

■ Artigo

doi: 10.4013/fsu.2026.271.16

Does knowledge circulate? Fleck, reliabilism, and the social status of knowledge

O conhecimento circula? Fleck, o confiabilismo e o status social do conhecimento

Marco Antonio Azevedo* 

Pontifícia Universidade Católica do Rio Grande do Sul, Porto Alegre, RS, Brazil.

E-mail: mazevedogtalk@gmail.com

Editores responsáveis:

Inácio Helfer

Leonardo Marques Kussler

Luís Miguel Rechiki Meirelles

*Pesquisador do CNPq

ABSTRACT: The claim that scientific knowledge cannot be exhaustively analyzed as an individual mental state has gained prominence in recent social epistemology. This paper advances that claim by examining empirical cases and engaging critically with a recent proposal to reconcile the two dominant contemporary accounts of epistemic justification—reliabilism and evidentialism. I argue for a conception of knowledge as a socially embedded epistemic state, one that aligns with ordinary intuitions in cases where individuals lack the justificatory access demanded by the classical account of knowledge as justified true belief. Drawing on Ludwik Fleck's theory of thought collectives, I examine shifting positions regarding the effectiveness of cloth masks during the COVID-19 pandemic to illustrate how scientific knowledge is historically situated, dynamically maintained, and circulates bidirectionally between esoteric and exoteric communities. The discussion concludes with exploratory remarks on the metaphysical dependence of individual knowledge states on their collective counterparts.

Keywords: social epistemology; reliabilism; evidentialism; Ludwik Fleck; COVID-19; scientific knowledge; thought collectives.

RESUMO: A afirmação de que o conhecimento científico não pode ser analisado exaustivamente como um estado mental individual ganhou destaque na epistemologia social recente. Este artigo avança essa afirmação examinando casos empíricos e dialogando criticamente com uma proposta recente para reconciliar as duas principais concepções contemporâneas de justificação epistêmica: o confiabilismo e o evidencialismo. Defendo uma concepção de conhecimento como um estado epistêmico socialmente inserido, que se alinha com as intuições comuns nos casos em que os indivíduos não possuem o acesso justificatório exigido pela concepção clássica de conhecimento como crença verdadeira justificada. Com base na teoria dos coletivos de pensamento de Ludwik Fleck, examino as mudanças de posição em relação à eficácia das máscaras de tecido durante a pandemia de COVID-19 para ilustrar como o conhecimento científico é historicamente situado, dinamicamente mantido e circula bidirecionalmente entre comunidades esotéricas e exotérmicas. A discussão conclui com observações exploratórias sobre a dependência metafísica dos estados de conhecimento individuais em relação às suas contrapartes coletivas.

Palavras-chaves: epistemologia social; confiabilismo; evidencialismo; Ludwik Fleck; COVID-19; conhecimento científico; coletivos de pensamentos.

Como citar:

Azevedo, M. A. 2026. Does knowledge circulate? Fleck, reliabilism, and the social status of knowledge. *Filosofia Unisinos*, 27(1):1-13, e27116. doi: 10.4013/fsu.2026.271.16



► 1 Knowing without justification? Ordinary intuitions and the puzzle of common knowledge

We've known for some time that AIDS is caused by a virus. Many of us have known that it's caused by a type of virus, a retrovirus. To arrive at what we know today, much research and scientific debate have taken place. The disease was recognized in 1981, but it wasn't until a few years later that it was concluded that it was a viral illness. The virus that causes AIDS was named HIV (Human Immunodeficiency Virus) in 1983. It is believed to have originated from a simian immunodeficiency virus (SIV) found in chimpanzees in Central Africa and to have been transmitted to humans, likely through the slaughter of chimpanzees for bushmeat, sometime in the early or mid-20th century. The virus then spread throughout human populations, and the first cases of AIDS were identified in the early 1980s.

Human knowledge of AIDS has expanded significantly from 1981 to the present day. There have been moments of intense debate and even disagreement. But nowadays, it's rare to find anyone who disagrees with the statement that AIDS is a viral disease. People generally believe this, so the statement:

“AIDS is caused by a virus”

can be taken as a general or common opinion. However, this is not just any opinion. People have opinions on many subjects and differ on them. But the fact that AIDS is caused by a virus is no longer a debatable fact. It is a fact taken as undeniable. The statement is certainly an element of scientific knowledge, because, after all, the fact that AIDS is caused by a virus was a scientific discovery. This discovery occurred in the labs of a few laboratories, and the statement circulated and was debated by scientists for a long time. But this scientific fact is not only recognized as such by today's virologists or infectious disease specialists. This is a widely recognized fact, such as the fact that “Earth is a planet in the solar system” or even more childish and widely shared statements like “Cooking food prevents disease.” Thus, it seems beyond doubt that:

We know that AIDS is caused by a virus, meaning that this is common knowledge, something that everyone KNOWS.

Now consider John, an ordinary person whose formal education doesn't go beyond elementary school. Suppose we asked him about the cause of AIDS and whether he agrees that AIDS is caused by a virus. John says he agrees with the statement. But if we press him to explain it to us, to justify it, perhaps he would say, if he were a good-natured and insightful person, something like: “I think it's better for you to ask that question to a doctor, not me.” Should the fact that John can't provide us with justifications for why he believes AIDS is caused by a virus be taken by us as evidence that he doesn't actually know this, but merely BELIEVES it? That only scientists know, and not him, an ordinary person?

Consider Mauro's case. He's a pediatrician. Mauro knows many things, including that many cases of acute otitis media are caused by viruses, not bacteria. He also knows that ear infections are more common in young children and that one theory for this high prevalence is that these children are more susceptible to viral infections (because they lack a developed specific immunity) and that their facial anatomy does not favor the clearance of mucus through the auditory tubes. But of course, these explanations are merely simplifications of something much more complex. Mauro knows a lot about pediatric practice, but if we ask him about more subtle aspects of virology (what are the most common viruses that cause such diseases in children?), he may not be able to answer. At least not immediately. But if he searches for

this information (nowadays online), he can quickly provide more precise information. Even so, there are some details that he might say: “It’s best to consult a virologist in this case.” Consider Tina’s case. She took her 3-year-old daughter, Lana, to Dr. Mauro. She was experiencing nasal congestion and had been in contact with schoolmates with flu-like symptoms. The doctor diagnosed nasopharyngitis and noted the presence of acute otitis media. However, he associated the otitis with a viral upper respiratory infection. He went on to tell her that Lana had acute otitis media, that there was secretion and a mild inflammation. In his words, she has “an acute middle ear effusion” (MEE), but that since the child didn’t complain of pain, had no fever, and the signs of inflammation were mild, he preferred to wait, maintain contact, and reevaluate later. He gave Tina the option of starting antibiotics, but stated that the progression toward improvement would likely be the same if they chose not to.¹ Tina was convinced by Dr. Mauro’s explanation that it was better to wait a few days and that there was no need to give Lana antibiotics. Dr. Mauro asked her to bring Lana back a week after the appointment and to contact him earlier if she complained of pain or has fever (something she wasn’t experiencing at the time). Tina then arrives home and her mother, Lana’s grandmother, Rosa, asks: “What did the doctor say, anyway?” Tina says Lana has a “common virus,” that she has an ear inflammation, but that it was best to wait. Rosa is surprised: “Why didn’t he prescribe an antibiotic?” Tina says she understood what Dr. Mauro said, but she can’t quite explain why Lana doesn’t need antibiotics. However, she says she trusts the doctor, and, ultimately, thinks it’s bad to give antibiotics unnecessarily.

Considering this, would it make sense to say that Tina *knows* her daughter has an ear infection? She certainly believes this, largely because she trusts Dr. Mauro. But does she know what an ear infection is? Does she know her daughter has a mild *acute otitis media* (AOM)?² Well, perhaps not, but she somewhat understood what the doctor explained. In any case, the attribution to Tina of the belief that her daughter has serous otitis depends on the fact that her daughter’s doctor concluded this during a clinical examination and communicated this information to her.

I think this example shows how knowledge spreads, forming a “web” of beliefs among many people, and how in a way it “circulates” among them in collectives with distinct thought styles, and that the attribution we make of beliefs and knowledge to these individuals depends not only on the strength and type of evidence, but also on contextual issues, including the bond of trust generated between people. Knowledge attribution also depends on the binding force generated by the “threads” of the web: the stronger these bonds, the more certain we are of the shared conviction among people within and between these distinct collectives. And this is independent of the fact that each person, within their collective, represents the subject through non-equivalent mental contents or talk about them with distinct terminologies. A plausible conclusion is that the concept of “otitis,” for example, although different when considered by the specialist, the physician, and the layperson, possesses a unity that is not reduced to the individual mental representations. On the contrary, the conceptual unity of “otitis” presents itself as a network of linked thoughts, uniting esoteric and exoteric representations.

Regarding AIDS, I suggest we accept that the fact that AIDS is caused by a virus is a piece of shared knowledge not only among scientists and doctors, but among the vast majority of ordinary people today. In this case, if we agree that even John knows that AIDS is caused by a virus, the question now

¹ The physician, in addition to relying on his experience, relied on recommendations from a clinical update literature popular among general practitioners, stating that “acute symptoms and signs of AOM” (of acute MEE) “usually resolve within three days, regardless of whether children are treated with antibiotics” (https://www.uptodate.com/contents/acute-otitis-media-in-children-treatment?search=bulging%20acute%20otitis%20media%20&source=search_result&selectedTitle=5~150&usage_type=default&display_rank=5). The recommendation is based on clinical evidence cited in the update text (Rovers et al., 2006; Thompson et al., 2013; Venekamp et al., 2023).

² Laypeople often use the term “acute” with a different meaning than doctors. “Acute” to a doctor is the opposite of “chronic.” Laypeople often call “acute” something strong or intense.

is how we can reconcile this statement with what is still much debated in epistemology in philosophical circles. After all, if knowledge is, as the so-called Traditional Theory of Knowledge (Ayer, 1956) argues, a true and justified belief, how can we attribute this to John? Would John have a true and justified belief that AIDS is caused by a virus? Well, John is incapable of providing us with an acceptable justification for why he believes this. But let's say he simply tells us: "Well, everyone knows this today, that AIDS is caused by a virus!" Wouldn't this statement be enough for us to affirm that he also knows this? After all, the following inference, borrowing a concept from Wilfrid Sellars (1953),

- It is now known that AIDS is caused by a virus.
- Therefore, John also knows that AIDS is caused by a virus.

This seems to be an example of a materially valid inference, because if we know that AIDS is caused by a virus, *ceteris paribus*, John (who is one of us) also knows.³

If we accept this material inference as valid, then the fact that John is unable to give us justifications for why he believes this does not serve to refute the conclusion that he knows this, just as the vast majority of people do today. Even so, we certainly wouldn't say that it is known something if there were no justifications for believing it. But these justifications need not represent mental states of John, or even of Dr. Mauro. Such justifications are not exactly opinions either. They are assertions, with different versions, details and complexities, which circulate among individuals, from different epistemic groups.

An additional question is: how does this "circulation" of beliefs and justifications among individuals from such different epistemic groups occur? One could argue that there is a dynamic in knowledge of this kind (which depends on scientific knowledge or discoveries) that begins with the activity of groups of scientists and only later establishes itself as something common. In the case of the knowledge that AIDS is a viral disease, this seems historically true. But couldn't laypeople influence the activity of scientists as well?

Consider the following chain of opinions on a topic that gained prominence during the recent COVID-19 pandemic.⁴ At the height of the pandemic, the WHO needed to provide precise guidance on certain preventive measures, including whether cloth masks could be used by the public as a protective measure against the virus. Thus, based on a study published in 2015 (<https://bmjopen.bmj.com/content/5/4/e006577>), which concluded that cloth masks increased the risk of infection, particularly from influenza, in healthcare workers, the WHO, on April 6, 2020, discouraged people in general from using cloth masks to reduce the risk of infection by the novel coronavirus (<https://apps.who.int/iris/handle/10665/331693>). This was the WHO's conclusion at the time:

Studies of influenza, influenza-like illness, and human coronaviruses provide evidence that the use of a medical mask can prevent the spread of infectious droplets from an infected person to someone else and potential contamination of the environment by these droplets.¹³ There is limited evidence that wearing a medical mask by healthy individuals in the households or among contacts of a sick patient, or among attendees of mass gatherings may be beneficial as a preventive measure.¹⁴⁻²³ However, there is currently no evidence that wearing a mask

³ In fact, there are two different material inferences related to the conclusion that John knows that AIDS is caused by a virus. First, the inference from "We know that AIDS is caused by a virus" and, second, the inference from "It is known that AIDS is caused by a virus". This second inference matters mostly for my argument.

⁴ I owe this example, and indeed much of these reflections, to my friend Ivan França Jr. In 2020, Ivan gave a brilliant presentation at the Unisinos Philosophy Colloquium, which took place during one of the most critical moments, with social isolation rules in full effect. In this presentation, based on Fleck, he demonstrated how the opinions of authorities and laypeople on the use of cloth masks changed throughout that year.

(whether medical or other types) by healthy persons in the wider community setting, including universal community masking, can prevent them from infection with respiratory viruses, including COVID-19.

The WHO's position was heavily criticized by people who didn't represent scientific groups, but by people we now call "social influencers." Many opinions circulated on social media, but the prevailing ones were those encouraging people to use homemade masks, particularly in a context of limited supply of professional masks (such as surgical masks and non-woven masks). Within a few months, the WHO changed its mind, to the point of rectifying its position with the following statement:

The use of masks is part of a comprehensive package of the prevention and control measures that can limit the spread of certain respiratory viral diseases, including COVID-19. Masks can be used either for protection of healthy persons (worn to protect oneself when in contact with an infected individual) or for source control (worn by an infected individual to prevent onward transmission) (WHO, 2020).

The true benefits of using cloth masks are still a topic of debate in the scientific community. However, as the WHO has highlighted, the evidence we had on the subject until 2020 was quite limited and of low quality. Even so, the conclusion reached in June was the opposite of that advocated in April. It could be argued that there was weak evidence that cloth masks could protect the general public from infection. Nevertheless, after some public debate, the health authorities became convinced that there was more reason to believe in an overall benefit, which led them to change their minds and begin encouraging people to use cloth masks, given the limited supply of professional masks. An interesting development was a publication at the end of 2020 (<https://bmjopen.bmj.com/content/5/4/e006577.responses#COVID-19-shortages-of-masks-and-the-use-of-cloth-masks-as-a-last-resort>) by some of the authors of the main randomized trial that informed the WHO's April 2020 opinion (MacIntyre, Seale & Dung et al, 2015), in which they explicitly acknowledge that the conclusion drawn from their 2015 study that cloth masks were harmful was of limited epistemic value:

It is important to note that some subjects in the control arm wore surgical masks, which could explain why cloth masks performed poorly compared to the control group. We also did an analysis of all mask wearers, and the higher infection rate in cloth mask group persisted. The cloth masks may have been worse in our study because they were not washed well enough – they may become damp and contaminated. The cloth masks used in our study were products manufactured locally, and fabrics can vary in quality. This and other limitations were also discussed. (...) If health workers choose to work using cloth masks, we suggest that they have at least two and cycle them, so that each one can be washed and dried after daily use.

Ultimately, the example shows that what we call "scientific knowledge" includes not only the mental states that philosophers strictly understand as "knowledge," since the body of scientific knowledge includes a web of statements supported by evidence of varying quality (and reliability). The specific recommendations made by experts are not, in effect, based on isolated statements, but on a set of statements supported by evidence of distinct qualities and relativized within contexts and considerations surrounding shared preferences and values. This is a fundamental element of what we ordinarily understand as *scientific* knowledge. This modality of knowledge appears to be a collective state, comprising states attributed to groups besides the particular *mental* states attributed to the individuals that comprise these groups.

The examples I presented above suggest that there is something incomplete in traditional conceptions of knowledge applied to what we call “scientific knowledge”, if we fail to consider the social role of justifications and the mutual dependence between epistemic groups. Before exploring this social character of general knowledge in more depth, it is necessary to revisit the foundations of classical epistemology and understand how belief, opinion, and knowledge have traditionally been distinguished.

► 2 Between belief, opinion and knowledge

The distinction between “believing” and “knowing,” as we know, dates back to the Socratic tradition. A common reference is Plato’s *Theaetetus*. Since then, traditional analysis has defined knowledge as justified true belief. It was only recently, with Edmund Gettier’s article (1963), that the traditional analysis began to be considered flawed. Gettier aimed to show that certain ordinary intuitions do not align with the traditional definition that justified true belief represent necessary and sufficient conditions for attributing knowledge to someone. His examples presented circumstances that satisfied the three criteria (S believes that p; p is true; and S is justified in believing that p), but which intuitively did not represent knowledge. One conclusion is that to claim that someone knows something requires more stringent conditions. But the examples I gave above show, on the contrary, that we are ordinarily less demanding than Gettier. John certainly knows that AIDS is caused by a virus, although he cannot provide justification for it. And something can be a piece of knowledge in a general sense even if some people don’t believe it or even when the justifications for asserting it are still limited. As Guyatt and Djulbegovic (2015) argue, in today’s medicine, what we include in the list of knowledge includes statements or conclusions to which we give greater or lesser assent, supported by different types of evidence, some more, some less, “strong.”

There is a certain discrepancy between what we ordinarily call opinion, belief, and knowledge and the concepts employed in analytic philosophy. Peter Hacker published a short article (2004) that shows us that our ordinary uses of “belief,” “opinion,” and “knowledge” differ from the concepts stipulated by analytic epistemologists (Hacker, 2002). Ordinarily, believing in something is weaker than knowing about it. Therefore, by saying that I believe something, I may be illocutionarily implying that I see good reasons for believing it, admitting, nevertheless, that such reasons are not yet strong enough to convince everyone. But when someone says they know something (viz. that AIDS is caused by a virus), they are saying, on the contrary, that it is not a mere opinion or even just a “belief.” In the context of already recognized knowledge, if someone says their opinion is different (that AIDS is not caused by a virus), one of the consequences is that we do not take that opinion seriously. In other words, it’s hard to believe that this person doesn’t actually believe this. This could imply that they actually know that AIDS is caused by a virus, even if they vociferously argue against it. As has been shown since Freud, what someone brings to their consciousness or even what someone expresses as their opinion may, in some situations, not even represent what that person actually believes or knows. Even though, much of classical knowledge theory shares the mentalist assumption that knowledge is, fundamentally, an internal state of the subject—a true belief supported by justifications accessible to one’s consciousness or reflection. This conception was profoundly shaken by the emergence of externalist critiques. Externalism, broadly speaking, rejects the thesis that the conditions for someone to know something must be exclusively internal to the subject. Justifications can depend on external factors—such as the reliability of belief-formation processes—and still constitute knowledge.

One of the most influential formulations of externalism is reliabilism, defended by authors such as Alvin Goldman (1986). Reliabilism holds that a true belief is knowledge if it was produced by a reliable

cognitive process, even if the subject lacks introspective access to that reliability. Thus, John can know that AIDS is caused by a virus if this belief was formed reliably—for example, by hearing it from trustworthy sources, such as doctors, teachers, or well-informed media outlets—even if he himself cannot explain why this belief is true or justify why their sources are reliable. What matters, in this case, is that the process that generated John’s belief tends to reliably produce true beliefs with sufficient frequency.

This perspective challenges the internalist requirement of reflective access to justification, showing that many of our ordinary attributions of knowledge—including those made to ordinary people—seem to dispense with such access. Let’s return to Tina’s example: she believes her daughter has an ear infection, and her belief is true and was formed reliably—she heard it from an experienced pediatrician who examined the child. From a reliabilist perspective, Tina knows this, even if she cannot explain how the doctor arrived at such a diagnosis. Her knowledge, in this case, depends on the reliable integration between different agents and epistemic institutions.

With this, we see the emergence of a non-mentalist conception of knowledge, in which some epistemic states need not necessarily be “inside the head” of an individual subject. This represents an important conceptual opening for understanding how knowledge can be distributed, shared, and sustained socially. This transition from an epistemology centered on the individual to one centered on intersubjective processes and institutional practices is what begins to pave the way for the conception of knowledge as a social epistemic state, which will be explored below.

► 3 Reliabilism, evidentialism and Comesaña’s hybrid proposal

An example of knowledge widely recognized as authoritative is medical knowledge. Medical knowledge, nevertheless, is anchored in multiple sources supported by evidence of varying methodological quality. As a result, many began to argue that doctors needed to adopt a critical view of what they considered established knowledge. One of the most influential initiatives to establish rigorous criteria for evaluating medical evidence is the Evidence-Based Medicine (EBM) movement, developed since the 1990s. EBM (in what we could understand today as its most mature form) proposes that clinical decisions be based on the best available evidence, especially that derived from controlled clinical trials and systematic reviews. Although this proposal has gained widespread support, it has also been criticized, particularly for its potential reductionism (Loughlin, 2009) and for neglecting contextual and ethical aspects of clinical practice (see, for example, Tonelli, 2006; Goldenberg, 2009; also Miles, Polychronis & Grey, 2006).

At first glance, one might think that evidentialism represents the natural epistemological foundation for the EBM movement. After all, EBM proposes that clinical decisions be made based on the best available evidence, which seems to align with the idea that a belief is justified if supported by good evidence. However, this approach is not as straightforward as it seems.

Evidentialism, formulated in its most influential version by Richard Feldman and Earl Conee (1985), holds that the epistemic justification of a belief depends exclusively on the evidence the subject possesses. In other words, what makes a belief justified is the degree to which it is supported by evidence accessible to the agent. However, this principle faces practical limitations, especially in the clinical context. Many healthcare professionals do not directly access primary studies, personally evaluate the data, or independently apply the statistical criteria that support the hierarchy of evidence. They rely on consensus, protocols, systematic reviews, and institutional guidelines.

Authors such as Gordon Guyatt—who coined the term “evidence-based medicine”—and Benjamin Djulbegovic have recognized the tension between the evidentialist ideal and actual clinical practice, which involves varying levels of access, understanding, and application of scientific evidence. The practice of

EBM relies on a collective structure of validation and dissemination of knowledge, suggesting that the epistemology of EBM can be better understood from an externalist and social perspective.

Guyatt and Djulbegovic, in fact, suggest that a hybrid approach better expresses the epistemological foundations of EBM. It takes not only evidence to assert something epistemically sustainable in medicine, but also reliable evidence. A hybrid approach is the one proposed by Juan Comesaña (2020). He proposes an interesting articulation between reliabilism and evidentialism, seeking to overcome the rigid opposition between the two approaches. Comesaña argues that an epistemic system can be reliable precisely because it structures, filters, and distributes evidence in a way that facilitates the formation of true beliefs among its members. Rather than requiring every agent to have direct access to primary evidence, what matters is that they be embedded in a reliable epistemic environment, in which evidence is properly processed and communicated.

This proposal is particularly useful for explaining how knowledge circulates in complex societies. Rather than relying exclusively on individual epistemic capacities, knowledge in a collective is sustained by networks of trust, reliable institutional practices, and social mechanisms that validate and disseminate information. Comesaña's evidentialist reliabilism thus offers a promising theoretical framework for understanding how specialized knowledge becomes common knowledge, even when the justification is not directly accessible to all individuals. He combines Goldman's reliabilism with Feldman and Conee's evidentialism into a single approach. However, regardless of these efforts, the person who truly best presented the foundations for understanding medical knowledge as social knowledge was the Polish physician and philosopher of science Ludwik Fleck.

► 4 Fleck and the social ontology of knowledge

The conception of knowledge as something socially sustained gains depth in the work of Ludwik Fleck, especially in his book *Genesis and Development of a Scientific Fact*, originally published in 1935 (1979). Fleck proposed that scientific knowledge is not the exclusive product of individual subjects, but results from a collective process of construction within a "thought collective" — a group of individuals who share the same style of thinking, that is, a common way of perceiving, interpreting, and validating phenomena.

For Fleck, a scientific fact is not born as such, but is constituted throughout a process of conceptual stabilization within shared practices, institutions, and languages. This process involves phases of resistance, adjustment, and sedimentation, in which different versions of a phenomenon are gradually tested, criticized, reconfigured, and ultimately adopted as valid by an epistemic community. Scientific "truth," in this view, is inseparable from the history of its constitution, as well as from the specific modes of observation, inference, and validation cultivated by each collective. Fleck faced harsh criticism in his time, as the assertion that scientific facts are constituted through practices and are not objective elements of the world independent of both human subjectivity and activity contradicted, and for some still contradicts, common scientific opinion. But perhaps this stems from a misunderstanding. For, on the one hand, several philosophers agree that the way we describe the world depends on the conceptual framework we employ to describe it, and this framework is not immutable and fundamentally depends on our conceptual practices. Thus, it is quite plausible that the facts we describe as "scientific" facts can only be stated within the scope of the scientific practices that make them amenable to description. It makes no sense to talk about viruses without starting from the conceptual framework of modern microbiology. Atoms have been described differently over time, from Democritus to modern quantum

physics, and it is quite likely that the entities and the states of affairs described according to these conceptual schemes are not equivalent.

If this makes sense, then what is understood as scientific knowledge is not equivalent to any individual state or set of mental states, nor is it “universal” in the sense attributed by classical objectivist doctrines. Scientific knowledge is, by its social nature, contextual and historically situated. It includes the mental states of individuals within a specific or general collective, but is not limited to them. Scientific knowledge fundamentally emerges from the interaction between subjects, practices, and institutions, and is maintained through processes of transmission and consolidation that transcend the immediate experience of individuals. In other words, in the domains of science (and perhaps in several other domains too), what we call “knowledge” is often the result of shared beliefs and practices, transmitted and validated collectively.

This perspective broadens and deepens the externalist thesis, as it not only shifts justification beyond the individual mind but also inserts knowledge (at least scientific knowledge) into specific sociocultural dynamics. When we affirm that John knows that AIDS is caused by a virus, we are affirming that he shares the epistemic positions common to individuals in his thought collective (in this case, the more general collective represented by laypeople), whose thought style differs from that of other groups represented by experts and scientists (virologists, for example), even though there are connections between the thoughts of all these groups. This dynamic interaction between the thought disseminated among distinct thought collectives is based on institutions trusted by all groups and is sustained by intersubjective validation practices that operate within and between these groups. With Fleck, we understand that common knowledge is not an impoverished version of scientific knowledge, but rather one of its most robust expressions: one that stabilizes and circulates across different layers of society, shaping not only what we know, but also how we live, decide, and act.

► 5 Epistemic dynamics and the transition between thought collectives

Fleck’s ideas offer a decisive starting point for understanding knowledge not as an isolated mental state, but as shared states of mind that result from a process historically constituted, and sustained by collective practices. From this, it is possible to go further: not only recognizing that different social groups share their own styles of thought, but also that there is transit, exchange, and disputes between these styles. Knowledge does not remain restricted to a single thought collective—it circulates among them, undergoes reinterpretation, and sometimes radically transforms when appropriated in new contexts.

This “traffic” is particularly evident in contemporary societies, marked by intense specialization of knowledge and, at the same time, by mechanisms for the widespread dissemination of information. Knowledge produced in research centers and by experts becomes part of media discourse, educational practices, institutional protocols, and even everyday conversations. This process is not neutral: it requires conceptual translation, simplification, adaptation, and, in many cases, the overcoming of resistance. Thus, what is disseminated is not just information, but a set of specific forms of interpretation and trust—knowledge that becomes common because it is transmitted, validated, and reiterated by sources considered reliable.

Circulation between epistemic collectives is facilitated by social institutions such as schools, the media, health services, and legal systems, all of which operate as vectors that make certain knowledge accessible and legitimate to broader audiences. The case of John—who knows that AIDS is caused by a virus without being able to justify it in technical terms—illustrates this point precisely: his knowledge

is sustained by a chain of institutional and social trust that connects laboratories to school textbooks, newspapers to conversations with health professionals. It is a distributed epistemic network.

This model allows us to understand why certain types of knowledge stabilize as “common knowledge” even if some of its foundations remain inaccessible to most people. This is not a degradation of knowledge, but rather a reconfiguration of its form: common knowledge is the result of the stabilization, through circulation and repetition, of initially restricted epistemic contents, which become plausible and reliable through shared validation practices. Knowledge, in this case, is less a matter of individual cognitive mastery and more of socially embedded and enacted states and practices.⁵

This conception has important philosophical consequences. It challenges epistemological models centered on the autonomous subject and the internal examination of justifications, proposing a relational, distributed, and dynamic approach to knowledge. It shows that knowledge is not something possessed in isolation, but something shared—and that the ability to share depends on trust, a common language, and institutions that make it possible. By shifting the focus of epistemology from the interior of the mind to the realm of collectives, and from these to the interior of historically situated societies, we take a decisive step toward understanding how knowledge ultimately circulates.

► 6 The metaphysics of social knowledge

The reliabilist proposal, especially in Alvin Goldman’s formulation, establishes that the epistemic justification of a belief depends on its being produced by a reliable process—that is, a process that, under normal conditions, tends to produce true beliefs. In contrast to evidentialism, which requires the subject to have access to the evidence that justifies their belief, reliabilism holds that the objective reliability of the process is sufficient for there to be knowledge, even if the subject cannot introspectively retrieve the reasons for their belief.

This seems to explain how people like John—who cannot explain why AIDS is caused by a virus—can still know that this is true, as long as their beliefs originate from reliable epistemic processes: listening to a doctor, trusting health institutions, studying through peer-reviewed educational materials. However, Goldman’s classic reliabilist model struggles to explain how the recognition of new scientific knowledge is constituted. After all, an epistemic process can only be considered reliable if there are, first, shared criteria for evaluating it as such. But if every criterion is a shared item, how could a single individual confidently certify that their claims satisfy them? It seems, therefore, that any certification requires peer recognition.

Let’s imagine that Nancy is a scientist who, during the COVID-19 pandemic, conducted a systematic review of studies on the use of cloth masks as a preventive measure against the spread of the novel coronavirus. Her study, although methodologically limited because it was largely based on observational studies, concluded that there was moderate evidence of benefit in the protection offered by cloth masks in community settings. Even so, at the time, her conclusion lacked the endorsement of health authorities or recognition among her peers as sufficiently robust evidence.

Let’s assume that Nancy was correct and that a well-designed experimental study would reach the same conclusions as her meta-analysis of observational studies. However, even if her conclusion were correct, her belief—that cloth masks would help reduce transmission—still could not be considered, strictly speaking, a case of scientific knowledge. This is because the epistemic status of her study’s conclusion as social (and therefore recognized) knowledge also depends on a collective validation

⁵ I do not intend here to argue for the connection between the thesis I am supporting and the view defended by enactivism in philosophy of mind, but I think that there is certainly conceptual proximity between this view and the social conception that I am defending here.

process: peer review, institutional acceptance, and integration with the broader body of available evidence. Only after public debate, the accumulation of new evidence, and the shift in positioning by international organizations could the statement “cloth masks help reduce contagion” be recognized as reliable knowledge.

The internalist would argue that, in this case, Nancy gained knowledge on the topic, advancing what was known before her study was completed. For the externalist, however, it is questionable to assert that Nancy knows that cloth masks protect against COVID-19 if all her peers still disagree. Perhaps, the externalist would argue, her conclusion qualifies for recognition as knowledge in the sense that it is based on reliable methods and presents a more robust justification than previous studies—and is therefore potentially closer to the truth than the opposing claims of her colleagues. This is even if she and her peers acknowledge methodological limitations in their review. Science embraces this heterogeneity: we are more convinced of some conclusions than others, but all claims—both those supported by robust evidence and those supported by weaker foundations—remain epistemically linked. Together, they comprise what is recognized as “knowledge” in a given field. Rather than considering this set as mere individual beliefs, it is more appropriate to understand it as a web of interconnected claims that seek to progressively approximate the truth—a network of interconnected beliefs, in which individual claims support each other, directly or indirectly (Quine & Ullian, 1970). These claims are based on processes and sources with varying degrees of reliability, whose methodological transparency is regulated by epistemic practices shared within a “thought collective,” as Fleck proposed.

The evolution of this system is dynamic: beliefs recognized at a given moment can be revised, restructuring individual and collective knowledge. Perhaps Thomas Kuhn is right in defending the notion of “normal science” as a period of consensus and accumulation of results under a single paradigm (Kuhn, 1962). Yet, even within these periods of stability, there is room for debate, internal revisions, and epistemically significant dissent. The example I gave, although hypothetical, aims to show that, even in a reliabilist model, what becomes disseminated within a collective as knowledge depends on intersubjective recognition that often transcends and even depends on acceptance in other circles. As Fleck argued, what is accepted and established as knowledge in an esoteric circle of scientists also suffers the effects of the reception of these same ideas in other circles, including exoteric ones (as shown by the example of the debate over the benefits of cloth masks). It’s not enough for a single individual to be right. Furthermore, it’s not enough for their correct belief to be subjectively supported by a reliable source. Their correctness must also be reliably recognized as such within a collective and dynamic process of evaluation, and this recognition plays a constitutive role in transforming beliefs into knowledge.⁶

► 6 Conclusion: social epistemology revisited

The analysis developed thus far allows us to propose, in an exploratory manner, a series of statements that articulate a relational and distributive conception of knowledge:

- 1) Individual knowledge depends on reliable collective practices of collective knowledge production, but does not require (in the case of individual knowledge of collectively known statements) individual introspective access to the justifications for these beliefs;
- 2) human knowledge diffuses and circulates among distinct epistemic collectives, with different styles of thought (as Fleck argued);

⁶ I think this is in tune with Wittgenstein’s view on the agreement between judgments as a normative criterion of correctness for individual actions.

- 3) scientific knowledge is something essentially collectively produced and mastered, historically situated and sustained by practices and institutions;
- 4) the distinction between belief and knowledge involves the role of epistemic trust and the intersubjective stabilization of beliefs in human collectives;
- 5) peer recognition is constitutive of the epistemic status of a belief as scientific knowledge;
- 6) reliabilism, combined with evidentialism, allows for a plural approach to valid forms of justification in different contexts.

These ideas suggest that an epistemology adequate to knowledge in complex societies must be sensitive to the collective dimension of justification. This does not imply relativism, but it does require attention to how truth, justification, and trust operate in concrete social networks. Knowledge not only circulates: it transforms, stabilizes, and is socially instituted. It is, therefore, an epistemic reality that is simultaneously ontological, institutional, and pragmatic.

► References

ROVERS, M. M.; GLASZIOU, P.; APPELMAN, C. L.; BURKE, P.; McCORMICK, D. P.; DAMOISEAUX, R. A.; GABOURY, I.; LITTLE, P.; HOES, A. W. 2006. Antibiotics for acute otitis media: a meta-analysis with individual patient data. *Lancet*, 368(9545): p. 1429-35. doi: 10.1016/S0140-6736(06)69606-2. PMID: 17055944.

AYER, A. J. 1956. *The Problem of Knowledge*. London: Macmillan.

COMESAÑA, J. 2020. *Being Rational and Being Right*. Oxford: Oxford University Press.

FELDMAN, R.; CONEE, E. 1985. Evidentialism. *Philosophical Studies*, 48: 15–34.

FLECK, L. 1979 [1935]. *Genesis and Development of a Scientific Fact*. Chicago: University of Chicago Press.

GETTIER, E. L. 1963. Is Justified True Belief Knowledge? *Analysis*, 23(6): 121–123.

GOLDENBERG, M. J. 2009. Iconoclast or Creed? Objectivism, Pragmatism, and the Hierarchy of Evidence. *Perspectives in Biology and Medicine*, 52(2): 168–187.

GOLDMAN, A. 1986. *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.

GUYATT, G.; DJULBEGOVIC, B. 2015. Evidence-Based Medicine and the Theory of Knowledge. In: GUYATT G., RENNIE D., MEADE M.O., COOK D.J. (eds.) *Users' Guides to the Medical Literature: A Manual for Evidence-Based Clinical Practice*, 3rd ed. McGraw-Hill Education. Disponível em: <https://jamaevidence.mhmedical.com/content.aspx?bookid=847§ionid=69031459>

HACKER, P. M. S. 2002. Philosophy: A Contribution, Not to Human Knowledge, but to Human Understanding. *Royal Institute of Philosophy Supplement*, 51: 1–28.

HACKER, P. M. S. 2004. On the Ontology of Belief. In M. Siebel & M. Textor (Eds.), *Semantik Und Ontologie: Beiträge Zur Philosophischen Forschung* (pp. 2–185). Frankfurt: Ontos Verlag.

LOUGHLIN, M. 2009. The Basis of Medical Knowledge: Judgment, Objectivity and the History of Ideas. *Journal of Evaluation in Clinical Practice*, 15(6): 935–940.

- MacINTYRE C. R.; SEALE, H. DUNG, T. C.; *et al.* A cluster randomised trial of cloth masks compared with medical masks in healthcare Workers *BMJ Open*, 5:e006577. doi: 10.1136/bmjopen-2014-006577.
- MacINTYRE, C. R., et al. 2015. A cluster randomised trial of cloth masks compared with medical masks in healthcare workers. *BMJ Open*, 5(4): e006577.
- MILES, A.; POLYCHRONIS, A.; GREY, J. 2006. The Evidence-Based Health Care Debate—2006: Where Are We Now? *Journal of Evaluation in Clinical Practice*, 12(3): 239–247.
- SELLARS, W. 1953. Inference and Meaning. *Mind*, 62(247): 313–338.
- THOMPSON, M.; VODICKA, T. A.; BLAIR, P. S.; BUCKLEY, D. I.; HENEGHAN, C.; HAY, A. D.; TARGET Programme Team. Duration of symptoms of respiratory tract infections in children: systematic review, *BMJ*. 2013 Dec 11;347:f7027. doi: 10.1136/bmj.f7027. Erratum in: *BMJ*. 2014;347:f7575. PMID: 24335668; PMCID: PMC3898587.
- TONELLI, M. R. 2006. Integrating Evidence into Clinical Practice: An Alternative to Evidence-Based Approaches. *Journal of Evaluation in Clinical Practice*, 12(3): 248–256.
- VENEKAMP, R. P.; SANDERS, S. L.; GLASZIOU, P. P.; ROVERS, M. M. 2023. Antibiotics for acute otitis media in children. *Cochrane Database Syst Rev*, 11(11): CD000219. doi: 10.1002/14651858.CD000219.pub5. PMID: 37965923; PMCID: PMC10646935
- WORLD HEALTH ORGANIZATION. 2020. *Advice on the use of masks in the context of COVID-19: interim guidance*. 6 April 2020 & 5 June 2020. Disponível em: <https://iris.who.int/handle/10665/332293>

Declaração de Disponibilidade de Dados:

Todo o conjunto de dados que dá suporte aos resultados deste estudo foi publicado no próprio artigo.

Submetido em 06 de agosto de 2025.

Aceito em 03 de março de 2026.