular permit. I know that there is a crossing, and if you want to cross the road you push a button. And so on. Any other person who can see and read can also obtain this information.

But *how* is this information shared? I submit that knowledge, or at least knowledge of *this* kind, is causally obtained. As I walk up the street, I am bombarded with information about gas prices, parking zones, book offers and so on. I am not given any justification for this information,

---

<sup>14</sup> Lewis, D., *Convention*, Harvard University Press 1969, p 42.

<sup>15</sup> Ibid, pp 26-27.

<sup>16</sup> Searle would call this a collective intentionality, and I think the concepts are virtually identical for this issue at least.

<sup>17</sup> However, conventions are not always perfect, and we can become unwittingly locked into them. A bridge may not be the most practicable place to meet if it is raining.

but I assume that there is some sort of underlying cause for the information I see. If there were no books, then there would be no book offer. So the fact that there is a book offer is causally dependent on the existence of books, and the sale of the books is causally dependent on people's knowledge of the book offer. But I personally have not checked if there are books or not; it is an assumption I make based on the common knowledge that generally, when people say that they are offering books, they really are offering books.

But, regardless of whether or not I know the originating cause of a piece of information in the public domain, I can say that there is some sort of cause of my acquisition of that piece of information. Suppose that you know that ravens are black, but I have never seen a raven before. You tell me that ravens are black and are of a particular size. You show me lots of empirical data about ravens from ornithologists and photographs of black birds. On the basic requirements of knowledge, I now have *some* sort of true, justified belief about the blackness of ravens, even though I have never seen a raven before<sup>18</sup>. And if I were to see a black bird of such a size and shape, with a particular set of behavior patterns and so on, then were I to say "That bird is a raven", then I am saying that "I know the bird fits the description of a raven", and chances are that I would be right. So I can say that I *do* know that ravens are black. But what is more is that my knowledge of the blackness of ravens is counterfactually dependent on your knowledge. And so knowledge is passed on in this fashion.

This is true of many things. I can assert that I know that Napoleon was the emperor of France, although Napoleon is now dead and I have never met him, and nobody alive that I know has ever met him. I can also say that I know that Mount Everest is the tallest mountain in the world, even though I have never been to Mount Everest or measured the height of any mountains generally. But there are people who *have* been to Mount Everest and other mountains, and have said that there is no mountain taller than Everest. If I believe their testimony, I can say that I know that Mount Everest is the tallest mountain in the world.

Likewise, the content of what we know depends on the causal history of that particular piece of information. We can have reliable knowledge even without justifications. However, in order for the knowledge to stay within the public sphere, it must be *reaffirmed* regularly. The vast majority of us would be inclined to say that we know the earth is round. But what is the justification for that claim? Perhaps I heard it from other people. Even so, if I knew who I heard it from, then what is *their* justification for the claim? Over time, as our knowledge is shared, our justification for that knowledge becomes weaker. So, every now and then, we have to rejustify and reinforce our knowledge.

For example, I may already know that the earth is round from other people's testimony. But when I am on a boat and I see the land disappear over the horizon, I might figure that the reason I cannot see the land anymore is due to the curvature of the earth. And *eureka!* I realize that I know the earth is round. But it is absurd to think that I did not *know* this before. I did know it, but the very fact that I have seen the effects of the earth being round for myself indicates that

---

<sup>18</sup> I am aware that there is much debate about whether knowledge is really justified true belief but this is not the issue I am looking at here: Gettier, E., "Is Knowledge justified, true, belief?" *Analysis* 23 (1963).

my knowledge has been confirmed. And so my link in the causal chain of knowledge about the earth's roundness is strengthened, and any of the subsequent links after mine are also stronger.

***Expert testimony: the pyramid and lattice structures***

So, sharing knowledge is a process of dividing epistemic labor<sup>19</sup>. We have experts in different fields who obtain (or reconfirm) facts and pass them into the public domain. But it is not that simple. The experts we rely on are also relying on other experts! Scientists are permitted to *assume* such things such as 212°F being the boiling point of water and so on; they do not have to go and empirically test this out every time to justify their experiments, so scientists base their findings on presumptions established by *other* scientists, and those scientists on other scientists, and eventually if we dig down enough through the levels of knowledge we will end up with an originating cause of a piece of publicly known information and there would be many levels of knowledge, like a pyramid.

However, public knowledge does not have a unilateral direction of fit. It is very rare for an expert (or anyone else) to rely on sources without proffering some reciprocating information. I might go to a doctor because I am unwell, and the doctor has knowledge as to my symptoms and how to alleviate them. But although the doctor has more medical knowledge than I do, a doctor cannot assess my symptoms and suggest a remedy without relying on my at least describing how I am feeling. A doctor relies on testimony from his or her patient in order to know what ailment the patient is suffering from. I am relying on my doctor to tell me about my symptoms, but in order for him (or her) to do so, he or she must first know dip into the pool of public knowledge to ascertain what the symptoms indicate. But I am part of that pool of knowledge, because the doctor relies on what I know about how I feel!

So perhaps the pyramid model for knowledge is not exactly correct. The simplistic pyramid structure has been widely criticized because it does not account for the mutual reinforcement between an explanation and what it explains<sup>20</sup>. My knowledge fits in with my other beliefs and knowledge, but those other beliefs have been selected because they also fit with my knowledge. A more suitable version is a lattice or web structure, in which we have an interlocking set of shared facts. It is very rare that one would have to rely on the knowledge of someone who has not (at some other point) relied on oneself. Knowledge is pervasive and the mutual expectations we have of each other exist at all levels, between experts and non-experts.

***Causation and the lattice structure***

How does the causal chain fit the lattice structure? I have indicated that causation has a unilateral direction of fit, so how do we obtain mutual support for a causal picture? I think the answer lies in our conventions<sup>21</sup>. Convention by definition has symmetry to it, based on *mutual*

---

<sup>19</sup> This is analogous to the division of linguistic labor that Putnam refers to in his paper: "Meaning and Reference", *Journal of Philosophy*, **73** (1970).

<sup>20</sup> Quine W. V., and Ullian, J., *The Web of Belief* (2<sup>nd</sup> Ed.), Random House, 1979, p 79.

<sup>21</sup> See Lewis, *Convention*, *ibid*.

*expectation*<sup>22</sup>. Take the previous example of my walking down the street. I suppose there is a cause for the signs that I see. This is one direction that the causation and expectation run in. Now, suppose I own a bookstore, and I want to sell books. A person wanting books goes to a bookstore. A person wanting to sell books wants people who want books to go to their bookstore. So we have established a mutual expectation of selling/being sold books. But where is the causation? One cause of books being sold is advertising. I might put up a sign that says “books for sale”. Assuming others can read and use the same language, whether or not I sell books (and how many) counterfactually depends on what sign I put up (if at all). So human action is more causal than logical, because there may not be any logical reason for selling the books, but there is a causal reason arising out of social convention.

Similarly, in many cases, we are mutually the causes of each other’s knowledge. We are a community of knowers, and it is commonplace for us to learn from each other<sup>23</sup>. Quine suggests that the “totality of our knowledge or beliefs... is a man-made fabric which impinges on experience only along the edges”<sup>24</sup>. This suggests that the causal relation exists at the periphery between knowledge and experience. However, a “thread” is causally related to the other “threads” by virtue of being interwoven, and if a “thread” is changed at the periphery, the change also has an effect on the other threads.

A well-known illustration is the Ship of Neurath<sup>25</sup>. All sides of the boat mutually support other sides and rely on a mutual configuration of the ship. The set of planks in the ship have a relationship between them and the whole set of planks (that is, the ship) is held up by the water. The water as a whole holds up the entire hull of the ship in virtue of the two-way forces between the planks in the hull. Likewise, human knowledge is supported by reality as a whole, not one “plank” at a time. Each item depends on others for support. The interrelated causal structure between the planks is also causally related to the world.

However, within the lattice structure, there are sections which may be more unilateral than others. A kindergarten teacher may be able to impart a lot of knowledge to her students, but not receive much in return. Conversely, a university professor and a class of graduate students may have more balance in knowledge and therefore a more lattice-like information exchange. Nonetheless, the overall structure of our shared knowledge is lattice-like, and such “pyramidal” instances are uncommon.

### **Conclusion**

The traditional view of the Cartesian skeptic is mistaken. We do not actually behave in the way that the skeptic supposes, as we have a large pool of shared knowledge. Our social background tells us what the relevant knowledge to draw from the pool is in any given case. We already have a common language between us so the content of our knowledge can be expressed,

---

<sup>22</sup> Again, we could use Searle’s term “collective intentionality”.

<sup>23</sup> Kitcher, P., *The Nature of Mathematical Knowledge*, Oxford University Press 1983, pp 4-5.

<sup>24</sup> Quine, W. V., *From a Logical Point of View*, Harper Torchbooks, 1961, pp 42-44.

<sup>25</sup> Quine, W.V., *Word and Object*, MIT Press, 1960 pp 3-4.

and it can be understood to be within a particular context. This is due to our social conventions governing language and social behavior.

Our knowledge is shared through the causal history of the facts at hand. We pass knowledge to one another and occasionally reaffirm what we know in the public domain. Expert testimony is particularly relevant as it shows our readiness to rely on facts that we have not empirically proven ourselves. However, there is interdependence in the division of epistemic labor and as such the imparting of knowledge takes place on a lattice-like, and not a pyramidal, structure. In sharing facts we are also commonly mutual causes of each other's knowledge. And so it is through this combination of mutual dependence and sharing that we arrive at a theory of knowledge which truly reflects the society in which we live.

### References

- Coady, C. A. J., *Testimony*, Oxford University Press 1992.
- Descartes, *Meditations on First Philosophy*, edited by Cottingham, J., Cambridge University Press 1996.
- Gettier, E., "Is Knowledge justified, true, belief?" *Analysis* 23 (1963).
- Kitcher, P., *The Nature of Mathematical Knowledge*, Oxford University Press 1983.
- Kripke, S., *Naming and Necessity*, Blackwell publishers 1980.
- Lewis, D., *Papers in Metaphysics and Epistemology*, Cambridge University Press 1999.
- Lewis, D., *Convention*, Harvard University Press 1969.
- Nozick, R., *Philosophical Explanations*, Harvard University Press 1981.
- Putnam, H., "Meaning and Reference", *Journal of Philosophy*, **73** (1970).
- Quine W. V., and Ullian, J., *The Web of Belief* (2<sup>nd</sup> Ed.), Random House, 1979.
- Quine, W. V., *From a Logical Point of View*, Harper Torchbooks, 1961.
- Quine, W.V., *Word and Object*, MIT Press, 1960.
- Searle, J., *The Construction of Social Reality*, Free Press 1995.
- Unger, P., *Ignorance*, Claredon Press 1975.
- Wittgenstein, L., *Philosophical Investigations*, 1953.