1 Introduction

This Element is about the social dimensions of scientific knowledge. Section 2 asks in what ways scientific knowledge is social. Section 3 develops a conception[1] of scientific knowledge that accommodates the insights of the first section and is consonant with mainstream thinking about knowledge in analytic epistemology. Section 4 asks under what conditions we can tell, in the real world, that a consensus in a scientific community amounts to shared scientific knowledge, as characterized in the second section, and how to deal with scientific dissent. Section 5 reviews the ways in which epistemic and social elements interact to coproduce scientific knowledge.

This Element engages with literature from philosophy of science and social epistemology, especially social epistemology of science, as well as Science, Technology, and Society (STS) and analytic epistemology. These disparate scholarly traditions have hardly engaged with each other, and this Element strives to bring them into interaction. I recommend that nonphilosophers, especially from STS, read the sections in reverse order. The Element focuses on themes and debates that date from the start of the second millennium,[2] and on nonformal approaches to social epistemology.[3]

2 The Social Dimensions of Scientific Knowledge

Science is a social enterprise. Researchers collaborate, socially interact with each other, divide scientific tasks among them, and form social structures. While a primary aim of science is generating scientific knowledge, rarely do we find a developed conception of *scientific knowledge*, particularly one that takes seriously its social dimensions. Section 3 proposes such a conception. This section discusses the desiderata from it.

[1] While the view that knowledge is a social phenomenon is prevalent in philosophy of science, feminist, and social epistemology, it's underappreciated in orthodox analytic epistemology (McKenna 2022). An argument against it is that "infants and animals have knowledge … thus the nature of knowledge cannot depend on language use, social needs, or distinctively human values" (Gardiner 2025). But as Gardiner argues, the aim of a social *conception* of knowledge (such as the one I develop in this Element) is to *ascribe* knowledge; namely, to correctly apply the concept of knowledge to cases. Whether knowledge is a social phenomenon or not, "knowledge ascriptions function to meet human needs, and the fundamental features of knowledge ascriptions stem from their roles in human social life. Ascription behaviour is infused with human interests and values" (Gardiner 2025). Additionally, animal knowledge, which exceeds the scope of this Element, also has social dimensions.

[2] Earlier debates about the social dimensions of scientific knowledge focused on patterns of rational theory change; see Lakatos and Musgrave (1970) and Nickles (2021).

[3] For computational social epistemology, see O'Connor (2023).

2.1 Scientific Knowledge: Propositional and True?

Knowledge is recognized in philosophy of science as a central aim of science (Bird 2022). A celebrated tradition in philosophy of science debates the conditions for the "growth of knowledge" (Lakatos and Musgrave 1970), yet one rarely finds a developed account in philosophy of science of that very thing the growth of which philosophers of science debate. While philosophers of science have developed detailed accounts of scientific explanation, scientific understanding, and scientific laws, they have largely neglected scientific knowledge.[4]

Philosophers of science, so it seems, have been content to leave the analysis of knowledge to analytic epistemologists. According to standard accounts of knowledge in analytic epistemology, an epistemic subject S knows that *p* only if

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

The subject S is a single human person, and *p* is a proposition (which is, roughly, the content of a factual claim). For example, Jacob knows that there is an apple on the table only if Jacob believes that there is an apple on the table, his belief is justified; namely, he has visual evidence that supports this belief or his belief was reliably generated by his cognitive system, and his belief is true; namely, there is an apple on the table, rather than, say, a realistic decorative apple made of wax. Standard analyses typically add bridging conditions between (2) and (3), or substitute (3) with a modal (counterfactual) truth-tracking condition or a condition that credits the truth of S's belief to S's epistemic virtue.[5]

Leaving the analysis of knowledge to analytic epistemology, however, may have been a mistake. Standard analytic accounts of knowledge don't smoothly square with how philosophers of science, scientists, science teachers, science students, and the public often think about *scientific* knowledge. For philosophers of science, rather than justified or truth-tracking true belief, scientific knowledge typically means something like "the descriptive content of our best [scientific] theories and models and the skills and concepts required to understand this content" (Chakravartty 2022, 5).

What challenges does scientific knowledge pose to the traditional analyses of knowledge? Regarding the first condition, philosophers of science (notwithstanding Bayesians) are less concerned with a scientist's personal beliefs, and more with what that scientist publicly accepts for the sake of research. As Kent Staley and Aaron Cobb (2011), for example, write:

[4] Notable exceptions are Longino (2002), Roush (2005), and Suppe (1993).
[5] See further discussion in Section 3.4.1.

> While beliefs are certainly relevant to actions particular scientists perform, including the activity of endorsing particular experimental conclusions, the assertion and endorsement of such conclusions in these forums can be distinguished from an individual scientist's beliefs about these assertions. Although not conclusive, these considerations suggest pursuing the idea that scientific knowledge should not be understood as essentially a species of belief. At any rate, whether or not this is the case, one seeking to understand epistemic justification in the sciences and the practices that produce it would be better off looking to what is asserted in the appropriate social contexts than worrying about underlying beliefs, as it is through the interaction of communicative acts that the corpus of scientific knowledge is formed. (p. 479)

Additionally, scientific knowledge isn't necessarily propositional. "Traditional epistemology formulates its problems and answers by thinking of knowledge as primarily propositional. This presupposition should be scrutinized in the light of historical and sociological analyses of cognitive performance and in the light of contemporary theories of human cognition" (Kitcher 1992, 80–81). On top of propositions, scientific knowledge consists of abstract models, pictures, diagrams, graphs (Perini 2012), and even material models (Baird 2004, ch. 2).[6] Sometimes the information contained in them cannot be fully expressed in verbal propositional form, and even when it can, scientists don't bother to do so (Perini 2005). It might be objected that such representations don't constitute knowledge, but only propositions about them. But that isn't how the concept of knowledge is commonly used. When a neat philosophical theory doesn't square with messy reality, we shouldn't reject reality; rather, we should revise the theory.

The nonpropositional form of some scientific knowledge bears on the truth condition. Truth is a property or state of propositions, not of models, diagrams, and such. There are "multiple forms of semantic success, including truth, but also isomorphism, similarity, approximation, and others" (Longino 2022, 176). Unlike truth, which is binary, other semantic relations of fit are gradational, and their adequacy is determined in light of the purposes of the representation (Frigg and Hartmann 2020; Parker 2020).

Another difficulty with truth is that scientific knowledge is tentative: revisable and changeable. Yet scientists aren't quick to deny some past theories the status of knowledge although these theories have been refuted or revised. Scientists similarly don't deny present empirically successful theories the status of knowledge although these theories may be revised or refuted. Newtonian mechanics, for example, is, strictly speaking, a false theory, yet it's still taught

[6] For an overview of these differences, see Kitcher (1992; 1994), Longino (2002, ch. 5), Giere (1988, ch. 3), and Gerken (2019).

and used, since it's a good enough approximation for a wide variety of purposes. It's still part and parcel of the corpus of scientific knowledge.

A further difficulty with the truth condition concerns unobservable entities and processes. Much of scientific knowledge is about unobservable entities, such as protons, or about slow, large-scale processes, such as evolution, which cannot be directly perceived by a human observer. Antirealist philosophers are skeptical or agnostic about the truth of scientific claims about unobservables. Even realists don't think that *all* scientific knowledge about unobservables is true, but only a subset of scientific claims about unobservable entities. Often, they talk about "approximate truth" rather than truth simpliciter (Chakravartty 2017).

More radical worries about truth exist. Some argue that truth, understood as a mind-independent relation of correspondence between a representation and its target system, is irrelevant or unhelpful to the analysis of scientific knowledge. I will address these worries in Sections 3.5 and 3.6.

2.2 Scientific Knowledge: Distributed and Communal

Traditional analyses of knowledge are individualistic. They describe an isolated individual subject without mentioning their social standing. They regard facts about the subject's social standing as irrelevant to the analysis of knowledge. By contrast, social epistemologists regard facts about the subject's membership in an epistemic community; the evidence available within the community; the norms of inquiry, reasoning, and testimony that prevail in that community; and perceptions of risk that stems from error, as having substantive epistemic relevance to any account of scientific knowledge.

A common view in the social epistemology of science is that in some substantive sense, scientific knowledge is communal, and in some respects, communal knowledge is more fundamental than individual knowledge. Namely, an individual's epistemic dependence on members of their epistemic community goes beyond mere reliance on them as informants. Rather, whether that individual's personal beliefs count as knowledge depends on substantive constituents of knowledge that aren't or never have been theirs but are located in other members of their epistemic community or the community itself. An account of scientific knowledge should accommodate the following thesis.

The Epistemic Dependence Thesis

> Whether a subject S knows *p* at time *t* may depend on epistemically substantive elements that other member(s) of S's epistemic community or the community itself possess at *t*, or an epistemically relevant event(s) undergone by some of them, individually or collectively, prior to *t*.

There are different versions of the *Epistemic Dependence Thesis*. The substantive epistemic elements that other members of the community or the community itself possess may be background beliefs (Nelson 1993; Longino 2002), evidence (Hardwig 1985; Miller 2015), shared epistemic standards, shared assessments of the severity of inductive risks (risks that stem from making a wrong epistemic judgment) (Wilholt 2009), or shared models of supportive reasoning (Kuhn 1970; Kusch 2002). Epistemically relevant events undergone by other members of the epistemic community may be belief-generating cognitive processes (Goldberg 2010), exercises of epistemic virtues (Green 2016), or critical scrutiny (Longino 2022). Such events aren't merely prior, necessary enabling conditions for the existence of knowledge, but substantive epistemic conditions that need to be met for scientific knowledge to obtain.

2.3 The Arguments for the Epistemic Dependence Thesis

The Argument from Distributed Cognitive Labor: cognitive processes that generate a single subject's beliefs and confer justification or the status of knowledge on these beliefs may occur largely outside the subject's own cognitive system and be distributed among multiple human subjects that epistemically depend on each other; scientific knowledge cannot be produced without such distribution of cognitive labor (Hutchins 1995; Giere 2002; Longino 2002, ch. 4; Goldberg 2010). Some versions of this argument extend cognition to technological artefacts as well as human subjects (Giere 2003; 2012; Heersmink 2016).

The Argument from Extended Justification: whether a subject's beliefs are sufficiently justified to constitute knowledge may depend on justificatory elements (e.g., evidence, segments of reliable belief-forming processes, cognitive virtues) of other members of their epistemic community whom the subject must explicitly or implicitly trust (Hardwig 1985; Faulkner 2018; Goldberg 2010; 2021; Miller 2015; Miller and Freiman 2020; Pritchard 2010; Green 2016; Dragos 2021; de Ridder 2014; 2019).

The Argument from the Sheer Vastness of Evidence: for some scientific claims, the evidence or arguments that are needed to justify them are so vast, that no individual on their own can review them; different individuals review small parts of evidence, and the substantive evidence by virtue of which such claims constitute knowledge is located only within a social collective (Habgood-Coote and Tanswell 2023; Hardwig 1991; Kukla 2012).

The Argument from Underdetermination of Theory by Evidence: more than one theory can fit the same evidence. Communal background assumptions, which aren't any particular individual's beliefs, stabilize individuals' justified beliefs in, or acceptance of a particular theory among the various possible theories. Additionally, when more than one available theory can accommodate the available data, shared communal values determine which one is accepted and certified as knowledge (Barnes and Bloor 1982; Kuhn 1970; Nelson 1993; Longino 2002; 2016; Potter 1996; compare Section 5.3).

The Argument from the Transformative Role of Criticism: typically, on scientific matters, even a highly reflective, open-minded, individual subject cannot free themself of all their biases and reach knowledge-level justification for their beliefs or claims on their own. To reach knowledge-level justification, hypotheses must undergo communal critical exchange that weeds out individuals' biases and examines a full range of alternative interpretations of the same evidence (Longino 2002, ch. 5; Nagel 1979, 13–14).

The Argument from Shared Reasoning Patterns: reasoning patterns, which must be followed for reasoning conclusions to constitute knowledge, are communal model solutions to cognitive problems (exemplars); for reasoning conclusions to constitute knowledge, the community must also be satisfied that these patterns were correctly followed (Kuhn 1970; Kusch 2002).

The Argument from Conventional Reliability Standards: epistemic reliability standards, which claims must pass to constitute knowledge, are communal conventions that reflect communal rather than individual weighings of inductive risks; without them, scientists cannot assess their peers' reliability and share their findings with each other (Wilholt 2009).

The Argument from Knowledge as a Common Good: some scientific knowledge is an inherently common good, produced and consumed collaboratively by many (Radder 2017).

The Argument from Theory-Ladenness of Observation: individuals' empirical observational beliefs are shaped by and acquire their meaning only through the mitigation of prior socially acquired and accepted theories and observational skills (Kuhn 1970).

2.4 Value-Ladenness and Pragmatic Encroachment

A second thesis this Element endorses is that *scientific knowledge is laden with social values*. Heather Douglas (2000; 2009, ch. 5) has influentially argued that science, specifically, the context of justification,[7] is neither free of social values nor should be.

[7] In the context of discovery, hypotheses are raised and research priorities are set, while in the context of justification, hypotheses are validated. It's less controversial that social values may legitimately influence the context of *discovery* by triggering hypotheses or setting research priorities. It's also acknowledged, however, that scientists should have autonomy to set research agendas. For science–society relations in the context of discovery, see Kitcher (2001; 2011), Keren (2013; 2015), and Jasanoff (2004c).

Douglas distinguishes two roles of social values in epistemic judgments: *direct* and *indirect*. In their direct role, values serve as reasons for making an epistemic judgment, such as theory acceptance. For Douglas, the direct role isn't legitimate in the context of justification, as it amounts to wishful thinking. In their indirect role, values determine the threshold level that evidence must meet for making a justified epistemic judgment by determining the levels of inductive risks we are willing to tolerate. Douglas identifies two types of inductive risks: wrongly accepting a false hypothesis, and wrongly rejecting a true hypothesis. There is an inherent trade-off between them. The more we expose ourselves to one type of risk, the less we expose ourselves to the other type.

Douglas argues that the indirect role is legitimate and required, because social values determine acceptable risks in a given context, and different social circumstances may legitimately require different balances between types of errors. When we value the possible consequences of a risk as mild, we lower the threshold level of evidence required for making a justified judgment, and when the consequences are acute, we raise it. (Section 5.4 reviews other ways values adjust evidential weights.)

Douglas (2009, 50–55) draws on Richard Rudner's (1953) paper "The Scientist *qua* Scientist Makes Value Judgments." The words "*qua* scientist" are important. Nobody denies that in some contexts, scientists make value judgments, for example, when they vote *qua* citizens. Rudner's point is that scientists make value judgments as an integral part of normal scientific practice within the context of justification. Rudner's argument has met criticism. One reply to Rudner is that scientists can explain inductive risks to relevant decision makers who decide which theory to apply. Even when such decision makers are scientists, they don't make them *qua* scientists, but *qua* politicians, civil servants, executives, and so on (Gundersen 2018).

This reply assumes that research consequences evaluation can be deferred to the context of application when inquiry is finished; thus, scientists, who work in the context of justification, needn't evaluate social risks. Douglas argues that this reply is inadequate because a strict separation between basic and applied science, or the contexts of justification and application, is unsustainable. Values penetrate deep into the context of justification and affect various stages of research, such as study design, data analysis, evidence characterization, and evidence interpretation. Analyzing a case study involving scientific research of the carcinogenic effects of dioxin, Douglas illustrates the inevitability of making value judgments in various stages of inquiry prior to theory acceptance. Douglas discusses a series of studies that exposed rats to dioxin. Slides with the rats' liver tissues were observed and characterized to determine if they had developed cancer. Four different studies that used the same slides as data characterized some of them differently, which led to different results regarding the carcinogenicity of dioxin.

Douglas argues that this case illustrates the necessity and inevitability of the indirect role. Values shouldn't influence the characterization of clear evidence. Clear cases of diseased tissues should be characterized as such, and clear cases of healthy tissue should be characterized as such. But values ought to influence the characterization of borderline evidence. In a society mostly concerned with the dangers of cancer, borderline slide cases should be characterized as diseased, and in a society mostly concerned with the economic burden of overregulation, they should be characterized as healthy. Such a practice reflects the types and levels of inductive risk that society is willing to take (Douglas 2000; 2009, 124).[8]

Douglas' claims about an indirect role for values in epistemic judgments in science correspond to, and support a view of *knowledge* known as Pragmatic Encroachment, the proponents of which also draw on Rudner's (1953) classic paper (Fantl and McGrath 2011). Pragmatic Encroachment states that facts about a subject's practical interests regarding a certain content are relevant to determining whether their belief or acceptance of this content passes the threshold of knowledge-level justification. Specifically, if the subject has high stakes regarding the content, they are *ceteris paribus* in a worse position to know it than if they have low or no stakes regarding it. The logical relations between knowledge, justification, and interests according to Pragmatic Encroachment are analogous to the relations that Douglas draws between an epistemic judgment, evidence, and social values in their indirect-role capacity (Miller 2014b).

It might be objected that there is an objective, invariant threshold of justification for knowledge, independent of pragmatic factors. Perhaps such a standard is certainty or near certainty. Surely, if in daylight conditions I clearly see a cup on my desk, I can be certain, hence know that there is a cup on my desk. Such self-contained simple examples of perceptual knowledge of macro objects, however, lack essential features of knowledge that relies on complex real-world inquiry and involves dependence on others (Code 1993). Scientific knowledge often lacks the certainty of knowledge of a cup on a desk, but it's knowledge nevertheless.

Conee and Feldman (2004) suggest that an invariant, nonpragmatic standard for knowledge-level justification would be "along the lines of the legal standard for conviction in criminal cases, proof beyond a reasonable doubt" (p. 296). They argue such a standard would strike the right balance between not being too lax or too stringent. It's unclear, however, why a standard that reflects the

8 For the arguments concerning the indirect role for values in science, see Elliot (2022, 22–28).

weighing of inductive risks in a criminal trial (better exonerate a guilty person than convict an innocent one) would be suitable for all knowledge, including scientific knowledge. As Conee and Feldman (2004), write, their notion of proof is "weaker than a mathematical proof" (p. 296). So why would it do, say, for mathematical knowledge?

Similarly, Gregor Betz (2013) argues that values needn't set a threshold of justification because "there is a class of scientific statements which can be considered – *for all practical purposes* – as established beyond reasonable doubt" (p. 218; emphasis added). But the caveat "for all practical purposes" explicitly states that practical concerns *are* relevant to justified theory acceptance. Betz simply weighs values such that they set the threshold for knowledge very high. Betz only shows that a class of scientific hypotheses trivially satisfies the practical conditions for their justified acceptance; namely, the risks associated with their justified acceptance in conceivable contexts are negligible. It doesn't follow that practical conditions are irrelevant for making justified scientific epistemic judgments (Miller 2014b).

Gerken (2019, 125–127) accepts that pragmatic factors are relevant to theory acceptance but argues that scientists ascribe *knowledge* to each other based on evidence alone. Gerken's objection fails, however. First, the conception of scientific knowledge developed here includes, for good reasons, acceptance and not just belief. But even if we restrict scientific knowledge to belief, knowledge ascription requires determining whether the evidence passes a threshold of knowledge-level justification. As argued, it's hard to see on what nonarbitrary, nonpragmatic basis a threshold that will cover all instances and uses of scientific knowledge can be formulated. Third, on complex scientific matters, such as global climate change, there is no candidate for *belief* other than best *accepted* scientific theories, which are saturated with the value judgments and trade-offs that were made in the inquiry process leading to their acceptance (Miller 2014b).

As I argue in Section 5.4, pragmatic factors can lower and raise evidential thresholds only within a limited range. Thus, while Gerken (2019, 126) is right that it would be weird, for example, to ascribe to a researcher knowledge that molecules X and Y bind *only because* nobody cares (assuming that the evidence is flimsy), when the evidence is stronger, pragmatic factors may make a difference in whether some content amounts to knowledge or not.

Only social values offer a nonarbitrary, principled, and relevant basis to decide the various dilemmas that arise during inquiry and influence its outcomes. Proponents of the Value-Free Ideal have yet to provide a persuasive argument for how one can set an evidential threshold for a justified epistemic judgment without appealing to social values, or how one can defer all these judgments to outside/after the normal process of scientific research. While the