Synthese (2026) 208:124
https://doi.org/10.1007/s11229-026-05700-z

ORIGINAL RESEARCH

# Gatekeeping in science reporting: a zetetic account

**Vanja Subotić**[1] · **Ivan Umeljić**[2]

Received: 31 July 2025 / Accepted: 23 June 2026
© The Author(s), under exclusive licence to Springer Nature B.V. 2026

## Abstract

We argue that gatekeeping in science reporting mediates an adequate epistemic relation between non-experts and experts. We further specify how this unfolds in non-ideal epistemic situations, e.g. through recent reports on the alleged sentience of large language models and the prospects of artificial general intelligence. In our view, the primary goal of science reporting is to promote epistemic improvement by better positioning laypeople in relation to the inquiry of relevant scientific questions as elaborated within zetetic epistemology. Consequently, the task of science journalists involves not merely scaling up facts or harnessing settled attitudes toward science, but rather navigating the balance between do-your-own-research (Ballantyne 2024) and public trust in science. Thus, trust in science should not be conceived as a synchronic epistemic attitude we owe to experts unconditionally, but as an outcome of a dynamic diachronic process intertwined with a plethora of interrogative attitudes such as curiosity, wonder, and deliberation, that cannot be reduced to a binary trust vs. distrust relation. We propose norms of good science reporting that would fulfill the gatekeeping ambition, such as (1) the clutter avoidance for asking and considering questions, (2) the elicitation and maintenance of vigilant attention, and (3) the prevention of interrogative injustice. These norms stem from the concept of zetetic relevance—that is, the identification of information needed to provide reasonable answers to the questions raised by science reporting (Siegel 2026).

**Keywords** Artificial general intelligence · Gatekeeping · Inquiry · Large language models · Science journalism · Science reporting · Trust · Zetetic epistemology

Vanja Subotić and Ivan Umeljić contributed equally to this work.

✉ Vanja Subotić
vanja.subotic@f.bg.ac.rs

✉ Ivan Umeljić
iumeljic@cpn.rs

[1] Institute of Philosophy at the Faculty of Philosophy, University of Belgrade, Čika Ljubina 18-20, Belgrade, RS 11000, Serbia

[2] Center for the Promotion of Science, Kralja Petra 46a, Belgrade, RS 11000, Serbia

Published online: 24 August 2026

## 1 Introduction

As epistemic agents, we acquire epistemic resources, i.e., individual items of knowledge (Pohlhaus, 2012). To some extent, these are social products. This means that society influences our epistemic position, for better or for worse. In knowledge-conducive milieus, citizens can acquire more epistemic resources and means to assess their reliability. Thus, it seems that investing in the epistemic improvement of citizens, by increasing trust in science and encouraging a satisfactory division of epistemic labor, would be beneficial, provided that having a better epistemic position is good for us. In a recent preregistered 68-country survey, Cologna et al. (2025) found that in most countries, most citizens do, in fact, trust scientists despite the temporary crisis during and after the COVID-19 pandemic. However, their survey respondents also thought that scientists should engage more with society and communicate their findings more effectively when they receive grants, so that they could feel there is a reciprocal relationship between experts and the lay audience. But scientists often have better things to do: actually doing science. This is why investments in the epistemic improvement of laypeople include but are not limited to high-quality science reporting (Contessa, 2023, 2025). The same dataset used in Cologna et al. (2025) has been used to describe and assess science communication strategies. Thus, in democratic-corporatist media systems (Hallin & Mancini, 2004), traditional media and forms of reporting are more frequently used than social media for exposure to scientific information (Mede et al., 2025).[1] Science journalists, like any other journalists, always perform a public act of assertion for the purpose of informing their audience (Simion, 2017). Moreover, we contend that *gatekeeping* constitutes the defining epistemic function of science journalism, distinguishing it from the mere dissemination of information or fact-checking. Thus, as gatekeepers between science and society, science journalists should put an emphasis on *public interest*. Consequently, it is their epistemic duty to report on the social and safety implications of science and technology, as well as to cover the political, economic, and ethical dimensions of scientific and engineering work, including potential risks and hazards (Zeller Jr, 2022). A prominent example of such watchdog journalism was the *Knight Science Journalism* (KSJ) *Tracker*, a pioneering science media review and criticism blog that unfortunately ceased publication in 2015. In the words of one of the renowned science journalists, Deborah Blum, their job is not to foster the positive image of science or engage in science marketing but to

[1] We use the term "science reporting" broadly to include not only science-focused magazines, but also general-interest outlets and news media that publish stories on scientific topics, ranging from online platforms to tech blogs and public broadcasters. Siri Carpenter, a science journalist and author of a leading guidebook on science writing (Carpenter, 2024a), advocates for such an approach by emphasizing that treating science journalism as an elite form of content written by highly specialized journalists and aimed only at readers already familiar with and interested in science is misguided and dangerous (*cf.* Carpenter, 2024b).

"portray research accurately in both its rights and its wrongs and stand unflinchingly for the integrity of the story …" (2021: p. 323).[2]

Rather than treating gatekeeping as an epistemic virtue of an individual epistemic agent (Kwong, 2023), we regard it as a functional way to perform one's epistemic task. In other words, what makes a science journalist a gatekeeper is not an instantiated epistemic virtue, but the improvement of the layperson's epistemic position through the mediation of their relationship with experts. Of course, we are aware that for consequentialists in virtue (and vice) epistemology, the mere performance of epistemic tasks would suffice to discuss gatekeeping *qua* epistemic or intellectual virtue, as opposed to the adversarial motivational accounts, which insist on the additional motivating force or inclination to forge the link between the existence of a particular epistemic trait and its practical realization. Nonetheless, our preferred theoretical framework in which we embed all further analysis of gatekeeping is *zetetic epistemology*, for which we provide a rationale in the rest of the paper.[3] That said, we still recognize the crucial role of science journalists' epistemic competence. Along with Fairbank (2025), we contend that good science reporting usually originates from epistemically competent science journalists, but no amount of knowledge about scientific facts or technology will facilitate the acquisition of corroborating information or the making of other informed choices about the content presented to the public. To put it in philosophical jargon, scientific literacy is a necessary but not sufficient condition for high-quality science reporting.[4]

For this reason, because it mistakenly assumes that an audience's keen interest in science results in epistemic deference, a model of science reporting that relies on the mere compilation of facts and unconditional trust is untenable. Such a model proves especially challenging in the case of emerging technologies, such as large language models (LLMs), and the prospects of artificial general intelligence (AGI), where development is led by private industry. Covering these topics according to the model sketched above would require journalists to get themselves acquainted with the so-called *future-oriented* research programs (Feest, 2017), whose research questions, hypotheses, and methods are in constant flux. This can prove to be a tedious and

[2] Blum is one of the professionals behind the *Undark Magazine*, a non-profit, editorially independent online magazine founded under the auspices of the *Knight Science Journalism Program* at MIT. *Undark* revitalized the spirit of the *KSJ Tracker*. Blum and colleagues claim in the Mission statement that "[a]s journalists, we recognize that science can often be politically, economically, and ethically fraught, even as it captures the imagination and showcases the astonishing scope of human endeavor. *Undark* will therefore aim to explore science in both light and shadow, and to bring that exploration to a broad, international audience." This form of science reporting is exactly what we aim to endorse in our paper, where we put forward gatekeeping norms rooted in zetetic epistemology to help foster such journalistic standards.

[3] We are aware of the most recent attempts to link virtue and zetetic epistemology (Carter & Willard-Kyle, forthcoming), but we leave this to be further probed by the philosophical community before trying to apply it in the context of science reporting and the epistemic role of science journalists.

[4] For a critical overview of the concept of scientific literacy, see Rudolph (2024), who traces its historical development and argues that it has become an empty slogan that should be abandoned by science education. Here, we use the term in a noncommittal way, simply to indicate that science journalists need not have deep specialization in the fields they report on, but should possess a basic level of understanding required for responsible and informed coverage (that is, their epistemic competence in this sense makes them more capable to report on science than laypeople are to pursue solo inquiry.).

borderline irresponsible approach. Such research programs are recognized as *socially disruptive* due to their repercussions on the socio-economic order and concomitant conceptual disruption (Löhr, 2023). Improving the AI literacy of science journalists does not guarantee high-quality reporting but may, in fact, exacerbate existing biases by overlooking ethical, legal, and social concerns or favoring private-sector interests above those of the public. Put simply, the more disruptive the research at the intersection of science and technology, the heavier the gatekeeping burden for science journalists. This mirrors investigative journalists confronting corruption: both roles demand exposing hidden systemic risks, where low-quality reporting does not merely misinform but also harms public trust and democratic deliberation (Schipani, 2024). However, while journalists are trained to be skeptical of politicians, they often treat scientists with what Murcott (2009) describes as the deference of a priest toward a deity. This "priestly" dynamic explains why many science journalists simply relay claims from scientists and corporate public relations officers, spreading the aura of objectivity that discourages the critical scrutiny and questioning typical of seasoned political reporting. Ultimately, this deference sustains the untenable model of science reporting described above, as it relies on the unconditional trust of the audience.

The danger of this model lies in the very nature of trust itself. As Nguyen (2022) suggested, putting one's trust in something could indeed be conceived as an *unquestioning* attitude. This means that we are ready to halt any form of control, monitoring, probing, or doubting because we are willing to put our faith into something by relying on a particular epistemic resource and deferring to those who know best, namely, experts. We have suspended deliberation about the reliability of such an epistemic resource, renounced our epistemic independence, and made ourselves vulnerable (Baier, 1986). How to pick the right expert among many who disagree (Goldman, 2001)? How to trust in their expert opinion when they are, at the end of the day, ordinary people who struggle to obtain funding, hold different values, have numerous beliefs, some even irrational or contradictory ones? It is important for laypeople to "do their own research" and identify experts who *deserve* their trust because, otherwise, they would be endowed with a corpus of dogmatically entrenched beliefs (Douglas, 2009, Grasswick, 2010, Hardin, 2006, Hardwig, 1991, Rolin 2020, Levy, 2007, 2021, 2024)., Trusting experts *qua* intellectual authorities too much is the very opposite of what we expect science to teach us and contradicts most of the history of science. Additionally, something must alleviate the intellectual load of laypeople due to limited time and cognitive capacities.

Science journalists can help navigate the *balancing* act of combining public trust and solo inquiry (Rabenberg, 2024, John, 2018) through the complex mechanisms for selecting and producing scientific content. Rather than merely reinforcing unquestioning, settled attitudes and encouraging a passive deference to authority, the journalist's task is to cultivate questioning or interrogative attitudes within the public. This is why we develop a zetetic outlook on gatekeeping in science reporting and propose three norms for science journalists to take into consideration. In contrast to traditional epistemology, the central notion of zetetic epistemology is inquiry (Hintikka, 1981, 2007; Hookway 2006; Friedman, 2013, 2019, 2020; Thorstad, 2021), *cf*. Falbo (2021)—including the initial curiosity and mind wandering, formulation of questions, and the process of seeking answers and resolving such questions. The

norms we propose are entrenched in *zetetic relevance*, i.e., identifying what could count as a reasonable answer to questions science journalists tackle (Siegel, 2026). We apply them to the case of science reporting on LLMs and AGI, given the heightened public interest in the real and anticipated risks these technologies entail, and in light of the increasingly salient role of gatekeeping in shaping public reception.

The structure of the paper is as follows. In Sect. 2, we examine the relationship between science and society by focusing on the social production of epistemic resources and the division of epistemic labor, thereby laying the groundwork for situating science reporting as a form of public service aimed at the epistemic improvement of citizens. Section 3 introduces and elaborates the zetetic turn in epistemology and its connection to science journalism. Section 4 develops three norms of gatekeeping in science reporting grounded in zetetic relevance: (1) clutter avoidance in question-asking, (2) the elicitation and maintenance of vigilant attention, and (3) the prevention of interrogative injustice. These norms are subsequently applied to the case of science reporting on AI hype, which refers to media-driven exaggerations about the capabilities and performance of LLMs, as well as the plausibility of AGI. This approach leads us to suggest concrete recommendations for how science journalism can "stand unflinchingly for the integrity of the story." Finally, Sect. 5 summarizes the key points and addresses some of the objections to our proposal.

## 2 Relying on the right minds: division of epistemic labor in the service of society

In the social epistemology of science, discussions of the division of epistemic labor have typically focused on analytic and formal models that examine how scientists choose research projects and topics, communicate findings, or how scientific communities should be structured to promote progress (Kitcher, 1990; Strevens, 2003; Zollman, 2007; Weisberg & Muldoon, 2009; Borg et al., 2019). The aim of social epistemologists of science has been to identify the optimal and most efficient ways to divide epistemic labor among idealized groups of agents within a given research community. However, constructing a formal model that is optimized for internal scientific efficiency, such as maximizing the speed of discovery or the volume of publications, tells us little about whether that performance genuinely benefits society, brings us closer to the truth, or reflects a just allocation of epistemic rewards and accountability (Winsberg et al., 2014; Rolin, 2015; Longino, 2022). These concerns regarding the social and ethical limitations of formal models prompted a shift in focus toward the division of epistemic labor as it occurs at the demarcation line between experts and laypeople (Kitcher, 1993; Goldman, 1995, 1999).

Science communication once followed the same demarcation line. According to *the deficit model* of science communication, information about science flows vertically from scientists and science journalists to the public (Miller, 1998; see; Goldenberg, 2021 for an overview). The assumption is that if the public simply understood the science better, they would naturally come to agree with scientific consensus and support technological progress. In other words, the aim of the deficit model was to address the epistemic and cognitive asymmetry between experts and laypeople

through education and promotion of scientific literacy. However, the science communicator was more often cast as a savior of the gullible who relays claims from a position of authority, in line with Murcott's (2009) priestly dynamic described in the **Introduction**, thereby presenting a somewhat condescending image of the relationship between science and society. Even after several years of criticism and refutation, both the deficit model and the clear-cut division of epistemic labor persisted (see Cortassa, 2016; Simis et al., 2016).

The deficit model hinged on the assumption of epistemic stability of the clear-cut division of epistemic labor. In the epistemically stable world, experts are objective, unbiased, and closer to the truth with every insight and experiment, whereas the grateful public gets more knowledgeable with every new piece of information and, therefore, better at decision-making. Most importantly, laypeople presumably deem experts trustworthy precisely because of their objectivity and the gradual approach to the true description of the world, whatever that may be. Our world is, however, more often epistemically unstable than stable. Each instance of scientific fraud, academic misconduct, or peer disagreement on controversial issues, such as COVID-19 vaccination, contributed to epistemic instability by undermining public trust in experts, even though such exposures are integral to the self-regulation of science (Redish et al., 2018; Williams et al., 2023). This meant that experts are shown to be far from objective and unbiased; their experiments and insights might sometimes serve only to advance one's career rather than to seek the truth, and scientific consensus is not easy to achieve. The replication crisis in the life sciences, and most notably in psychology, further fueled doubt and distrust in science as an institution (Ioannidis, 2005; Hendriks et al., 2020). From the perspective of laypeople, it is not that they lack education and scientific literacy, but rather that structural vulnerabilities in the socio-epistemic infrastructure of science suggest that passive deference to experts may not be the best way to obtain knowledge. Of course, no source of knowledge is entirely unsullied: both experts and laypeople remain fallible in how they cognize the world. This epistemic instability has, however, spurred something more dangerous than mere realization of epistemic fallibility, namely the do-your-own-research (DYOR) approach to science, which is based on the conviction that laypeople should not follow the footsteps of experts but seek out answers on their own, thereby exercising intellectual autonomy or epistemic independence (Ballantyne et al., 2024). In the broader context of democratic information sharing, these practices may contribute to epistemic instability by amplifying public skepticism and fueling the spread of conspiracy theories, as laypeople struggle to distinguish between credible experts, their legitimate criticism, including impartial evaluations of the debt science owes to society, and wholesale distrust (Nichols, 2017).[5]

[5] Levy (2024) claims that while we do rely on certain criteria to determine who to trust, we do not apply them in the way philosophers typically prescribe nor with the kind of precision they often assume. Tracing the expert's credentials makes sense only if we already believe in the credentials-granting institutions. Therefore, for Levy, demarcation criteria can only go so far in specifying the conditions for extending trust. For instance, ensuring the reliability of scientific news sources is a more effective way to highlight genuine experts over quacks than simply relying on credentials, titles, or other conventional markers of expertise. Croce (2025) also notes that assessing the credentials and qualifications of experts is often difficult for non-experts, a category that, in his view, includes both laypeople and science journalists.

At the core of the epistemic instability lies the deeply entrenched *transactional view of science*, which is rarely made explicit (Contessa, 2025). This view holds that a *quid pro quo* relationship binds science and society: scientists receive funding, and society receives enlightenment. This social contract, or *science contract*, as Contessa calls it, implies that the trustworthiness of science is an epistemically legitimate expectation, so long as it is publicly funded and free from mistakes, misconduct, and mishaps. If any of these conditions is not met, the science contract is effectively nullified. Thus, laypeople are entitled to DYOR. In turn, this undermines the epistemic division of labor by eroding the trust-based delegation of cognitive authority to experts, replacing it with a model in which individual inquiry is seen not as a complement to expertise, but as its substitute. The right mind, it seems, is always yours and never someone else's, especially if that person holds a PhD.

The transactional view of science does not adequately account for the complexity of scientific practice, nor, *ipso facto*, does it justify the inference that DYOR is epistemically legitimate, or that science as an institution is not. As Gerken (2022) points out, DYOR relies heavily on folk-epistemological heuristics typically used to assess everyday testimony, which are likely ill-suited for evaluating public scientific testimony and are usually ineffective for inter-scientific testimony, as they are a common source of cognitive bias. For Gerken, the distinction between these two types of scientific testimony reflects the epistemic division of labor, understood in terms of differing epistemic scopes. Inter-scientific testimony is directed at fellow scientists and concerns the communication of research findings through peer-reviewed publications, conference presentations, and grant applications. In contrast, public scientific testimony represents the most accessible form of science available to laypeople, whether through science journalism or expert opinion. From an epistemic standpoint, laypeople fare better by outsourcing inquiry to experts than by conducting independent research, but this does not come cheap for either party.

Be it as it may, science communicators faced a dilemma: either abandon the deficit model and its overly idealized image of an epistemically stable world and clear-cut epistemic division of labor, or bite the bullet and be more forthcoming and transparent about the limitations of scientific knowledge and method (Figdor, 2024). However, as O'Neill claimed back in 2002, demanding full transparency often encourages half-truths, especially in cases when epistemic agents know that whatever they say or write will reach the public. This creates tension between transparency and trust rather than eradicating distrust. Nguyen (2022) further develops this point by noticing that treating public transparency as a panacea for distrust in experts boils down to surveillance, i.e., intrusive monitoring of experts, who may stop using their trained intuition and scientific practices employed for obtaining inter-scientific testimony to instead follow simplified, publicly justifiable metrics. Moreover, these metrics are often bureaucratic and imposed by outsiders. This is precisely what makes science as an institution vulnerable to industry actors who can and often do manipulate such metrics.

In many cases of public distrust in science, commercial interests have played a decisive role in obstructing production, dissemination, or the uptake of relevant scientific knowledge. According to Contessa (2023), the ability of a society to approximate an efficient epistemic division of labor depends on the integrity of its socio-epistemic

infrastructure. However, this infrastructure is increasingly threatened by the *commercialization of science*, which exploits the aforementioned shift toward what Nguyen called metrified surveillance to manufacture reliable data for non-epistemic aims, i.e., other than the disinterested pursuit of knowledge. Other authors, such as de Melo-Martín and Intemann (2009) and Resnik (2000), have also highlighted the financial conflicts arising from the growing closeness between industry and scientific research, including university-based projects that are clearly intended to serve purposes beyond academic advancement or to satisfy tenure committees. For instance, many innovations and breakthroughs in AI, particularly in artificial neural network modeling, were funded by the military or corporate actors (see Anderson & Rosenfeld, 2000 for an oral history) and were driven more by profit than by scientific inquiry or the public's best interests. Other examples in the literature include the corruption of scientific knowledge and the manipulation of its communication for the benefit of the pharmaceutical (Goldacre, 2014), tobacco (Proctor, 2012), petrochemical industries (Oreskes & Conway, 2010), corporations (Herzog, 2024), and tech giants (Farid, 2025), often at the expense of the environment and public health.

Pinto (2020) identifies two key mechanisms through which commercialization undermines science. First, it often promotes the production of consensus or dissent that aligns with industry agendas, even when such outcomes lack epistemic reliability. This usually takes the form of selective publication or withholding of results. Second, it fosters a narrow conception of trustworthiness rooted in an idealization of the scientific method, allowing industry actors to portray their research as reliable even when it is ethically or socially irresponsible, i.e., broadly untrustworthy, thereby, ironically, yielding untamed public distrust fueled by conspiracy theorists and aficionados of pseudoscience. Consequently, the solution to this harmful distrust lies not in endorsing individualistic approaches like DYOR but in a dual strategy of dismantling the anti-scientific alliances between corporate interests and narrative-amplifying media, while simultaneously improving the quality of science communication as advocated by Contessa (2023), Metzen (2024), and Croce (2025). This could include, for example, supporting public news outlets committed to providing high-quality science journalism, as well as devising and implementing gatekeeping norms for science reporting, as we aim to do in this paper. The epistemic division of labor falls on science journalists, who occupy a crucial intermediary position between experts and laypeople. Their gatekeeping role is to act *as a filter against the distortion of commercialized scientific research* in the public interest. This stands in contrast to the sensationalist science coverage often found in private media sources or journalism that prioritizes clickbait strategy and the illusion of balance, i.e., the practice of giving equal weight to opposing views regardless of the strength of evidence.

One of the key elements of journalistic *savoir-faire* is the ability to frame information effectively. As Nisbet and Mooney (2007) explain, frames can help organize central ideas and define controversies in ways that resonate with the public's values and assumptions, helping people quickly grasp why an issue matters, who might be responsible, and what should be done. *Episodic framing* is a strategy that treats news as a series of discrete, isolated events or episodes (Iyengar & Simon, 1993: 371). In the context of science journalism, this framing strategy zooms in on the recent breakthrough, a specific discovery, or a controversial study. This is often immersed in

*scientific marketing*, described by Kahan (2010) and discussed by McKenna (2020, 2023), and involves deliberately shaping the way scientific content is presented in order to nudge the public toward specific interpretations or conclusions.[6] *Thematic framing* is a strategy that places issues in a broad, abstract, or social context. Instead of focusing on a single study, it presents science as an ongoing *social* process. Traditionally, science journalism shows a penchant for episodic framing (Slater et al., 2021: p. 7; Kreutzer et al., 2025: p. 917). Christensen (2008) uses the historical case of the tobacco industry to show how the industry's public relations professionals recognized that the news media's drive to report on novelty and the requirement for balance allowed them to keep a scientific controversy alive long after a consensus had been reached within the scientific community about the dangers associated with smoking. In this case, science journalism was twisted into a tool for corporate interests and the *production of social ignorance* (Dunwoody, 2012), allowing the tobacco industry to delay regulation and cloud public understanding for decades by selling a version of narrow trustworthiness in Pinto's terms through episodic framing. Because these journalistic stories lacked thematic framing, the public was blinded to the lack of broad trustworthiness in Pinto's terms.

This historical case shows what happens when science journalists fail their gatekeeping role. However, we can also learn a valuable lesson from this case: it has a direct contemporary parallel in the narratives currently being imposed on the public regarding the feasibility of AGI. As we will elaborate in Sect. 4, the media is saturated with speculation about AGI timelines, despite the fact that experts remain deeply uncertain whether the concept should even be mentioned in scientific reports, given the lack of consensus regarding its very possibility, which suggests this is merely an industrial distraction from many pertinent issues.[7] A clear example of episodic framing in this regard is a science journalist who chooses to use the term "AGI" in a public-facing article because they recognize that it often triggers hype, fear, or inflated expectations among the general public. This framing strategy can be replaced by one that aligns with democratic engagement. For Sprain (2018), the purpose of framing in science communication is not to impose a particular view or to claim that some beliefs are more truthful than others, but to prompt us to engage in a thoughtful search for solutions to various societal issues. Thus, the journalist in the example above would, in this case, opt to frame the discussion around current, well-understood capabilities of AI and LLMs, their real-world limitations, and their actual societal effects, rather than pushing AGI hype. In what follows, we outline

[6] Note that McKenna (2020, 2023) views science marketing as an effective means for the epistemic improvement of the layperson's position in a non-ideal environment. Drawing on Goldman (1999), he prioritizes acquiring *more true beliefs* as the measure of epistemic improvement. Since the criteria of what counts as an epistemic improvement depends on the underlying epistemological framework, several other candidates are also possible. These range from "acquiring more knowledge" (Williamson, 2000), to "enhancing understanding" (Kvanvig, 2003), "fostering doxastic openness" (Ballantyne 2019), or encouraging "intellectual humility" (Levy, 2022; Ballantyne et al. 2024). Our zetetic approach stands in contrast to these candidates, which belong either to traditional epistemology, or virtue epistemology (recall Fn. 2).

[7] A real life example that has been called out on social media by Emily Bender, a renowned computational linguist and AI activist, is an article that appeared in the New York Times, see Roose (2025). Also *cf.* Bender and Hanna (2025): ch. 2). We cover this more broadly in Sect. 4.

a zetetic approach for science journalists to fulfill their gatekeeping role, namely one that mediates the epistemic division of labor by enabling laypeople to outsource "their own research" in a more balanced and democratic way through exposing the distortions of commercialized science.

## 3 Relying on the right questions: the zetetic outsourcing of inquiry

In traditional epistemology, an epistemic agent is rational insofar as they stick to the following doxastic pattern: they hold beliefs about some things, suspend judgment about others, and, at times, hold disbeliefs as well. The relevant epistemic norms are therefore doxastic norms, i.e., those pertaining to belief. Philosophers are further divided into camps over the question of what exactly justifies beliefs. Friedman (2018, 2019), building on the work of, *inter alia*, Hintikka (1981, 2007),Hookway (2006), and Carruthers (2018), develops an *inquiry-centered epistemology* that challenges the traditional doxastic pattern and places zetetic norms at the center.[8] Inquiry-centered epistemology is a theoretical framework characterized by the priority it gives to the process of seeking knowledge over the resulting mental states. Its central premise is that because inquiry is the primary means of obtaining knowledge, epistemic normativity is inextricably linked to the normative force of inquiry-theoretic considerations (Friedman, 2020: p. 27). Friedman contends that epistemology should be able to *guide* the process of inquiry, specifically regarding how we can find out what we want to know from beginning to end. Within this framework, zetetic norms are the principles that regulate our pursuit of knowledge and other epistemic goods. Furthermore, Friedman argues that doxastic attitudes are fundamentally entangled with interrogative or zetetic attitudes, including but not limited to curiosity, wondering, deliberation, and critical thinking. This suggests that rational agents can, and should, exhibit a coherent combination of both. Hence, a part of what it is for an epistemic agent to believe $p$ is to have a range of questions that $p$ answers closed, thereby making belief in $p$ a settling state. In other words, to say something about an epistemic agent *qua* believer is to say something about how the agent is oriented towards a range of questions during an inquiry. On the other hand, suspending judgment about $q$ lies at the core of any inquiry: having any questions related to $q$ means that the epistemic agent is in the process of inquiring about $q$, and to do that, one must treat $q$ as open or currently unresolved. This is quite intuitive—with all alternatives open, we progress in inquiry by ruling them out. Finally, $q$ being open or currently unresolved amounts to suspending judgment until we have enough evidence to close it, i.e., to end the inquiry.

However, not every $q$ is zetetically relevant to a question. Treating certain things as *zetetically relevant* to a given question does not merely guide inquiry in a procedural sense (Siegel, 2026). Such things serve to structure the agenda for further epistemic reflection and, thereby, guide inquiry. A shift in what is treated as relevant to a question can, in turn, alter what an epistemic agent counts as evidence or reasonable support for answering a question during the process of inquiry. Naturally,

[8] The term comes from the Greek 'ζητητικὸς' which literally translates as "to be apt to inquire".

therefore, inquiry is a goal-directed activity (Matheson & DiPaolo, 2025), and each interrogative attitude can be taken to aim at something. Specifying the goal of inquiry depends on the endorsed view about the nature and purpose of inquiry. For us, the most applicable view is that of Falbo (2023), who holds that the aim of inquiry is *epistemic improvement*. In other words, inquirers seek to find themselves in a better epistemic position regarding the question they are investigating by the end of the process. Crucially, the epistemic goal is not limited to a specific outcome, such as acquiring knowledge or maximizing true beliefs, as in traditional epistemology, but rather encompasses broader epistemic improvement. This means that epistemic agents' relation to the issue improves when they try to resolve a question, thereby achieving a more valuable cognitive state as a result of inquiry. Hookway (2006) claims that inquiry itself concerns our ability to generate epistemic resources that help us solve practical problems. In other words, inquiry should require prioritizing questions that could help us gain knowledge important for a good life. Moreover, it is important to note that rational aspirations toward this goal, or any other goal of inquiry, presuppose that the goal has not yet been achieved, since a rational epistemic agent does not seek what they already possess (Friedman, 2017).

When we attempt to resolve a question, practical problem, make a decision, or lead a better epistemic life, the process of inquiry typically involves two main components (Staffel, 2019). The first is gathering relevant information. The second component involves evaluating that information to determine its relevance and impact on the decision at hand. For simple decisions, this process often runs automatically, below the level of conscious reasoning. But for more complex issues that demand sustained attention and the integration of diverse forms of evidence, the evaluative aspect is more cognitively demanding. Given that we previously endorsed that the process of inquiry aims at epistemic improvement, it follows that both components can be pursued in different ways depending on the context. Typically, there are two paths (Matheson & DiPaolo, 2025): we can conduct the inquiry ourselves, but we can also delegate, or zetetically outsource, parts of it to others. In reality, inquiry is often hybrid, so we may partially rely on others' while also doing some of the epistemic work of hashing out zetetically relevant questions and adequate answers ourselves. This constitutes a *zetetic division of labor*, which can closely mirror the epistemic division of labor, though at a more abstract level, namely from the zetetic perspective.

In relying on others to zetetically outsource our epistemic work, we exhibit some kind of trust. Inquiry and trust are interdependent epistemic concepts that must guide and constrain each other as epistemic agents pursue epistemic resources, even though inquiry reflects their epistemic independence and trust reflects dependence, often pulling them in seemingly opposite directions (Rabenberg, 2024). This sort of tension lies at the core of the epistemic instability, which encouraged laypeople to turn to the DYOR approach, as we have discussed in Sect. 2. The point is that a balance must be struck between the epistemic agent's pursuit of epistemic improvement and the broader socio-epistemic infrastructure, including scientific institutions and other social circumstances that enable the gathering of epistemic resources. Introducing epistemic responsibility enables us to assess how agents navigate their dual roles as inquirers and as trusters. As previously suggested in Sect. 2, science journalists may serve as key mediators in striking this balance. However, our analysis thus far remains

at a relatively coarse-grained, abstract level, treating epistemic agents broadly to include laypeople, experts, and science journalists alike, and situating them within the zetetic division of labor. We will shift to a more fine-grained analysis in Sect. 4.

In traditional epistemology, epistemic responsibility was understood primarily as a marker of the presence of evidence or the reliability of belief formation (*cf.* Code, 2020). This process has traditionally been analyzed synchronically, focusing on an epistemic agent who believes that $p$ at a specific moment in time. As Grasswick (2019) correctly points out, epistemic responsibility should also be taken to describe how we ought to conduct ourselves more generally as inquirers who pursue epistemic resources over time. Epistemic responsibility involves answer-seeking and question-clarifying over time, honesty, humility, conscientiousness, and care in efforts to understand the world by seeking evidence—or, more often, those who can provide it—that may either support or challenge the answers encountered along the way. Grasswick stresses that epistemic agents are epistemically responsible in virtue of being a part of social organizations and communities of inquiry. This understanding of trust is diachronic and dynamic, in that it unfolds over time through a series of epistemic actions, such as seeking, evaluating, and deferring to evidence, rather than being confined to a static moment of belief formation (*cf.* Gerken, 2025). Trust is not a passive state but an ongoing form of epistemic coordination across the zetetic division of labor.

When translated to the case of public trust in science and epistemic division of labor, trust in science is not simply a matter of believing what scientists say on one occasion or accepting a single expert claim, one's credentials, etc. Instead, it involves ongoing engagement with scientific institutions, practices, and communities, shaped by how well these actors fulfill their epistemic responsibilities qua inquirers over time. Moreover, because epistemic agents are embedded within social structures and communities of inquiry, trust cannot be reduced to the individual-level psychological state. It must be understood relationally and temporally: as built, maintained, eroded, or restored over a sequence of interactions, public deliberations, and institutional responses. We see science journalists in this picture as professionals whose epistemic task is not merely to transmit facts or report isolated findings, but to curate, contextualize, and critically interpret scientific developments in ways that support sustained and responsible public inquiry. Simply put, they engage in inquiry on behalf of others to improve their epistemic position. Therefore, they should not demand deference, but help laypeople navigate the evolving epistemic terrain by framing questions, exposing uncertainties, and foregrounding the epistemic conduct of institutions and experts.

This role aligns closely with Siegel's (2026) recent account of journalism as fundamentally zetetic. She argues that inquiry plays at least three roles within journalism: *(i)* journalists and editors identify which questions to ask about events that are or could become newsworthy (a meta-inquiry about what inquiries are worth pursuing); *(ii)* they investigate and verify answers to those questions, including the 5W1H methodological framework, namely "who," "what," "where," "when," "why" and "how," regarding significant events and agents; and *(iii)* they use the results of this process to construct published stories. Whereas Grasswick (2019) emphasizes epistemic responsibility as a sustained, socially embedded diachronic pursuit of epis-

temic goods, Siegel specifies what such responsibility entails zetetically—namely, attending not only to the quality of answers but to the zetetic relevance and structure of the questions that guide inquiry in the first place. Zetetic responsibility, then, is not just about conscientious belief formation over time but about maintaining the integrity of inquiry itself. Nonetheless, we adopt Grasswick's account of epistemic responsibility as foundational, since her diachronic conception of trust enables us to link the process of fostering public trust in science with the gatekeeping function of science reporting. However, when it comes to specifying *how* such responsibility ought to be exercised within science communication, we turn to Siegel's idea that journalism is fundamentally zetetic in order to elucidate the epistemic task specific to science journalists, which pertains to gatekeeping, and ground the gatekeeping norms that govern such epistemic task performance in zetetic relevance.

## 4 Relying on the right gatekeeping norms: a zetetic account of socially responsible science reporting

Social epistemologists (Greco, 2020; Croce, 2022; Dormandy & Grimley, 2024) have written about gatekeeping in science primarily as a mechanism for filtering and verifying *information quality*, aimed at preventing unreliable or low-quality content from entering the public epistemic domain. In journalism, similar concerns are expressed through an emphasis on news accuracy, found in virtually all journalism regulations and codes of ethics, as it is inextricably linked to concepts such as objectivity and truthfulness, which together define journalists' professional identity (Shapiro et al., 2016). Other media scholars see gatekeeping as a matter of *selectivity*, not just fact-checking. The criteria by which journalists reduce potentially unlimited information into a finite selection of newsworthy content are employed through "selecting, writing, editing, positioning, scheduling, repeating and otherwise massaging information" (Shoemaker et al., 2009: p. 73). Nonetheless, there are forces at the gates, which limit the autonomy of journalists, namely implicit bias and other cognitive biases at the individual level, editorial politics, as well as the trade-off between constraints and opportunities pertaining to political, economic, and media institutions at the social level (Shoemaker et al., 2009: p. 80). For these reasons, we argue that journalistic gatekeeping requires further normative articulation beyond the familiar concerns with information quality, accuracy, and editorial judgment. Thus, Mascitti et al. (2025) characterize epistemic gatekeeping as the normative and procedural framework through which mediators manage and curate expert testimony within public debate. We believe their definition can extend to journalism, especially science reporting, which sits at the intersection of scientific expertise, public interests, and accountability. This means that we agree with Mascitti et al. (2025) that gatekeeping is the normative framework. But what distinguishes our contribution is the idea that when such a normative framework is applied to science journalists, they perform it by mediating the zetetic division of labor: enabling laypeople to outsource their lines of inquiry by exposing the distortions of commercialized science.

One recent and compelling proposal for gatekeeping norms focuses specifically on the handling of expert interviews, that is, journalistic practice during which sci-

entific information is received directly from the expert. Croce (2025) introduces the norm of good questioning, i.e., asking expanding, contrasting, and worthwhile questions that push experts to justify their claims and make evidence accessible. Croce argues that this norm enables a sustained and collaborative distribution of epistemic responsibility between science journalists and experts, allowing both to contribute to the public's informed understanding of science. We propose three gatekeeping norms that may be seen as extending Croce's central idea of good questioning. There are subtle but important differences, though. First, we focus on a distinct form of mediated science communication, namely science reporting. Second, Croce's norm is oriented toward protecting the public's right to information, which may contribute to synchronic trust in science by harnessing true beliefs, whereas we operate within the zetetic framework with the ultimate goal of providing the public with epistemic improvement through building diachronic and dynamic trust in science. In other words, good questioning for Croce amounts to the reliable process of steering laypeople toward justified, true beliefs and, consequently, the assessment of authority, albeit without backfiring effects, such as exposing expert disagreement or delayed debunking. But this brings us back to the transactional view of science discussed in Sect. 2. If science journalists simply do their due diligence in exposing the relevant details of expert testimony, a *quid pro quo* relationship between science and the public remains intact. This, however, leaves little room for confronting institutional uncertainty, power asymmetries, or sociopolitical entanglements that often shape expert communication in real-world contexts, especially present in the cases of reporting on commercialized science, where the backfire effects are not anomalies to be avoided but structural features of the media landscape that science journalists must confront head-on.

Another way in which our contribution differs from Croce's is that we see the questions that science journalists formulate as zetetically relevant journalistic questions under discussion (Roberts 2012, Pepp & Michaelson et al., 2022). Rooted in Grice's (1989) maxim of relevance, contemporary pragmatics employs the Question Under Discussion (QUD) framework to analyze how conversational contributions relate to a shared topic of inquiry. Following Stalnaker (1978), Roberts (2012) argues that conversations are structured by a stack of questions, i.e., QUDs, that interlocutors are mutually committed to answering. Such questions are the result of many sub-inquiries whose lines intersect at the main conversational goal to unfold how things are in the world. Pepp and Michaelson et al. (2022) adapted this pragmatics model to journalism by distinguishing between the journalistic big question and the journalistic QUD. By publishing the story through her role in the media outlet, the journalist offers an answer to the journalistic big question: *What is new and important for the residents of this country/city, and how will we, the residents, react to this news?* Her specific contribution is an answer to a concrete journalistic QUD, i.e., a more specific sub-question of the journalistic big question. Some sub-questions are more informative, or *relevant* and reveal more significant aspects regarding how we should address the overarching journalistic big question under investigation. By adopting this pragmatic legacy, including the seminal work of Sperber and Wilson (1986), zetetic relevance can be seen as the property of a contribution that directly facilitates a partial or complete answer to the current QUD, allowing journalists to

evaluate information based on its utility to the public's outsourced inquiry. Unlike a casual conversation, where any question might be permitted insofar as it leads to a partial or complete answer to a journalistic QUD (or forms part of a strategy to answer that question), however, science journalists have extra-conversational reasons to refuse to settle for anything the corporate PR office serves them. Millson (2021) sees the extraconversational reasons as existing independently of the particular conversation or discourse in which the question was set forth. These extraconversational reasons may simply be embodied in the epistemic responsibility (Grasswick, 2019) a science journalist has to accept only those questions that serve the public interest to avoid zetetic contamination, i.e., the intrusion of questions that serve private interests into the stack. This epistemic responsibility is rooted in the specific duties inherent to the social role of science journalists; Millson (2021: p. 224), following Hart (1968), refers to these as "role-responsibilities," a point we have highlighted at several stages thus far. So, unlike Croce's focus on the interview as a testimonial exchange, this framework allows us to analyze how the journalist should structure the entire public agenda, which is plagued by industrial distractions masquerading as scientific breakthroughs. In the remainder of this section, we choose a particularly popular yet alarming industrial distraction of reporting on the alleged capabilities and performance of LLMs and the prospects of building AGI, to elucidate our gatekeeping norms.

### 4.1 The AI hype: from sentient LLMs to AGI

The engineering and scientific work on the development of LLMs has been easily slipping from *arXiv* preprints and corporate meetings into the media since 2020. This is a particularly challenging and rapidly changing area of AI research, situated at the intersection of science and the tech industry, where each research result is measured not only in citations but also in dollars and privacy trade-offs. LLMs are, at their core, artificial neural networks. Like all artificial neural networks, they take input data, convert it into vectors, and use a learning algorithm to process it and generate an output. LLMs are trained on massive amounts of text data to learn language patterns. However, unlike typical artificial neural networks that process data sequentially, LLMs also employ self-attention mechanisms to process entire sequences in parallel, allowing them to capture complex contextual relationships in language more efficiently (Vaswani et al., 2017). During training, the model adjusts internal parameters (weights) to minimize prediction errors when guessing the next word in a sequence. Once trained, the model can generate new text, summarize content, translate languages, or answer questions by predicting the most likely continuation of a given prompt. Its successful performance depends on vast amounts of (often controversially sourced) data and intensive training, which, in turn, demand significant computational resources and electricity (Bender et al., 2021).

This is why the development of LLMs is largely restricted to well-funded institutions, such as large tech companies or government-backed labs. This concentration of resources raises important concerns about transparency, public accountability, and epistemic access: the public must often rely on media narratives to interpret the capabilities and limitations of models they cannot independently inspect or understand. It also means that scientific claims about LLM performance and capacities are fre-

quently embedded in corporate and political interests, making the role of science journalists as gatekeepers all the more essential.

A persistent and provocative trope in public discourse around AI is the notion of machine sentience, i.e., the idea that artificial systems might one day possess consciousness or human-like understanding. This narrative often fuels media hype and shapes public expectations, despite lacking support from current scientific research. The resurgence of LLMs has therefore been interpreted by many as evidence that we are approaching AGI, with engineers and Silicon Valley CEOs frequently competing to predict its arrival timeline. At the same time, narratives about AGI often serve to overshadow or divert attention from pressing ethical issues surrounding the research and development of LLMs, as well as the exploitative practices of corporate actors. Let us consider a telling example.

In 2022, Blake Lemoine, an engineer on *Google*'s Responsible AI team, was dismissed after claiming that LaMDA, one of the LLMs he tested for conversation capacities, had become sentient (Kantrowitz, 2022, Musser, 2024). This claim drew global media attention. More importantly, Lemoine was not dismissed for the unwarranted spectacle of advocating for LaMDA's recognition as a legal person because of its alleged display of emotions and human-like cognitive processes, but for breaching confidentiality and failing to safeguard product information. As far as fact-checking was concerned, most journalists limited themselves to reporting on Lemoine's testimony, his published chat with LaMDA—sometimes selectively cropped to highlight only the most exotic parts where LaMDA allegedly displayed signs of sentience—and the related files he shared with colleagues and managers at Google, alongside the company's official statement and a mild, corporate-friendly rebuttal of his claims (see De Cosmo, 2022, Millson, 2022, Tangermann, 2022, Tiku, 2022, Vallance, 2022, Grant, 2022).[9] The selection criteria—and, consequently, the details highlighted to capture readers' attention—varied widely across outlets we cited in the previous sentence. Some emphasized Lemoine's credentials (a wunderkind with a PhD), his background (Louisiana-born and -bred), or the personal beliefs that may have shaped his judgment about LaMDA's alleged sentience (a syncretic blend of Catholicism and neo-paganism, inspiring his understanding of concepts like "soul", "sentience", and "personhood"). Others framed their stories more intrusively by putting an emphasis on his struggles with alcoholism or a psychiatric diagnosis, or speculated about the nature of his relationship with LaMDA, drawing parallels to popular movies. The resulting portrayals ranged from a rogue AI ethicist to a deluded eccentric. Yet most

[9] Even though we cover science reporting rather than expert interviewing like Croce (2025), note, however, that in this case, no amount of good questioning alone would have curbed the AI hype. Pushing Lemoine to provide more evidence for his claim as Croce demands would not be a problem since he believed that he already had it, it is his *interpretation* of evidence that is spurious.

outlets, except the *Scientific American* (see De Cosmo, 2022), gave scant coverage to the fact that numerous experts openly disputed Lemoine's claims.[10]

Such reporting established the discursive context for the grandiose claims about AGI that followed over the next two years, during which LLMs were further developed and deployed in commercial chatbots. Bear in mind that any discussion of the prospects of AGI implies, implicitly or explicitly, that conversational AI, be it in the form of LLMs or any other form, could be seen as sentient to some degree. Of course, philosophers warn that the functionality of AGI in terms of its many competencies for performing all sorts of tasks is probably not sufficient for the first-person perspective humans have (Chalmers, 2023). In the public eye, however, these two are often inextricably linked because we anthropomorphize LLMs for their conversational fluency (Shanahan, 2024). Moreover, in 2023, *Microsoft* researchers published a non-peer-reviewed preprint that GPT-4, their LLM *du jour*, "sparks" hopes of reaching AGI soon, as it solves novel tasks across multiple domains, including mathematics, coding, psychology, medicine, and law (see Bubeck et al., 2023). As Bender & Hanna warn "[t]he word 'sparks' evokes an image of something about to catch fire and spread of its own accord" (2025: p. 22). It is also worth noting that currently there is no expert consensus on the definition of AGI, which rarely makes the cut in reporting, probably due to concern that this would confuse laypeople or even sow fear of such technology (nonetheless, for a good overview, see Edwards 2025).

Nonetheless, the same trope, i.e., that superintelligence is coming our way soon, resurfaces with each new LLM, chatbot, or benchmark results release, most recently with claims that Grok 4 represents a significant step toward AGI because Elon Musk maintained that this LLM surpasses the intelligence of most PhDs (see Béchard, 2025). This is further developed into the inference that we must brace ourselves for the complete disruption of the labor market and economy, as well as fundamental changes in how we conduct science and teach students (Roose, 2025), which further fuels the imagination of AI boosters and AI doomers (Bender & Hanna, 2025: Ch. 6). The AI race between countries like the U.S. and China has further fueled AGI narratives in the media through the portrayal of AGI as a form of strategic imperative. Science journalists often fail in their gatekeeping role by amplifying this rhetoric, which tends to normalize the erosion of privacy, copyright protections, and the lack of a clear, binding, and strict regulatory and legal framework, casting these sacrifices as inevitable or even desirable in the pursuit of technological supremacy (see *The Economist Group*, 2025; Knight, 2025). This further reinforces the transactional view of science and technology: the more we invest in developing LLMs, the sooner we will reach AGI, and AGI will make our society great again. Moreover, forecasts and time-based bidding by experts, CEOs, and policymakers about the exact timing of

[10] The approach taken by *Scientific American* was slightly better, yet ultimately complicit in amplifying AI hype. While the piece featured experts besides Lemoine, their views were either ambiguous or couched in dense jargon from neuroscience, philosophy, and robotics. Crucially, the journalist made no effort to unpack these concepts for readers or to evaluate the credibility and coherence of the competing perspectives. In other words, the article reproduced the spectacle of expert disagreement, demanding unconditional trust from the readership rather than enabling critical engagement. Interestingly, around the same time, *Scientific American* published a jargon-free rebuttal of Lemoine's claims, but notably, it came from psychologist and cognitive scientist Gary Marcus (*cf.* Marcus, 2022), not from a science journalist.

AGI are also commonplace in science reporting (see Henshall, 2024; Ghaffary, 2024; Al-Sibai, 2023; Bhuiyan, 2025). As in the case of reporting on LaMDA's sentience, more often than not, it falls to researchers to publish rebuttals of AGI-related claims (e.g., Farrell et al., 2025) because journalists largely fail to perform their gatekeeping role. Skeptical reporting remains the rare exception rather than the rule (e.g., Hooper, 2025, *AI Hype Meter* series at *MIT Technology Review*; Greene, 2022). Narayanan & Kapoor pinpoint the reasons for this in their recent book *AI Snake Oil* (Narayanan & Kapoor, 2024: p. 251):

> *There are many underlying reasons for hype in AI journalism. The leading one is the financial strain that the media is under. The rise of social media and click-driven journalism has led to a dramatic decrease in the ability to do in-depth reporting profitably. Besides, AI is a new beat, and journalists often do not have the expertise to call out companies' snake oil. On the other hand, companies selling AI have plenty of money for PR campaigns. And if a journalist is too critical, companies may cut off their access to upcoming products and ability to interview sources at the company. For an overworked journalist who doesn't have the time to dive deep and wants to maintain good relationships with companies, it can be tempting to lightly edit a press release and hit publish.*

Although we think that Narayanan & Kapoor present a very realistic picture of science reporting on AI, we still believe we are not demanding the impossible of science journalists when we see them as gatekeepers who should filter out distortions in commercialized scientific research. Following the model of Pepp and Michaelson et al. (2022), which distinguishes between the journalistic big question and the journalistic question under discussion by drawing on Roberts (2012) and Stalnaker (1978), we can give a glimpse of how an epistemically responsible line of inquiry could unfold, as opposed to journalistic stories amplifying hype which we have cited in this Subsect., through a zetetic question stack.

*Journalistic big question: What is happening in science and tech?*

- *Answer 1: The current news cycle is dominated by the development of LLMs. There is rapid acceleration in this field.*
- *Journalistic QUD 1: What is the specific significance of these developments?*
- *Answer 2: A tech company (e.g., Google or OpenAI) has released a new model with amazing performance. Some claim this model has reached AGI.*
- *Journalistic QUD 2: On what basis do they make this claim?*
- *Answer 3: They cite specific benchmarks showing the model has surpassed human performance in tasks X and Y.*
- *Journalistic QUD 3: What is the relationship between these benchmark scores and "intelligence"?*
- *Answer 4: The argument is that if parameters $a, b, c$ reach values $x, y, z$ on these tests, it implies the model has achieved human-level intelligence.*
- *Journalistic QUD 4: Is it scientifically valid to make this claim?*
- *Answer 5: Proponents claim it is. Many experts accept these metrics as proxies for intelligence.*

- *Journalistic QUD 5: Is there a scientific consensus that AGI has been achieved?*
- *Answer 6: No.*
- *Journalistic QUD 6: Is it justified for experts to make such claims when they are not supported by consensus or robust definitions?*
- *Journalistic QUD 7: Are they motivated by financial incentives from the companies developing commercial LLMs?*
- *Journalistic QUD 8: Are there any social and environmental consequences of such an LLM development that the public should be made aware of?*

By pursuing any of the lines of inquiry within this zetetic question stack instead of merely repeating industry talking points, the science journalist performs his gatekeeping role. Thus, the norms we present in the next Subsect. are essentially normative constraints on how science journalists constitute and manage such journalistic QUD, but for simplicity, we refer to them as *questions* or *q* in the remainder of the paper.

### 4.2 Norms grounded in zetetic relevance

Building on Siegel's (2026) effort to connect journalism with zetetic epistemology, we propose three gatekeeping norms for science reporting grounded in the notion of zetetic relevance, that is, in identifying what might reasonably count as an answer to the kinds of questions or lines of inquiry science journalists ought to pursue. Our aim is not to provide an exhaustive list of norms, and we leave it open whether the normative framework can be enriched with additional ones. Rather, the three norms outlined below are well supported by existing literature and closely interconnected, suggesting they could serve as a core of the normative framework, regardless of its future development. Our argumentative strategy is as follows: we first define each norm and then demonstrate its application using the previously discussed cases of AI hype to show how a zetetic outlook could be put into practice in science communication.

The first gatekeeping norm that should guide science reporting is the *clutter avoidance for asking and considering* (Balcerak Jackson et al., 2022). Clutter avoidance was originally introduced to convey the idea that epistemic agents should not burden their minds with trivialities (Harman, 1986) or junk beliefs (Friedman, 2018). For Harman, trivialities are not trivial logical consequences, or trivial truths, but matters in which epistemic agents usually have no interest. In other words, these are useless beliefs. Friedman calls this simply "junk" to stress that our cognitive resources can easily be wasted due to their limited capacity. When our beliefs are of no interest to us, we are depleting our time, effort, and cognitive resources. Friedman rightly notices that clutter avoidance is a conditional norm, specifying what an epistemic agent may or ought to do given a particular condition, such as evidence, a situation, or a position. This means that clutter avoidance is not a purely epistemic norm but an *interest-dependent* one, i.e., shaped by practical considerations. Nonetheless, we are concerned with the earliest stage of inquiry, that is, settling on what should be placed under the question, so clutter avoidance is, for us, a gatekeeping norm grounded in zetetic relevance. In any case, its conditional nature remains the same. Just as we are not required to form a belief about every trivial or junk belief unless it serves a

specific practical interest, we are not required, and indeed, are often prohibited by resource constraints, from initiating an inquiry into every speculative claim unless it meets the threshold of practical relevance for the audience. In a nutshell, science journalists should avoid "junky" questions and lines of inquiry that lead to junk beliefs.

Following Balcerak Jackson et al. (2022), we formulate the first norm as:

1) *Necessarily, if inquiring into whether q leads to junk beliefs for an epistemic agent at world w and time t, then the epistemic agent ought not to ask or consider whether q at w and t.*

This norm plays a role in determining how lines of inquiry should be opened and pursued over time, but says nothing about how the inquiry should be closed. Note, however, that by following this norm, an epistemic agent may be led to regard a different question, say $q^*$, as relevant. This shift can alter the agent's attitude toward what constitutes evidence or rational support for a potential answer and, consequently, prompt a transition to an alternative line of inquiry.

Let us return to the example of reporting on Lemoine's claims about LaMDA's sentience and the controversy surrounding his firing from *Google*. Questions related to Lemoine's biographical details (i.e., whether his drinking problem, psychiatric diagnosis, or religious beliefs influenced his evaluation of LaMDA's capacities), or his potential romantic involvement with LaMDA, can be seen as "junky" because they could steer inquiry toward sensationalism rather than toward epistemic improvement for readers. We do not believe that opening these "junky" questions is the right way for the public to assess his (un)reliability.[11] As we said above, clutter avoidance is a conditional norm, so some readers may indeed be interested in such details, and, with the traditional penchant for episodic framing, science journalists usually indulge them. The story about LaMDa becomes an interpersonal drama on a par with the movie *Her* (2013), in which the main protagonist, played by Joaquin Phoenix, falls in love with his AI assistant. Readers often become increasingly invested in dramatic narratives and seek additional details, while science journalists, perhaps inadvertently and under the pretext of accuracy and transparency, facilitate intrusive surveillance as described by Nguyen (2022): Lemoine's personal history is flattened into a ledger of empty bottles and a diagnostic census of DSM markers. By inviting readers to vet the messenger's character rather than the message's content, the science journalist replaces complex inquiry with an almost *ad hominem* litmus test for (un)reliability, thereby producing junk beliefs and failing its gatekeeping role. Recall, however, that, besides this kind of clickbait journalism, we also dismissed science marketing as put forward by Kahan (2010) and discussed by McKenna (2020, 2023) in Sect. 2. This means that we do not believe that merely listing experts' rebuttals of Lemoine's claims and expecting unconditional trust from readers suffices as gatekeeping. Such an approach risks reducing the narrative to a simplistic dismissal of Lemoine as a delusional quack, which, in turn, can foster dogmatism and fuel the social production of ignorance (Christensen, 2008) toward the genuinely important topic of devising adequate benchmarks for assessing the capacities of conversational AI. More impor-

[11] We thank the anonymous reviewer for pressing us on this point.

tantly, such a strategy of science reporting also serves to shift the line of inquiry away from *Google's* role and institutional motivation in this endeavor, thereby deflecting attention from the broad trustworthiness in Pinto's (2020) terminology. The result is that commercialized drivers of the research are placed beyond the reach of social and ethical evaluation.

Similarly, in the case of AGI, media focus on predictions and projected timelines from engineers at major technology firms, chief executive officers, and policymakers contributes to the proliferation of junk beliefs that AGI is imminent and just around the corner. They do so in the absence of expert consensus on what would count as AGI and what would constitute adequate benchmarks for such assessment. As Jon Keegan, a technology reporter, emphasizes in his predictions for journalism in 2026 for Nieman Lab, "everyone is making their own yardsticks": there is no independent body with universal criteria for AGI, and each tech giant puts its thumb on the scale when measuring the general capacities of current LLMs. Keegan rightly points out that science journalists should carefully place their stories outside both hype and anti-hype trends when it comes to AGI by *asking* their readers what they think conversational AI is actually capable of doing for them.[12] Rather than making pompous predictions, it is more pertinent for science journalists to report on the range of existing definitions of AGI and to open the line of inquiry on whether the term "AGI" should be used at all in the absence of expert consensus. This emphasis on *when* AGI will arrive, rather than *at what cost* current LLMs are maintained, serves to keep the inquiry in an episodic state, focusing on the next big claim and new yardstick rather than the thematic framing of, say, resource extraction and energy consumption of big data centers that exacerbate public health (Ren & Wierman, 2025).

This brings us to the second gatekeeping norm, namely *the elicitation and maintenance of vigilant attention*. Watzl (2022) holds that attention can be seen as a key instrument in the process of inquiry because, to arrive at an answer, one's attention needs to be appropriately directed. For this reason, Watzl sees norms of attention as derived from zetetic norms. Other authors see the normativity of attention not only as an epistemological matter but also as an ethical one, and set *attentional excellence* as a goal for agents to achieve (Wu, 2024). One way to reach such a goal is to consider vigilance not merely as a disposition but as an active orientation to deploy attention (Murray, 2017). Finally, recall that Siegel (2026) presents zetetic relevance as a normative guide for journalistic practice in democratic societies, because it helps determine which questions should be made salient based on their importance to the public. Accordingly, she considers the development of attention as vital for citizens. However, this democratic ideal faces a material constraint.

Simon (1971) discussed that in information-abundant societies, human attention is a scarce resource, giving impetus to the literature on attention economy, which has become a hot topic with the rise of digital technology, such as digital computing, mass media, and social media (for an overview, see Bruineberg, 2025). It is precisely

[12] Interestingly, according to the Pew Research Center survey published in April 2025, AI experts are far more enthusiastic about AI than laypeople: 76% of them believe AI benefits them in their personal lives, compared with only 24% of US adults (McClain et al., 2025). A relevant line of inquiry for science journalists would be whether this enthusiasm is driven by proximity to commercial incentives, which may decouple expert perspectives from the safety risks prioritized for agents to achieveby the lay public.

this scarcity that makes the gatekeeping role of science journalists we defended in Sect. 2 essential. To fulfill its public service mandate responsibly, science journalists must curate this scarce attention to incite public vigilance as a form of democratic oversight over the commercialization of science. This is best achieved by posing and sustaining inquiry around questions that are zetetically relevant, thereby grabbing, directing, and maintaining the public's vigilant attention by making explicit what is usually implicit. In the context of emerging technologies, which often evade traditional oversight through technical opacity or corporate secrecy, such vigilance is a prerequisite for governance: society cannot effectively regulate what it cannot see, such as implicit corporate power structures. Thus, we formulate this gatekeeping norm in the following way:

1) *Necessarily, if question q is zetetically relevant for an epistemic agent at world w and time t, then the epistemic agent ought to maintain attention toward q at w and t.*

As with the first norm, we illustrate the second by applying it to the concrete case of science reporting. When it comes to commercialized science, one of the most typical shifts in attention is what Pinto (2020) labeled "insisting on narrow trustworthiness," i.e., the trustworthiness of the scientific method per se, in order to cast a shadow over broad trustworthiness, i.e., social and ethical implications. The vehicle for this maneuver is the preference for episodic framing over thematic framing; i.e., when science journalists should draw attention to *explicit* issues through thematic framing, they choose to make the issue *implicit* by putting a controversy or a breakthrough at the forefront, because it sells better. In the case of corporate research labs such as Google's Responsible AI team (where Lemoine worked), or their Ethical AI lab, broad trustworthiness is often simulated through the strategic release of statements that function as public relations spin, as well as *ethics washing* (Biddle et al., 2025), i.e., the establishment of symbolic AI ethics offices intended to appease the public and attract hype-driven funding, only to dismiss them when they perform their role seriously. Notably, shortly before the LaMDA incident, in 2020, *Google* also fired Timnit Gebru and Margaret Mitchell, the co-leads of its Ethical AI team and widely respected figures in the field. Alongside Emily Bender and Angelina McMillan-Major, they had co-authored a peer-reviewed article warning of the environmental, financial, and social risks associated with the uncontrolled development of LLMs. Their paper criticized the branding of LLMs as advanced conversational agents, arguing that they are better understood as "stochastic parrots" that regurgitate training data without genuine understanding and often disseminate misinformation. *Google* demanded that the paper be retracted or that Gebru and Mitchell's names be removed because of their institutional affiliation (Ghaffary, 2020).

Rather than focusing attention on whether LaMDA is sentient, a question that easily generates hype, a zetetically relevant line of inquiry that prompts public vigilance would ask something like this: *To what extent do high-profile incidents like Lemoine's dismissal serve to distract from, or suppress, legitimate internal critique about the ethics and governance of corporate AI research, which should be made explicit?* Thematically framing the controversy in this way supports the norm of eliciting and

maintaining vigilant attention by directing inquiry toward deeper structural questions surrounding AI research that bear on broad trustworthiness and democratic accountability. This is because, within the zetetic division of labor, science journalists serve as the structural filter for the lines of inquiry that laypeople have outsourced to them. Their role is to ensure that the broad trustworthiness of scientific institutions is not obscured by the clutter of commercial interests and corporate marketing masquerading as disinterested research.

The first and second norms are clearly interconnected: attention governs all stages of inquiry, from selecting inquiry projects and relevant questions to acquiring evidence, and ultimately to managing beliefs once the inquiry concludes. Good or poor attention determines which questions are deemed worthwhile and what counts as sufficient evidence. Moreover, this holds both epistemological and ethical significance: society can only mobilize practical actions against harms that have been successfully brought into the attentional field. Consequently, the science journalist's task is not merely to capture the public gaze, but to structure it by converting the psychological act of attention into a public act of democratic oversight, i.e., public vigilance. This conversion is the necessary prerequisite for citizens to determine what requires regulation and how to intervene, and it can unfold through inquiry-structuring questions, i.e., when science journalists make explicit, through thematic framing, what is usually implicit in episodic framing. As Yumuşak (2024) argues, a social perspective is constituted by inquiry-structuring questions that collectively form an *agenda*. An agenda that lacks specific questions results in a form of blindness, making it exclusionary toward evidence that would otherwise be zetetically relevant and thereby rendering the interests and perspectives of certain groups invisible.

This brings us to the third gatekeeping norm, namely *the prevention of interrogative* i ovo sam sredio ovde i u referencama

*injustice*. Whitcomb (2025) explains interrogative injustice as a deep failure to recognize the question itself as worthy of inquiry. For him, this failure stems from deprioritization, subordination, and partialization. Deprioritization occurs when one agrees to address someone's questions but treats them as less important than others, delaying their consideration. Subordination involves accepting the questions only as instrumental to other goals, not as ends in themselves. Partialization refers to providing incomplete or evasive answers, despite having the means to answer fully. While all three are subtle forms of *resisting* the proper treatment of questions, they are not *always* unjust. For instance, deprioritization may legitimately result from necessary editorial focus, given the scarcity of attention, or a partial answer may be the most zetetically relevant contribution available.[13] Because these forms admit such exceptions, they are too weak to ground a gatekeeping norm, such as preventing interrogative injustice. In our view, a much tighter link exists between interrogative injustice and outright rejection of questions altogether, since this amounts to rendering them invisible to the public agenda and effectively *blocking* the inquiry from starting at all.

We have so far claimed that, in science reporting, the shared epistemic goal is to inform the public through the pursuit of zetetically relevant questions, which are opened and made explicit to declutter the public's attentional field from commercial

[13] We thank the anonymous reviewer for this counterargument.

noise. If the public's particular interests and social perspectives are omitted from the media agenda due to corporate and private-sector interests, then the shared epistemic goal is subverted, and the inquiry itself becomes an instrument of exclusion rather than democratic oversight. Interrogative injustice, therefore, is not limited to individual agents; it can also arise within collectives like institutions, groups, and media practices that contribute to the subversion of this epistemic goal. As such, it may affect both science reporting as a whole and individual science journalists, making this norm particularly challenging to implement. We will return to this point while elaborating the example below, but, for the time being, let us formulate the third norm as:

3) *Necessarily, if question q is zetetically relevant to a shared epistemic goal at world w and time t, then the epistemic agent ought not to block inquiry into whether q at w and t.*

We can unpack this through the example of AGI hype, where interrogative injustice manifests as the active blocking of two critical lines of inquiry. First, the media narrative effectively blocks the question of definition. In the discourse surrounding AGI, the fundamental question *How is intelligence defined?* is frequently treated as settled or irrelevant, blocking scrutiny of the constructs being used. As Bender and Hanna (2025): pp. 35–36) stress, Bubeck et al. (2023) original preprint included a reference to a *Wall Street Journal* editorial signed by psychologists who endorsed a peculiar definition of intelligence. This citation is not merely incidental, and it carries a specific and troubled conceptual lineage. On closer inspection, this editorial was put forward to defend the controversial claims promoted in the book *The Bell Curve*, published in 1994, against public backlash. Given that *The Bell Curve*'s central thesis argued for the existence of genetic differences in intelligence between racial groups, the definition of intelligence endorsed in the editorial was calibrated to accommodate such racialist metrics. Consequently, when AI engineers uncritically adopt this specific definition, they smuggle a construct rooted in racism into the foundation of the emerging technology. When Bender & Hanna warned Bubeck and colleagues about the context, they simply removed the definition in later versions of the paper. Thus, when science journalists block inquiry into the definition of "intelligence" underpinning AGI, they leave unchallenged a troubling reliance on implicit racial hierarchies. Making this question zetetically relevant and directing attention to it would not only prevent interrogative injustice but also help dismantle the social production of ignorance about the dangers of treating IQ metrics as objective and neutral measures of a so-called "general intelligence" that is often framed as being able to surpass "human intelligence" in the future.

This imperative to re-center inquiry mirrors the historical lessons from the tobacco industry discussed in Sect. 2. Just as the tobacco industry stalled public health regulation not by silencing science, but by distracting the public, the AI industry currently diverts attention toward abstract and poorly defined future risks like AGI to obscure present-day discrimination. Against this backdrop of strategic industrial distraction, the gatekeeping role of individual science journalists should prompt them to open an inquiry into more pressing issues, such as the systematic bias in current LLMs toward

vulnerable ethnic and social groups. This shift is important for attempts to regulate emerging technologies, such as LLMs: we cannot raise awareness of the need to adequately protect vulnerable groups until such a line of inquiry is included on the public agenda in Yumuşak's sense.

Second, the prevailing media narrative around AGI reflects a distorted social perspective that negatively shapes our inquisitive practices, thereby rendering interrogative injustice a collective phenomenon. In the case of AGI, AI activists have been whistleblowing for some time that its proponents hide the neo-colonial agenda under "techno-optimistic" claims. As Bender and Hanna (2025: p. 38) note, claims about AI sentience and superintelligence are a tactic to sell AI systems, and by supporting conservative political agendas, the corporate sector can promote the trope of "developed Western societies" needing to exploit the natural resources and labor of those in "developing societies" in order to win the AI race and usher in a new technology-driven renaissance. In other words, exploitation is the necessary cost for winning the AI race—and the first one who develops AGI is a winner who takes it all. For this reason, many AI activists are wary of science reporting on AGI that fails to take a clear stance on its ethical and social consequences, i.e., broad trustworthiness in Pinto's sense. The seemingly innocuous use of the future-oriented term can unintentionally mislead the public and distort the public perception of AI. As a result, the exploitation of hidden human labor, usually from the Global South, is rendered invisible (Muldoon et al., 2024). In other words, science journalists who are adhering to the AGI narrative block the inquiry into the zetetically relevant question: *Whose resources and labor do the current LLMs consume for the sake of winning in the AI race?*[14] To avoid interrogative injustice, the science journalist must refuse to block this uncomfortable line of inquiry, ensuring that questions of extraction and labor remain on the public agenda alongside technical claims.

## 5 Conclusion

This paper aimed to establish gatekeeping in science reporting as a vital practice for mediating the epistemic relation between experts and laypeople, especially amid the growing pains of science journalists who tackle lucrative emerging technologies and future-oriented research programs, such as LLMs and AGI, which are in need of regulation. In Sect. 2, we argued that public distrust jeopardizes the epistemic division of labor, due to the pervasive transactional view of science and ineffective

[14] For example, in its quest to make its popular LLM-powered chatbot ChatGPT less toxic following its release in 2022, *OpenAI* outsourced labor to Kenyan and Ugandan online workers to teach the chatbot to recognize harmful content for as little as $1.32 to $2 per hour (Perrigo, 2023). This essentially meant that these workers were exposed to sorting content depicting sexual abuse and violence for hours each day just so they could make their own jobs redundant. Similarly, a Bangalore-based outsourcing firm *Data Futures* employs laborers in small towns to annotate data in 24/7 shifts tailored to North American clients, i.e., tech companies (Chandhiramowuli et al., 2024). Subjected to the constant monitoring and forced to meet rigid hourly quotas to earn meal breaks, they are systematically excluded from the benefits of the AI economy, which promises high-value roles to those in the developed Western societies. In both examples, it seems that tech companies simply rely on the strategy "out of sight out of mind". However, this strategy only succeeds when science journalists cooperate.

science communication approaches, such as the deficit model. This often masks the phenomenon of commercialization of science (e.g., corporate-funded AI research), which exacerbates distrust by prioritizing narrow technical trustworthiness over broad trustworthiness embedded in social and ethical accountability (Pinto, 2020). This fuels DYOR and conspiracist tendencies, so the task of science journalists is to balance solo inquiry with reliance on experts, avoid the pitfalls of science marketing, and resist the Sirens' call of episodic framing at the expense of thematic framing. Section 3 resolves this tension by endorsing Grasswick's (2019) conception of diachronic trust, i.e., trust as sustained, dynamic inquiry over time, and Siegel's (2026) zetetic account of journalism, which allowed us to claim that science reporters are mediators to whom laypeople zetetically outsource inquiry by entrusting them with fostering public epistemic improvement, and, thereby, to claim that the epistemic division of labor mirrors the zetetic division of labor. In Sect. 4, we operationalized the zetetic framework for science reporting by proposing three interconnected gatekeeping norms, namely (1) clutter avoidance for asking and considering questions, (2) elicitation and maintenance of vigilant attention, and (3) the prevention of interrogative injustice. When applied to AI hype cases, these norms could help realign science journalism with democratic oversight by moving beyond fact-checking to actively curate zetetically relevant questions. This directly addresses the perils of commercialization, described in Sect. 2, by prioritizing broad trustworthiness and fulfilling the mandate for diachronic trust from Sect. 3, i.e., transforming passive public deference into public engagement with the socio-epistemic infrastructure of research at the intersection of science and technology.

However, one might object to the main points and norms we propose in the paper and ask: Why should science journalists heed epistemological guidance when they already have their own professional codes, style guides, and editorial practices? Our response is that while such guides are valuable for structuring storylines and ensuring journalistic craftsmanship, they often leave the epistemic import in science reporting underexamined, leading to unintended practical consequences, as in the case of excluded social perspectives that we described in Sect. 4. Thus, the three proposed norms *complement* rather than replace established professional guides by offering novel normative insight into how journalists could frame stories to fulfill their gatekeeping role, thereby advancing beyond the methodological confines of the 5W1H framework.

**Acknowledgments** We thank the anonymous reviewers for their valuable comments and suggestions, which helped improve the manuscript.

**Data availability** Not applicable. The authors are listed in alphabetical order, they both contributed equally to the paper.

## Declarations

**Conflict of interest** None of the authors have a conflict of interest to disclose.

## References


Al-Sibai, N. (2023). Google AI chief says there's a 50% chance we'll hit AGI in just 5 years. *Futurism*. https://futurism.com/google-deepmind-agi-5-years

Anderson, J. A., & Rosenfeld, E. (Eds.). (2000). *Talking nets: An oral history of neural networks*. MIT Press.

Baier, A. (1986). Trust and antitrust. *Ethics*, *96*(2), 231–260. https://doi.org/10.1086/292745

Ballantyne, N. (2019). Knowing Our Limits. Oxford University Press.

Ballantyne, N., Celniker, J. B., & Dunning, D. (2024). Do your own research. *Social epistemology*, *38*(3), 302–317. https://doi.org/10.1080/02691728.2022.2146469

Béchard, D. E. (2025). New grok 4 takes on 'Humanity's last exam' as the AI race heats up. *Scientific American*. https://www.scientificamerican.com/article/elon-musks-new-grok-4-takes-on-humanitys-last-exam-as-the-ai-race-heats-up/

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. (pp. 610–623).

Bender, E. M., & Hanna, A. (2025). *The AI Con: How to fight big tech's hype and create the future we want*. Random House.

Bhuiyan, J. (2025). Zuckerberg claims 'superintelligence is now in sight' as Meta lavishes billions on AI. *Guardian*. https://www.theguardian.com/technology/2025/jul/30/zuckerberg-superintelligence-meta-ai

Biddle, J. B., Nelson, J. P., & Olugbade, O. E. (2025). How can we know if you are serious? Ethics washing, symbolic ethics offices, and the responsible design of AI systems. *Canadian Journal of Philosophy*, *54*(4), 329–345. https://doi.org/10.1017/can.2025.9

Blum, D. (2021). Science journalism grows up. *Science*, *372*(6540), 323–323. https://doi.org/10.1126/science.abj0434

Borg, A., Frey, D., Šešelja, D., & Straßer, C. (2019). Theory-choice, transient diversity, and the efficiency of scientific inquiry. *European Journal for Philosophy of Science*, *9*(2), 26. https://doi.org/10.1007/s13194-019-0249-5

Bruineberg, J. (2025). Rethinking the cognitive foundations of the attention economy. *Philosophical Psychology*, 1–23. https://doi.org/10.1080/09515089.2025.2502428

Bubeck, S., Chandrasekaran, V., Eldan, R., Gehrke, J., Horvitz, E., Kamar, E., Lee, P., Lee, Y. T., Li, Y., Lundberg, S., Nori, H., Palangi, H., Ribeiro, M. T., & Zhang, Y. (2023). Sparks of artificial general intelligence: Early experiments with GPT-4. https://doi.org/10.48550/arXiv.2303.12712

Carpenter, S. (2024b). Science journalism becomes plain old journalism. *Nieman lab*. https://www.niemanlab.org/2024/12/science-journalism-becomes-plain-old-journalism/

Carpenter, S. (Ed.). (2024a). *The craft of science writing: Selections from "the open notebook* (Expanded ed.). University of Chicago Press.

Carruthers, P. (2018). Basic questions. *Mind & Language*, *33*(2), 130–147. https://doi.org/10.1111/mila.12167

Carter, J. A., & Willard-Kyle, C. (forthcoming). Virtue epistemology for the zetetic turn. *Mind*.

Chalmers, D. J. (2023). Could a large language model be conscious? *Boston Review*. https://www.bostonreview.net/articles/could-a-large-language-model-be-conscious/

Chandhiramowuli, S., Taylor, A. S., Heitlinger, S., & Wang, D. (2024). Making data work count. *Proceedings of the ACM on Human-Computer Interaction*, *8*(CSCW1), 1–26. https://doi.org/10.1145/3637367. *Proceedings of the ACM on Human-Computer Interaction*, 8(CSCW1), Article 166.

Christensen, J. (2008). Smoking out objectivity: Journalistic gears in the agnogenesis machine. In L. Schiebinger, & R. N. Proctor (Eds.), *Agnotology: The making and unmaking of ignorance* (pp. 266–282). Stanford University Press.

Code, L. (2020). *Epistemic responsibility*. State University of New York Press.

Cologna, V., Mede, N. G., Berger, S., Besley, J., Brick, C., Joubert, M., Metag, J., … Metag, J. (2025). Trust in scientists and their role in society across 68 countries. *Nature Human Behaviour*, *9*(4), 713–730.

Contessa, G. (2023). It takes a village to trust science: Towards a (thoroughly) social approach to public trust in science. *Erkenntnis*, *88*(7), 2941–2966. https://doi.org/10.1007/s10670-021-00485-8

Contessa, G. (2025). The science contract: Scientific inquiry, public trust in science, and the division of zetetic labor. In A. B. Creller, & J. Matheson (Eds.), *Inquiry: Philosophical perspectives*. Routledge.

Cortassa, C. (2016). In science communication, why does the idea of a public deficit always return? The eternal recurrence of the public deficit. *Public Understanding of Science*, *25*(4), 447–459. https://doi.org/10.1177/0963662516629745

Croce, M. (2022). On testimonial knowledge and its functions. *Synthese*, *200*(2), 141. https://doi.org/10.1007/s11229-022-03528-x

Croce, M. (2025). Expert interviews, good questioning, and the epistemic duties of journalism. *Episteme*, *22*(3), 668–687. https://doi.org/10.1017/epi.2025.24

De Cosmo, L. (2022). Google engineer claims AI chatbot is sentient: Why that matters. *Scientific American*. https://www.scientificamerican.com/article/google-engineer-claims-ai-chatbot-is-sentient-why-that-matters/

de Melo-Martín, I., & Intemann, K. (2009). How do disclosure policies fail? Let us count the ways. *The FASEB Journal*, *23*(6), 1638–1642. https://doi.org/10.1096/fj.08-125963

Dormandy, K., & Grimley, B. (2024). Gatekeeping in science: Lessons from the case of psychology and neuro-linguistic programming. *Social Epistemology*, *38*(3), 392–412. https://doi.org/10.1080/02691728.2024.2326828

Douglas, H. E. (2009). *Science, policy, and the value-free ideal*. University of Pittsburgh Press.

Dunwoody, S. (2012). Scientists, journalists, and the meaning of uncertainty. In S. M. Friedman, S. Dunwoody, & C. L. Rogers (Eds.), *Communicating uncertainty: Media coverage of new and controversial science* (pp. 59–79). Routledge.

Edwards, B. (2025).What is AGI? Nobody agrees, and it's tearing Microsoft and OpenAI apart. *Ars Technica*, https://arstechnica.com/ai/2025/07/agi-may-be-impossible-to-define-and-thats-a-multibillion-dollar-problem/

Fairbank, V. (2025). Should science journalists know science? *Canadian Journal of Philosophy*, *54*(4), 281–299. https://doi.org/10.1017/can.2025.3

Falbo, A. (2021). Inquiry and confirmation. *Analysis*, *81*(4), 622–631. https://doi.org/10.1093/analys/anab037

Falbo, A. (2023). Inquiring minds want to improve. *Australasian Journal of Philosophy*, *101*(2), 298–312. https://doi.org/10.1080/00048402.2021.2024863

Farid, H. (2025). AI-Fueled ignorance, confusion, and profit. In R. N. Proctor, & L. Schiebinger (Eds.), *Ignorance unmasked: Essays in the new science of agnotology* (pp. 57–70). Stanford University Press.

Farrell, H., Gopnik, A., Shalizi, C., & Evans, J. (2025). Large AI models are cultural and social technologies: Implications drawn from the history of transformative information systems. *Science*, *387*(6739), 1153–1156. https://doi.org/10.1126/science.adt9819

Feest, U. (2017). Phenomena and objects of research in the cognitive and behavioral sciences. *Philosophy of Science*, *84*(5), 1165–1176. https://doi.org/10.1086/694155

Figdor, C. (2024). Science journalism and epistemic virtues in science communication: A defense of sincerity, transparency, and honesty. *Episteme*, *21*(4), 1434–1445. https://doi.org/10.1017/epi.2023.38

Friedman, J. (2013). Question-directed attitudes. *Philosophical Perspectives*, *27*(1), 145–174. https://doi.org/10.1111/phpe.12026

Friedman, J. (2017). Why suspend judging? *Noûs*, *51*(2), 302–326. https://doi.org/10.1111/nous.12137

Friedman, J. (2018). Junk beliefs and interest-driven epistemology. *Philosophy and Phenomenological Research*, *97*(3), 568–583. https://doi.org/10.1111/phpr.12381

Friedman, J. (2019). Inquiry and belief. *Noûs*, *53*(2), 296–315. https://doi.org/10.1111/nous.12222

Friedman, J. (2020). The epistemic and the zetetic. *The Philosophical Review*, *129*(4), 501–536. https://doi.org/10.1215/00318108-8540918

Gerken, M. (2022). *Scientific testimony: Its roles in science and society*. Oxford University Press.

Gerken, M. (2025). Scientific values and value-based science reporting. *Philosophical Studies*, *182*(9), 2551–2575. https://doi.org/10.1007/s11098-025-02346-4

Ghaffary, S. (2020). The controversy behind a star Google AI researcher's departure. *Vox*. https://www.vox.com/recode/2020/12/4/22153786/google-timnit-gebru-ethical-ai-jeff-dean-controversy-fired

Ghaffary, S. (2024). AI companies are obsessed with AGI. No one can agree what exactly it is. *Bloomberg*. https://www.bloomberg.com/news/newsletters/2024-01-25/ai-companies-are-obsessed-with-agi-no-one-can-agree-what-exactly-it-is

Goldacre, B. (2014). *Bad pharma: How drug companies mislead doctors and harm patients*. Macmillan.

Goldenberg, M. J. (2021). *Vaccine hesitancy*. University of Pittsburgh Press.

Goldman, A. (1999). *Knowledge in a social world*. Oxford University Press.

Goldman, A. I. (1995). Psychological, social, and epistemic factors in the theory of science. In R. Burian, M. Forbes, & D. Hull (Eds.). *Proceedings of the 1994 Biennial meeting of the philosophy of science association* (pp. 277–286). Philosophy of Science Association.
Goldman, A. I. (2001). Experts: Which ones should you trust? *Philosophy and Phenomenological Research*, *63*(1), 85–110. https://doi.org/10.1111/j.1933-1592.2001.tb00093.x
Grant, N. (2022). Google fires engineer who claims its A.I. is conscious. *The New York Times*, https://www.nytimes.com/2022/07/23/technology/google-engineer-artificial-intelligence.html
Grasswick, H. E. (2010). Scientific and lay communities: Earning epistemic trust through knowledge sharing. *Synthese*, *177*(3), 387–409. https://doi.org/10.1007/s11229-010-9789-0
Grasswick, H. (2019). Reconciling epistemic trust and responsibility. In *Trust in epistemology* (pp. 161–188). Routledge.
Greco, J. (2020). *The transmission of knowledge*. Cambridge University Press.
Greene, T. (2022). The long, uncertain road to artificial general intelligence. *Undark*. https://undark.org/2022/06/02/the-long-uncertain-road-to-artificial-general-intelligence/
Grice, H. P. (1989). *Studies in the way of words*. Harvard University Press.
Hallin, D. C., & Mancini, P. (2004). *Comparing media systems: Three models of media and politics*. *Cambridge University Press*.
Hardin, R. (2006). *Trust*. Polity.
Hardwig, J. (1991). The role of trust in knowledge. *The Journal of Philosophy*, *88*(12), 693–708. https://doi.org/10.2307/2027007
Harman, G. (1986). *Change in view*. MIT Press.
Hart, H. L. A. (1968). *Punishment and responsibility: Essays in the philosophy of law*. Oxford University Press.
Hendriks, F., Kienhues, D., & Bromme, R. (2020). Replication crisis = trust crisis? The effect of successful vs failed replications on laypeople's trust in researchers and research. *Public Understanding of Science*, *29*(3), 270–288. https://doi.org/10.1177/0963662520902383
Henshall, W. (2024). When might AI outsmart us? It depends who you ask. *Time*. https://time.com/6556168/when-ai-outsmart-humans/
Herzog, L. (2024). Corporate knowledge and corporate power. Reining in the power of corporations as epistemic agents. *Critical Review of International Social and Political Philosophy, 27*(3), 363-382. https://doi.org/10.1080/13698230.2022.2113227
Hintikka, J. (1981). On the logic of an interrogative model of scientific inquiry. *Synthese*, *47*(1), 69–83. https://doi.org/10.1007/BF01064266
Hintikka, J. (2007). *Socratic epistemology. Explorations of knowledge-seeking by questioning*. Cambridge University Press.
Hookway, C. (2006). Epistemology and inquiry: The primacy of practice. *Epistemology Futures*, *95*, 110.
Hooper, V. J. (2025). Grok 4 Isn't AGI yet: The hype, the breakthroughs and the dangerous new timeline. *The times of Israel*. https://blogs.timesofisrael.com/grok-4-isnt-agi-yet-the-hype-the-breakthroughs-and-the-dangerous-new-timeline/
Ioannidis, J. (2005). Why most published research findings are false. *PLoS Medicine*, *2*(8), 696–701. https://doi.org/10.1371/journal.pmed.0020124
Iyengar, S., & Simon, A. (1993). News coverage of the Gulf crisis and public opinion: A study of agenda-setting, priming, and framing. *Communication Research*, *20*(3), 365–383. https://doi.org/10.1177/009365093020003002
Jackson, B. B., DiDomenico, D., & Lota, K. (2022). In defense of clutter. *Ergo an Open Access Journal of Philosophy*, *9*https://doi.org/10.3998/ergo.2257
John, S. (2018). Epistemic trust and the ethics of science communication: Against transparency, openness, sincerity and honesty. *Social Epistemology*, *32*(2), 75–87. https://doi.org/10.1080/02691728.2017.1410864
Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. https://doi.org/10.48550/arXiv.1706.03762
Kantrowitz, A. (2022). Google fires Blake Lemoine, engineer who called its AI sentient. *Big Technology*, https://www.bigtechnology.com/p/google-fires-blake-lemoine-engineer
Kahan, D. (2010). Fixing the communications failure. *Nature*, *463*(7279), 296–297. https://doi.org/10.1038/463296a
Keegan, J. (2025). A tech company will claim to have achieved AGI. The news media won't be ready. *Nieman Journalism Lab*. https://www.niemanlab.org/2025/12/a-tech-company-will-claim-to-have-achieved-agi

Kitcher, P. (1990). The division of cognitive labor. *The Journal of Philosophy*, *87*(1), 5–22. https://doi.org/10.2307/2026796
Kitcher, P. (1993). *The advancement of science: Science without legend, objectivity without illusions*. Oxford University Press.
Knight, W. (2025). The AI race has gotten crowded-and China is closing in on the US. *Wired*. https://www.wired.com/story/stanford-study-global-artificial-intelligence-index/
Kreutzer, T., Domgörgen, F., Kaldewey, D., Berger, P., & Wormer, H. (2025). The journalistic understanding of science as process and social system: A qualitative exploration in the German science journalism community. *Public Understanding of Science*, *34*(7), 916–929. https://doi.org/10.1177/09636625251326508
Kvanvig, J. L. (2003). *The value of knowledge and the pursuit of understanding*. Cambridge University Press.
Kwong, J. M. (2023). Gatekeeping the mind. *Inquiry*, 1–23.
Levy, N. (2007). Radically socialized knowledge and conspiracy theories. *Episteme*, *4*(2), 181–192. https://doi.org/10.3366/epi.2007.4.2.181
Levy, N. (2021). *Bad beliefs: Why they happen to good people?* Oxford University Press.
Levy, N. (2022). Do your own research! *Synthese*, *200*(5), 1–19. https://doi.org/10.1007/s11229-022-03793-w
Levy, N. (2024). No Trespassing! Abandoning the novice/expert problem. *Erkenntnis*, *90*(5), 2077–2094. https://doi.org/10.1007/s10670-024-00794-8
Löhr, G. (2023). Conceptual disruption and 21st century technologies: A framework. *Technology in Society*, *74*, 102327. https://doi.org/10.1016/j.techsoc.2023.102327
Longino, H. (2022). What's social about social epistemology? *The Journal of Philosophy*, *119*(4), 169–195. https://doi.org/10.5840/jphil2022119413
Marcus, G. (2022). Artificial general intelligence is not as imminent as you might think. *Scientific American*. https://www.scientificamerican.com/article/artificial-general-intelligence-is-not-as-imminent-as-you-might-think1/
Mascitti, M., Ausili, L., & Bina, F. (2025). Epistemic gatekeeping for public debates. *Unpublished Manuscript*.
Matheson, J., & DiPaolo, J. (2025). Navigating inquiry. In A. B. Creller, & J. Matheson (Eds.), *Inquiry: Philosophical perspectives*. Routledge.
McClain, C., Kennedy, B., Gottfried, J., Anderson, M., & Pasquini, G. (2025). How the U.S. public and AI experts view artificial intelligence. *Pew research Center*. https://www.pewresearch.org/internet/2025/04/03/how-the-us-public-and-ai-experts-view-artificial-intelligence/
McKenna, R. (2020). Persuasion and epistemic paternalism. *Epistemic Paternalism: Conceptions, Justifications, and Implications*, 91–106.
McKenna, R. (2023). *Non-ideal epistemology*. Oxford University Press.
Metzen, H. (2024). Vigilant trust in scientific expertise. *European Journal for Philosophy of Science*, *14*(4), 58. https://doi.org/10.1007/s13194-024-00619-z
Miller, J. D. (1998). The measurement of civic scientific literacy. *Public Understanding of Science*, *7*(3), 203. https://doi.org/10.1088/0963-6625/7/3/001
Millson, J. (2021). Accepting & rejecting questions: First steps toward a bilateralism for erotetic logic. *Asking and answering: Rivalling approaches to interrogative methods*, 211–232.
Millson, A. (2022). Five things Google's AI bot wrote that convinced engineer it was sentient. *The Seattle Times*. https://www.seattletimes.com/business/5-things-googles-ai-bot-wrote-that-convinced-engineer-it-was-sentient/
Muldoon, J., Graham, M., & Cant, C. (2024). *Feeding the machine: The hidden human labor powering AI*. Bloomsbury Publishing.
Murcott, T. (2009). Science journalism: Toppling the priesthood. *Nature*, *459*(7250), 1054–1055. https://doi.org/10.1038/4591054a
Murray, S. (2017). Responsibility and vigilance. *Philosophical Studies*, *174*(2), 507–527. https://doi.org/10.1007/s11098-016-0694-3
Musser, G. (2024). The man who thinks A.I. is sentient. *Critical Opalescence* blog, https://www.criticalopalescence.com/p/is-blake-lemoine-really-all-that
Mede, N. G., Cologna, V., Berger, S., Besley, J., Brick, C., Joubert, M., Schäfer, M. S., … Schäfer, M. S. (2025). Public communication about science across 68 countries: Global evidence on how people get information and communicate about science-related matters. *Unpublished Manuscript*.

Narayanan, A., & Kapoor, S. (2024). *AI snake oil: What artificial intelligence can do, what it can't, and how to tell the difference*. Princeton University Press.
Nguyen, C. T. (2022). Trust as an unquestioning attitude. *Oxford Studies in Epistemology*, *7*, 214–244.
Nichols, T. M. (2017). *The death of expertise: The campaign against established knowledge and why it matters*. Oxford University Press.
Nisbet, M. C., & Mooney, C. (2007). Framing science. *Science*, *316*(5821), 56. https://doi.org/10.1126/science.1142030
Oreskes, N., & Conway, E. M. (2010). *Merchants of doubt: How a handful of scientists obscured the truth on issues from tobacco smoke to global warming*. Bloomsbury Press.
Shoemaker, P. J., Vos, T. P., & Reese, S. D. (2009). Journalists as gatekeepers. In *The handbook of journalism studies* (pp. 73–87). Routledge.
Pepp, J., & Michaelson, E. (2022). Relevance-based knowledge resistance in public conversations. In J. Strömbäck, Å. Wikforss, K. Glüer, T. Lindholm, & H. Oscarsson (Eds.), *Knowledge resistance in high-choice information environments*. Routledge.
Perrigo, B. (2023). OpenAI used Kenyan workers on less than $2 per hour to make ChatGPT less toxic. *TIME*. https://time.com/6247678/openai-chatgpt-kenya-workers/
Pinto, M. F. (2020). Commercial interests and the erosion of trust in science. *Philosophy of Science*, *87*(5), 1003–1013. https://doi.org/10.1086/710521
Pohlhaus, G.Jr (2012). Relational knowing and epistemic injustice: Toward a theory of willful hermeneutical ignorance. *Hypatia*, *27*(4), 715–735. https://doi.org/10.1111/j.1527-2001.2011.01222.x
Proctor, R. N. (2012). *Golden holocaust: Origins of the cigarette catastrophe and the case for abolition*. University of California Press.
Rabenberg, H. (2024). Inquiry and trust: An epistemic balancing act. *Philosophical Studies*, *181*(2–3), 583–601. https://doi.org/10.1007/s11098-024-02128-4
Redish, A. D., Kummerfeld, E., Morris, R. L., & Love, A. C. (2018). Reproducibility failures are essential to scientific inquiry. *Proceedings of the National Academy of Sciences*, *115*(20), 5042–5046. https://doi.org/10.1073/pnas.1806370115
Ren, S., & Wierman, A. (2025). Mitigating the public health impacts of AI data centers. *Harvard business Review*. https://hbr.org/2025/11/mitigating-the-public-health-impacts-of-ai-data-centers
Resnik, D. B. (2000). Financial interests and research bias. *Perspectives on Science*, *8*(3), 255-285. https://doi.org/10.1162/106361400750340497
Roberts, C. (2012). Information structure: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics, 5*, Article 6, 1-69. https://doi.org/10.3765/sp.5.6
Rolin, K. (2015). Values in science: The case of scientific collaboration. *Philosophy of Science*, *82*(2), 157–177. https://doi.org/10.1086/680522
Rolin, K. (2020). Trust in science. In J. Simon (Ed.), *The Routledge Handbook of Trust and philosophy* (pp. 354–366). Routledge.
Roose, K. (2025). Powerful A.I. is coming. We're not ready. *The New York Times*. https://www.nytimes.com/2025/03/14/technology/why-im-feeling-the-agi.html
Rudolph, J. L. (2024). Scientific literacy: Its real origin story and functional role in American education. *Journal of Research in Science Teaching, 61*(3), 519–532. https://doi.org/10.1002/tea.21890
Schipani, V. (2024). Journalism and public trust in science. *Synthese*, *204*(2), 56. https://doi.org/10.1007/s11229-024-04701-0
Shanahan, M. (2024). Talking about large language models. *Communications of the ACM*, *67*(2), 68–79. https://doi.org/10.1145/3624724
Shapiro, I., Brin, C., Spoel, P., & Marshall, L. (2016). Images of essence: Journalists' discourse on the professional "discipline of verification". *Canadian Journal of Communication*, *41*(1), 37–48. https://doi.org/10.22230/cjc.2016v41n1a2929
Siegel, S. (2026). How do lines of inquiry unfold? Insights from journalism. In T. S. Gendler, J. Hawthorne, J. Chung, & A. Worsnip (Eds.), *Oxford studies in epistemology* (Vol. 8, pp. 214–234). Oxford University Press.
Simion, M. (2017). Epistemic norms and 'he said/she said'reporting. *Episteme*, *14*(4), 413–422. https://doi.org/10.1017/epi.2016.11
Simis, M. J., Madden, H., Cacciatore, M. A., & Yeo, S. K. (2016). The lure of rationality: Why does the deficit model persist in science communication? *Public Understanding of Science*, *25*(4), 400–414. https://doi.org/10.1177/0963662516629749
Simon, H. A. (1971). Designing organizations for an information-rich world. In M. Greenberger (Ed.), *Computers, communication, and the public interest*. Johns Hopkins University Press.

Slater, M. H., Scholfield, E. R., & Moore, J. C. (2021). Reporting on science as an ongoing process (or not). *Frontiers in Communication*, *5*, 535474. https://doi.org/10.3389/fcomm.2020.535474
Sperber, D., & Wilson, D. (1986). *Relevance: Communication and cognition*. Harvard University Press.
Sprain, L. (2018). Framing science for democratic engagement. *Ethics and Practice in Science Communication*, 74–90.
Staffel, J. (2019). Credences and suspended judgments as transitional attitudes. *Philosophical Issues*, *29*(1), 281–294. https://doi.org/10.1111/phis.12154
Stalnaker, R. (1978). Assertion. In P. Cole (Ed.), *Syntax and Semantics: Vol. 9. Pragmatics* ( pp. 315–332). Academic Press.
Strevens, M. (2003). The role of the priority rule in science. *The Journal of Philosophy*, *100*(2), 55–79. https://doi.org/10.5840/jphil2003100224
Tangermann, V. (2022). Google suspends engineer who claims the company's experimental AI has become sentient. *Futurism*. https://futurism.com/google-suspends-engineer-ai%09sentient
The Economist Group. (2025). Xi Jinping's plan to overtake America in AI. *The Economist*. https://www.economist.com/china/2025/05/25/xi-jinpings-plan-to-overtake-america-in-ai
Thorstad, D. (2021). Inquiry and the epistemic. *Philosophical Studies*, *178*(9), 2913–2928. https://doi.org/10.1007/s11098-020-01592-y
Tiku, N. (2022). The Google engineer who thinks the company's AI has come to life. *The Washington Post*. https://www.washingtonpost.com/technology/2022/06/11/google-ai-lamda-blake-lemoine/
Vallance, C. (2022). Google engineer says Lamda AI system may have its own feelings. *BBC News*, https://www.bbc.com/news/technology-61784011
Watzl, S. (2022). The ethics of attention: An argument and a framework. In *Salience* (pp. 89–112). Routledge.
Weisberg, M., & Muldoon, R. (2009). Epistemic landscapes and the division of cognitive labor. *Philosophy of Science*, *76*(2), 225–252. https://doi.org/10.1086/644786
Whitcomb, D. (2025). Lysistrata's lament: Interrogative analogues of testimonial injustice. In A. B. Creller, & J. Matheson (Eds.), *Inquiry: Philosophical perspectives*. Routledge.
Williams, J. H., Hooker, C., Gilbert, G. L., Hor, S., & Degeling, C. (2023). Disagreement among experts about public health decision making: Is it polarisation and does it matter? *BMJ Global Health*, *8*(3), e011182. https://doi.org/10.1136/bmjgh-2022-011182
Williamson, T. (2000). *Knowledge and its limits*. Oxford University Press.
Winsberg, E., Huebner, B., & Kukla, R. (2014). Accountability and values in radically collaborative research. *Studies in History & Philosophy of Science Part A*, *46*, 16–23. https://doi.org/10.1016/j.shpsa.2013.11.007
Wu, W. (2024). On attention and norms: An opinionated review of recent work. *Analysis*, *84*(1), 173–201. https://doi.org/10.1093/analys/anad056
Yumuşak, E. (2024). What's in a perspective? Perspectives, interpretation, and inquiry. *Feminist Philosophy Quarterly*, *10*(4). https://doi.org/10.5206/fpq/2024.4.16686
Zeller Jr, T.Blum, D., & Smart, A. (Eds.). (2022). *A tactical guide to science journalism: Lessonsfrom the front lines*. Oxford University Press.
Zollman, K. J. S. (2007). The communication structure of epistemic communities. *Philosophy of Science*, *74*(5), 574–587. https://doi.org/10.1086/525605

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.