

THE 'EXTENDEDNESS' OF SCIENTIFIC EVIDENCE

Eric Kerr

National University of Singapore

Axel Gelfert

National University of Singapore

Abstract

In recent years, the idea has been gaining ground that our traditional conceptions of knowledge and cognition are unduly limiting, in that they privilege what goes on inside the 'skin and skull' (Clark 1997: 82) of an individual reasoner. Instead, it has been argued, knowledge and cognition need to be understood as embodied (involving both mind and body), situated (being dependent on the complex interplay between the individual and its environment), and extended (that is, continuous with, rather than separate from, the world 'outside'). Whether these various interrelations and dependencies are 'merely' causal, or are in a more fundamental sense constitutive of knowledge and cognition, is as much a matter of controversy as the degree to which they pose a challenge to 'traditional' conceptions of cognition, knowledge and the mind. In this paper we argue that when the idea of 'extendedness' is applied to a core concept in epistemology and the philosophy of science—namely, scientific evidence—things appear to be on a much surer footing. The evidential status of data gathered through extended processes—including its utility as justification or warrant—do not seem to be weakened by virtue of being extended, but instead are often strengthened because of it. Indeed, it is often precisely by virtue of this extendedness that scientific evidence grounds knowledge claims, which individuals may subsequently ascribe to themselves. The functional equivalence between machine-based gathering, filtering, and processing of data and human interpretation and assessment is the crucial factor in deciding whether evidence has been gathered, rather than the distinction between intra- and extracranial processes or individual and social processes (or combinations thereof). To prioritize biological processes here, and to assert the superiority of human cognitive capacities seems both arbitrary and unwarranted with respect to gathering evidence, and ultimately would lead to an unattractive skepticism about many of the methods used in science to gather evidence.

In other words, conceiving of scientific evidence as ‘impersonal’ (or at least not necessarily personal) not only better captures the character of evidence-gathering in practice, but also makes sense of a large amount of evidence-gathering that ‘personal’ accounts fail to either acknowledge or accurately describe. Whilst we suggest it is likely that all internally-distributed evidence-gathering processes are merely contingently internal processes, a significant number of externally-distributed evidence-gathering processes are necessarily externally-distributed. Some evidence can only be gathered by extended epistemic agents.

1. Introduction

In this paper we argue that the concept of ‘extendedness’, which has been controversially discussed within the philosophy of mind and epistemology, offers a natural way of thinking about a core concept in epistemology and the philosophy of science: the concept of *evidence*. In particular, we suggest that such an extended notion of evidence may be palatable even to those who, for various reasons, are unwilling to accept extendedness with respect to the mind or knowledge. We first (Section 2) note a tension that arises for individualist accounts of evidence when applied to scientific knowledge. Such knowledge requires extensive reliance on testimony, which is often the primary mode of accessing it. Yet its original evidence base is often highly distributed, in ways that make it impossible for any one individual to obtain sufficient first-hand evidence for knowledge of—potentially quite simple—scientific claims. In Section 3, we consider a disconnect between prevalent concepts of evidence in philosophy and how the concept is used in science. Philosophical conceptions of evidence and evidence-gathering typically operate at an individualistic level of description and take perception to be a paradigmatic case. By contrast, we identify a number of contexts, especially in science, where the gathering, filtering, and processing of data—and sometimes the assessment of its evidential significance—is outsourced to machines and algorithms (which presumably do not perceive and for which an individualistic level of description may not be appropriate). In Section 4, we turn to a problem in philosophy and cognitive science that motivates the hypothesis of extended cognition (HEC) and remains unresolved in epistemology, namely the delineation of agent and environment. In cognitive science, the problem concerns how to determine what parts of the environment (if any) count as parts of a cognitive process. In epistemology, the divide between internalist and externalist theories of knowledge can be characterized as a problem of how much weight to put on factors internal to an epistemic agent and how much to put on her environment. With respect to evidence, we suggest that,

just as there may be extended cognizers and extended knowers, there may be extended evidence-gatherers. In the light of such findings, we argue (in Section 5), for the epistemic parity between processes of evidence-gathering performed by groups of epistemic agents or hybrids of groups and artefacts ('extended epistemic agents' or 'extended evidence-gatherers' or EEAs) and those performed by individual epistemic agents (IEAs). We discuss solutions to the delineation problem that have been proposed in each discipline and assess whether arguments about extended cognition and extended knowers have any purchase on the idea of extended evidence. Previous attempts to apply the ideas of 'extendedness' and 'distributed cognition' (Brown 2009, 2011; Giere 2002a/b; Hutchins 1995; Magnus 2007; Nersessian et al. 2004) have either focused on science at large and scientific knowledge (as a socially distributed activity with an obvious cognitive function) or on the cognitive role of scientific models, instruments, and measurement devices. Regarding the latter, historians and sociologists of science have amassed numerous case studies demonstrating the indispensability of such objects and artefacts to the creation of scientific knowledge. However, arguing that science is, in practice or in most cases, too complex an activity to be pursued by any one individual is one thing; showing that evidence-based scientific knowledge—as the *output* of this complex activity—is 'socially extended' is quite another. Section 6 develops this line of argument further by arguing that whilst any evidence can, in principle, be gathered by EEAs, some evidence cannot be gathered by IEAs and can only be gathered by EEAs both in practice and in principle. This suggests that, contrary to traditional conceptions of evidence in philosophy, evidence is at best *contingently internal*—when the processes involved are integrated and contained within a single cognizer (as in the typical case of perceptual evidence)—whereas in most cases, evidence is distributed and extended across the cognizer's environment.

2. Epistemic Dependence as a Problem for Individualist Epistemology

Let us begin by noting a dilemma for individualist accounts of knowledge and evidence, which arises from the recognition that, as recent social epistemology has shown time and again, our dependence on others for knowledge and information runs deeper than has traditionally been acknowledged. Of the various forms of 'epistemic dependence' (Hardwig 1985), testimony as a source of knowledge has received by far the greatest share of philosophical attention. More recently, various indirect forms of epistemic dependence have been studied, notably our reliance on our social environment for 'epistemic coverage' (Goldberg 2010)—i.e. for keeping us abreast of important developments and general knowledge of the changing world around us. The term 'epistemic dependence' itself is due to John Hardwig's influential analysis of the role of trust in science. As Hardwig sees it, the kinds of claims we

typically regard as scientific knowledge—that ‘DNA encodes genetic information’, that ‘modern humans first evolved in East Africa’, etc.—are only the end product of a socially distributed process of inquiry and information-sharing. When we, as individuals, hold such claims to be true, we implicitly rely on the cognitive processes of one or more (other) epistemic agents for the formation, sustainment, or reliability of our beliefs. Direct and indirect forms of epistemic dependence—in the form of relying on scientific testimony and the spread of knowledge within the scientific community—thus lie at the heart of science. What warrants scientific claims is a complex web of processes of data gathering, filtering, and processing; aggregation of evidence; and—often late in the process—publication and debate among scholars.

This gives rise to a tension between the resolutely social nature of science *as an activity* and the seemingly unproblematic way in which we credit ourselves with *knowledge* of—sometimes arcane—scientific facts. For, no individual knower could realistically hope to secure any significant portion of scientific knowledge all on her own. Ironically, it is the fact that there is simply too much relevant scientific knowledge and evidence for a single human cognizer to process that precludes fully crediting any one person with it. As a result, Hardwig argues, it would seem that ‘there can no longer be knowledge in many scientific disciplines because there is now too much available evidence.’ (Hardwig 1991: 699) There has been considerable philosophical debate about the implications of this dilemma. Hardwig’s own response is to argue that one can ‘know vicariously’, i.e. without personally possessing evidence for what one knows, by trusting others for knowledge and information. It is only by acknowledging the ineliminable role of trust in science, and specifically the role of testimony as a (non-evidential) source of knowledge, that we can secure *individual* knowledge of (collectively produced) scientific facts. As Hardwig puts it, ‘for finite beings, epistemic interdependence is epistemologically *better* than epistemic independence.’ (Hardwig 1988: 319)

Others have tried to resist Hardwig’s conclusion by arguing that the testimony of others should be regarded as itself a form of evidence: if I encounter a colleague who asserts a particular scientific claim, then, although I lack direct evidence for or against the claim in question, at the very least I gain first-hand evidence *that so-and-so asserted it* (cf. Adler 2002, Ross 1986). Thus understood, testimonial evidence—i.e. evidence of someone’s testimonial assertion—is indeed something I can acquire first-hand. It would, however, be hasty to conclude that the availability of such first-hand evidence is sufficient to ground individual knowledge of whatever is being asserted. Consider coming across a media report of a scientific study on the health effects of caffeine:

‘The staunch individualist would thus insist that the most I can claim to know or rationally believe is that there was a report on NPR about an article linking coffee and heart disease. Unfortunately for me, however, the belief that there

was such a report on NPR is not a belief I'm much interested in. I'm interested in whether my coffee addiction will get me into trouble and, consequently, in whether my newly formed belief that it increases my risk of heart disease is a rational belief.' (Hardwig 1988: 311)

As Hardwig sees it, the attempt to redefine the notion of (first-hand) evidence to include testimonial evidence (in the sense above), is a thinly veiled—and ultimately unsuccessful—attempt to defend a narrow individualist thesis about the connection between evidence and knowledge. Best then, so the argument goes, to hold on to a strong ('first-hand') sense of evidence, as that which is cognitively available to a single human cognizer, and admit that scientific knowledge requires more than evidence, namely trust in others and in their evidential claims. Only by trusting others for *their* reports of evidence, can we 'pool' our collective evidential resources and vindicate the possibility of individual scientific knowledge. Although this account rightly emphasizes the social dimension of science and the importance of trust and testimony, it remains deeply indebted to an individualist conception of *evidence*. As Hardwig puts it, all evidence 'must have once been also *personal* evidence, possessed by at least one member of the [scientific] tradition' (Hardwig 1988: 318). It is this assumption we wish to challenge in the subsequent sections. The idea is not to broaden the categories of admissible evidence in order to vindicate an (obsolete?) picture of an autonomous cognizer bootstrapping her way to individualist knowledge on the basis of first-hand evidence; rather, we wish to suggest that much evidence *itself* extends across the realm of both human and non-human agency. Importantly, the fact of extendedness does not detract from its evidential significance; indeed, as we will see, it is often *in virtue of* its extendedness that scientific evidence grounds knowledge claims, which individuals may subsequently ascribe to themselves.

3. The Concept of Evidence in Epistemology and in Science

The centrality of the concept of evidence in matters of knowledge is difficult to doubt, even by those who reject W. K. Clifford's evidentialist dictum that 'it is wrong, always, everywhere, for anyone to believe anything on insufficient evidence' (Clifford 1879: 186). As a first gloss, something may be considered *evidence* for a proposition p if it counts in favour of the truth of p , makes it more probable that p , or simply confirms that p . For Kim, evidence is 'inseparable' from justification: 'When we talk of "evidence" in an epistemological sense we are talking about justification' (Kim 1993: 226). This initial characterization leaves ample room for different ways of specifying the precise nature of the relationship between evidence and what it supports (e.g., the hypothesis or theory in question). For example, in connection with perception—which is usually regarded as the paradigmatic

case of an ‘on-board’ (Craig 1990: 11) source of knowledge—we sometimes speak of the ‘evidence of the senses’, where the latter is constituted by sense data in the mind of an individual human cognizer. By contrast, when it comes to science, we often conceive of ‘scientific evidence’ in impersonal terms, as including theoretical results and empirical findings (whether these are acknowledged by an individual, group, or the scientific community at large) as well as data recorded by instruments and information stored in databases.

Much epistemological thinking about evidence is heavily indebted to an internalist understanding of what it means for something to ‘be evident’. The paradigmatic philosophical example of evidence is perhaps the perceptual evidence of the individual perceiver. This characterization of evidence posits that evidence is something that an agent possesses: *S* has evidence that *p*. For example, *S* has evidence that the swan is black (the evidence being that the swan appears black to *S*). In its most extreme form, internalism holds that, in order to determine whether one’s belief that *p* is justified, ‘one need only consider one’s own state of mind’ (Chisholm 1989: 76). Thus, Chisholm defines a proposition’s being evident in terms of its maximal (i.e. unsurpassed) reasonableness for a human cognizer. Specifically,

‘A proposition *h* may be said to be *evident* for a subject *S* provided (1) that *h* is reasonable for *S* and (2) that there is no proposition *i* such that it is more reasonable for *S* to believe *i* than it is for him to believe *h*.’ (Chisholm 1966: 22)

The thought that evidence is primarily a matter of how propositions relate to one another carries over to other contemporary definitions. Thus, for Williamson (2000), evidence is simply the totality of propositions that one knows. In philosophy of science, too, the propositional conception of evidence has traditionally had the ‘upper hand’, as is obvious when evidence is described as being *entailed* by a hypothesis—entailment being a logical relation between propositions. The same applies to those definitions that characterize evidence in terms of the effect that one proposition has on (our assessment of) the probability of another. For example, if the truth of proposition *e* would make proposition *h* more probable, *e* should be considered supporting evidence for *h*, whereas if it made it less probable, it would count as evidence against *h*. This idea can be made more precise using conditional probabilities: *e* is evidence for a hypothesis *h*, if and only if $P(h|e) > P(h)$ —i.e. if the probability of *h* given *e* exceeds the probability of *h* prior to knowledge of whether or not *e* obtains. Bayesian epistemology takes this way of thinking about evidence to another level, by modelling belief revision along probabilistic lines, with the ‘probabilities’ in question no longer being interpreted as observed frequencies or as objective probabilities—i.e. as independent of the reasoner—but instead as reflecting a reasoner’s degree of belief. Other recent theories reject the idea that evidence consists primarily of (more or

less) believed propositions, instead allowing for the 'evidence of the senses' to play a central role. Thus, Conee and Feldman hold 'that experiences can be evidence, and beliefs are only derivatively evidence [...] all ultimate evidence is experiential' (Conee and Feldman 2008: 87–88). According to this view, justification is simply a function of an individual's evidence, where evidence is construed as mental states. Such evidentialist *mentalism*, however, is again a strong expression of internalism—perhaps even more so than propositional accounts, given that, as Popper has argued, the latter hold out the promise of an impersonal account of 'objective knowledge' (Popper 1972).

Either way, epistemological definitions of evidence differ in peculiar ways from notions of evidence prevalent in science and in everyday life. The emphasis on the propositional nature of evidence and/or on internal mental representations of the world, appears to be at odds with ordinary and scientific usage, according to which 'evidence' is not exhausted by propositions, beliefs, sense data, and their ilk. Consider the case of forensic evidence: The bloody knife found at the scene of a murder is a concrete thing, not an abstract proposition or mental representation, and whatever information it contains (*qua* being a physical object) is not due to its expressing a proposition.¹ Scientific conceptions of evidence typically rely on externalist ideas. For the scientist, evidence is something objective, measurable and mind-independent in the sense that it is not constituted or constructed in any relevant way by the individual perceiver. This characterization of evidence posits that evidence is something that an agent *S* *acquires*. For example, we might say that '*S* acquires evidence that the fossil is of an Iguanodon'. Finally, there is the evidentiary status of those physical events and phenomena that may be considered 'natural signs'—and, in this sense, *evidence of*—their underlying causes: smoke is evidence of fire, a black cloud indicates the possibility of rain, tears are evidence of a person's being in distress, etc. In those cases, it would seem unnecessarily complicated—and, indeed, unnatural—to insist that the physical processes must first be translated into mental representations on the part of the inquirer before they can acquire the status of evidence.

When characterized in purely abstract terms—whether in terms of conditional probabilities, a proposition's 'being evident' to a specific cognizer, or the entailment of a proposition by a hypothesis—the concept of evidence does not appear to lend itself in any obvious way to extension beyond the realm of propositions and (individual) attitudes towards them. But a moment's reflection on the role of evidence in the sciences shows that an empirically richer conception of evidence is needed. In order for the concept of 'evidence' to be descriptively applicable to scientific practice, it should not gloss over the interplay between the gathering, filtering, and processing of data *as* evidence. It is perhaps telling that abstract accounts of evidence are typically, and more easily, illustrated using examples from the 'hard' sciences. However, while the notion of entailment (of a piece of evidence *E* by a

hypothesis *H*) may be an apt way of describing the relationship between, say, the fundamental (universal) laws of physics and specific observation statements, it would hardly be a good fit for those branches of science that deal with historically contingent complex processes—such as biological evolution—and their products (e.g. biological organisms) which may not strictly conform to relevant law-like regularities. Furthermore, as historians of science have demonstrated time and again, even seemingly clear-cut historical examples of confirmation of a theory by first-hand evidence upon closer inspection often reveal considerable ambiguity in the selection and interpretation of the initial data.² In contexts marked by interdisciplinarity and by the absence of universal laws, it will be even less plausible to reduce the complexities of evidential reasoning to simple relations of the sort envisaged by traditional accounts of evidence in epistemology and philosophy of science.

As an example of the latter kind of context, consider the case of archaeology. As Wylie (2011) has persuasively argued, accounts of evidence in archaeology need to confront ‘the inescapable fact that the evidential significance of archaeological data is an interpretive construct’ (2011: 376). Partly this is due to the kinds of relationships and states of affairs that archaeologists study; partly it is the result of the tenuous (and often underdetermined) causal link between the physical evidence that survives and the ephemeral nature of the processes and historical facts it supports. As Wylie puts it:

‘The nature of the subject domain—the contingencies of social meanings and actions, the complexity of how these relate to material culture and material traces, and the vagaries of preservation—precludes the possibility of establishing inferential premises that could secure interpretive conclusions with the degree of certainty [that would traditionally be required].’ (Wylie 2011: 378)

Similar conditions and trade-offs obtain in other fields of inquiry, not just in the social sciences, but wherever contingent and complex evolving systems are concerned. In areas such as geology, palaeontology, evolutionary biology, and the behavioural sciences, scientific practitioners face a dilemma: either they ‘must confine themselves to the pursuit of narrowly descriptive goals’ (as in what Wylie considers to be the failed case of positivist archaeology), or they ‘must be prepared to embrace the speculative horn of the dilemma’, along with the vulnerability ‘to the free play of contextual values’ that attends it (ibid.: 376–377). While this dilemma seems to present us with a stark choice between, on the one hand, a positivist approach that trades significance for empirical certainty and, on the other hand, freewheeling speculation, it clearly cannot be the last word on the matter. For, as Wylie notes, archaeological evidence is both ‘famously ephemeral and enigmatic, yet resolutely tangible and often epistemically consequential’ (ibid.: 372): Our knowledge of past civilizations and human prehistory *has* expanded vastly

over the past century, on the basis of the very evidence that, according to the dilemma, should be considered extremely fragile. On pain of scepticism about such knowledge, it follows that archaeologists—and, by extension, other scientists facing the same dilemma—must have developed strategies and mechanisms of gathering and handling evidence that render manageable the trade-off identified by the dilemma.

One possible strategy consists in relying on the division of cognitive labour among a group of inquirers, as well as on the checks and balances that operate in a well-integrated community of researchers. This not only reflects the fact that, as Susan Haack puts it, 'a scientist virtually always relies on results achieved by others, from the sedimented work of earlier generations to the latest efforts of his contemporaries' (Haack 2010: 255). It also offers a way of testing the—possibly biased and necessarily subjective—judgments of an individual researcher against those of the community, since 'by having several people make the same observation, they can discriminate the eccentricities of a particular individual's perceptions from what can be perceived by all normal observers' (ibid.: 269). Recognizing science as an essentially social enterprise opens up the possibility of a communitarian response to the problem of scientific evidence. On such a communitarian account, evidential judgments arise from active participation in close-knit communities of knowers, where the latter are based on mutual recognition between specific individuals (following Welbourne 1986). For the social-communitarian response to be plausible, an abstract conception of the 'scientific-community-at-large' will not do. As Haack notes, "'the" scientific community to which philosophers of science sometimes optimistically refer [. . . is really] a constantly shifting congeries of sub-communities' (Haack 2010: 268) and, ideally, agreement within 'the relevant scientific sub-community at a time' (ibid.: 274)—or at least some temporary shared assessment of the current evidence—will correlate with the objective justification of the claim in question: 'The processes by which a scientific community collects, sifts, and weighs evidence are fallible and imperfect [. . .] but they are good enough' (ibid.: 274).

The formation of close-knit sub-communities, whether in the form of research groups, collaborations, or informal networks, is an important strategy scientists have adopted in response to the demands of gathering and assessing evidence. It is, however, not the only available strategy. In the remainder of this section, we shall sketch an alternative approach that does not aim at ever closer integration among the members of a specific sub-community, but instead aims at facilitating a largely *impersonal* exchange of claims and data that carry evidential significance. In order to be able to appreciate the subsequent example, it is worthwhile to reflect on the distinction between (raw) 'data' and (interpreted, propositional) 'claims'. Whereas only the latter fit the traditional philosophical conception of evidence as propositional, arguably it is the former—'uninterpreted inscriptions, graphs recording variation over time, photographs, tables, displays', as Hacking puts it (1992: 48)—which

ultimately ground all our evidence-based claims about the world. Some of the traditional reluctance to grant ‘data’ the status of evidence may have had to do with the thought that evidence requires an element of selection and interpretation, which may seem lacking in the case of ‘raw’ data. However, it is important to realize that any demarcation along those lines is bound to be fuzzy and is eroding quickly, even in disciplines where it may have once been applicable. For example, all scientific detectors, in virtue of their design, embody choices about what kinds of events should count as (potential) evidence. In areas such as experimental particle physics, where particle accelerators produce vast quantities of information, algorithms sort incoming information into ‘noise’ and (potentially significant) ‘events’, with only the latter being recorded as ‘raw data’ deemed worthy of further (human or algorithmic) analysis.

With the advent of cheap DNA sequencing technologies, molecular biology has recently made the transition to a state of overabundance of data and evidence, ranging from snippets of genetic information to whole-organism genomes. Often, the generation of vast amounts of data exceeds what any close-knit subcommunity—such as a specific research group—can hope to analyze. As Sabina Leonelli has noted, ‘a large amount of data produced in the course of experiments is discarded without being circulated to the wider community’, and even when such circulation does take place—for example, via publication in a scientific journal—‘there is little chance that a researcher working in a different area or on a different claim will read the paper, see those data and thus be in a position to evaluate their relevance to their own projects’ (Leonelli 2008: 5–6). This renders Haack’s communitarian ideal of vigorous intra-group discussion as the basis of genuinely *shared* evidential assessments inapplicable. Rather than a single group *making the data their own*, what one finds is the—largely impersonal—standardization and circulation of data, not least via ‘the use of digital databases to gather, organise and distribute the heterogeneous mass of available data’ (ibid.: 2–3):

‘One radical move to “liberate” data from their local context of production has been the construction of public repositories that are available online and collect all data produced in a digital format (e.g. in shot-gun sequencing, micro arrays and in situ experiments), regardless of which of them are used as evidence in publications.’ (ibid.: 9)

Such repositories are routinely consulted by researchers, not primarily with the goal of evaluating the reliability of the information (let alone in order to replicate the experiments that resulted in the recorded evidence), but in order to ‘find out what work has already been done that could potentially inform their research goals’ (ibid.: 11). Datamining—i.e. the practice of extracting basic information from databases and publications—strips away many of the (‘temporal, personal, and social’; Haack 2010: 254) factors that a

social-communitarian account would consider important. Instead, datasets that are 'mined' will be tagged with a 'unique identifier' which is machine-readable and 'makes it possible for data to be computed and analysed through machines (when possible, in automated ways)' (Leonelli 2008: 14). Much of the gathering, filtering, and processing of data in these cases—from the initial sequencing to the algorithmic processing, datamining, and subsequent labelling—occurs automatically and without the critical involvement of a human cognizer. In its most extreme case, human involvement would be limited to the setting-up of the technological infrastructure, the programming of the relevant algorithms, and other auxiliary activities. Even in current implementations of datamining, which still require considerable human intervention (e.g., from human curators who decide on how to 'label' the data), much of this human activity is limited to procedural issues—for example, the standardization of descriptive terminology across subdisciplines (ibid.: 23)—and is not undertaken with an eye towards specific uses of the evidence. It stands to reason that, as the use of datamining becomes more widespread, more and more of this activity will eventually be outsourced to machines and performed by algorithms, rather than human cognizers.

Individualists about evidence might argue that such gathering, filtering, and processing of data by algorithms and machines, although an important part of scientific data-gathering, is itself devoid of evidentiary significance: it becomes evidence only when eventually interpreted by human cognizers. In other words, while delegating part of the processing of data to external machines and algorithms does not invalidate subsequent (human) interpretation and inferences on their basis, the final assessment and judgment of the data—and its use *as evidence*—lies with the end user. This, however, is to misunderstand the character of the automated processing of data and to underestimate the extent to which it constitutes a form of 'interpretation'. As Leonelli puts it,

'users accessing data through the database do not only get the prospective evidence that they need, *but also a specific interpretation of how the terms used as labels in the database refer to objects and processes in the world.* When extracting data from a database, users implicitly agree to use the labels found in the database when formulating claims about phenomena *for which those data serve as evidence.*' (Leonelli 2008: 24; italics added.)

The machine-based gathering, filtering and processing of data, one might say, is *functionally equivalent* to human interpretation and assessment of the evidence. Denying data that has been generated in this way the status of *scientific evidence* would reflect a stipulative, and ultimately unwarranted, attachment to the superiority of human cognitive capacities. It should also be noted that such a denial would deprive us of much of what we consider to be evidence-based knowledge—not least in light of the fact that the

methods outlined already provide the empirical basis of much cutting-edge biological research today. Furthermore, the automation and standardization of data ‘expands the evidential scope of data’ in important ways: ‘It makes data accessible to other research contexts and therefore potentially re-usable as evidence for new claims; and it associates data with a broader range of phenomena than the one to which they were associated in the context of production’ (ibid.: 30). Rather than merely granting that extended—e.g. machine-based and automated—gathering of data is *compatible with* subsequent human interpretation of data as evidence, we wish to suggest that it makes an *active contribution* to the evidential basis of science: one that cannot easily be replicated by limited human cognizers, but would undoubtedly count as a major achievement were it to be replicated in this way.

4. The Environment’s Role in Knowledge and the Agent-Environment Distinction

There is a stubborn problem in epistemology of delineating epistemic agents from their environments which also applies to the case of evidence. We suggest that this difficulty is symptomatic of a more general problem of defining precisely what counts as the agent and what counts as the environment. In recent decades the problem is most sharply visible in, but not limited to, the debate around internalism and externalism in epistemology. Further, whilst the delineation problem, as we shall refer to it, has caused a lot of problems for theories about cognition and knowledge, these problems can be more easily avoided with respect to a core aspect of knowledge acquisition, justification, and inquiry, namely evidence.

The internalism-externalism debate is partly motivated by Gettier cases which purport to show that in certain circumstances changes in an agent’s external environment can disqualify an agent from knowing a given proposition. Consider the following cases adapted slightly from Pritchard (2010):

Pressure Gauge

Barton is an engineer whose job it is to monitor the readings of a pressure gauge at the top of an oil well. If the pressure rises too high there could be an accident and so his job is to report to his supervisor if the pressure gauge displays a reading of x or more. One day the pressure gauge rises well above x and Barton immediately informs his supervisor. The pressure gauge is working well and its reports are accurate. In this case, we assume, we would have no hesitation in saying that Barton has evidence that the pressure in the well is higher than x .

*Pressure Gauge**

The second case is identical to the first but for two things. First, the pressure gauge is broken. It can still display readings but the readings are not reliable. Second, there has been a freak change in the atmospheric pressure that affects the behaviour of the gauge. Barton notices that the pressure has risen above x and immediately notifies his supervisor. Although the pressure gauge is not working as it should, the environment is causing the gauge to always display accurate readings.

Should we say that Barton has evidence in the second case? If we replace the word 'evidence' with 'knowledge', mainstream epistemology would say that Barton does not have knowledge since he has just been lucky that the atmosphere is behaving in this curious manner (See, e.g., Lewis 1996; Kvanvig 2004; Pritchard 2003). Since evidence is closely tied to knowledge and its cognates, we might expect our intuitions around such cases to be similar. And yet, we would suggest that it is more intuitive to say that Barton does have evidence in Pressure Gauge*. We submit that this is because describing evidence as more part of the environment than part of the organism seems more intuitive, *prima facie*, than describing knowledge in the same way. In other words, our concept of evidence is not as dependent on spatial, biological factors as our concept of knowledge. The connection between environmental luck and EEAs is explored in Carter (2013) who suggests there is a tension between two insights: the first being that, when the question is about knowledge, cases such as Pressure Gauge* do not count as knowledge because Barton's 'getting it right' is primarily a matter of luck and the second that privileging processes that take place within the skin and skull of an individual cognizer over those that do not, when the two are analogous in the right way, is unwarranted. This latter view is expressed in Clark and Chalmers' Parity Principle:

Parity Principle

If, as we confront some task, a part of the world functions as a process which, *were it done in the head*, we would have no hesitation in recognizing as part of the cognitive process, then that part of the world *is* (so we claim) part of the cognitive process. (Clark and Chalmers 1998: 8)

The tension is created because it is plausible to suppose that the Parity Principle applies not just to metaphysical questions about cognition but to epistemological questions. Carter calls this the Epistemic Parity Principle:

Epistemic Parity Principle

For agent *S* and belief *p*, if *S* comes to believe *p* by a process which, were it to go on in the head, we would have no hesitation in ascribing knowledge of *p* to *S*, then *S* knows *p*. (Carter 2013: 4203)

Given this, we would expect that we should have similar intuitions about cases where an individual comes to believe *p* by analogous internal processes. The problem is that this creates a tension between the Epistemic Parity Principle and mainstream views on cases of environmental luck. For Epistemic Parity, all that matters is the process by which an individual comes to know a proposition. Recall the two cases from Clark and Chalmers (1998): One in which a person with a normally-functioning biological memory, Inga, had to remember the time of an appointment and another in which an Alzheimer's sufferer, Otto, had to 'remember' the time of an appointment by consulting a notebook he kept. According to the Epistemic Parity Principle, if one grants that Inga has knowledge and that the processes by which Inga and Otto form their beliefs are analogous, then one should grant knowledge to Otto.

Now consider the Epistemic Parity Principle in relation to *Pressure Gauge*-type scenarios. Suppose that, in *Pressure Gauge*, Barton had so many numbers to remember over the course of the day that he keeps a notebook tracking changes in the pressure in the well. Suppose further that when Barton gets back to his office his supervisor asks him for a specific entry from his records. Barton had left his notebook at the well but fortunately could remember the specific entry his supervisor asked him. It is lucky that his supervisor asked for the one he could remember but it is not lucky that he remembered. Now consider another case, *Pressure Gauge***, à la Carter, in which Barton did remember his notebook but it had been sabotaged by a malicious colleague. The colleague changed all the times but accidentally left one untouched. In this case, the supervisor asks for a specific entry that Barton cannot remember. He consults his notebook and reports it to his supervisor. Fortunately, it is the one entry that the colleague neglected to alter. Now, according to Epistemic Parity, if *Pressure Gauge* is a case of knowledge then *Pressure Gauge*** is also a case of knowledge. However, in the first case mainstream thinking in epistemology would say that Barton does know the specific entry (since he is consulting a clear memory which could not easily have been wrong) but in the second case Barton does not know the specific entry (since the colleague could have easily changed that entry and Barton could easily have been wrong). From this perspective, it would seem that changes in an individual's extended environment do have epistemological implications. So if we do not wish to give up insights about epistemic luck we are forced to investigate more closely what is the difference between intracranial and extracranial processes that provide knowledge.

The Parity Principle itself makes a very low demand. All it asks is that if one process that takes place extracranially is analogous (in the right way) to an intracranial process then both are cognitive processes. At most, spatial considerations should not be the determining factor in deciding what is cognitive and what is not. What precisely is required for a process to be the right kind of extracranial analogue has been cashed out in a variety of ways whether it is in terms of 'continuous reciprocal causation' (Clark 2006: 24), 'glue and trust' (Clark 2006: 106; Clark and Chalmers 1998), 'coupled systems' (Clark and Chalmers 1998), whether the right kind of 'ongoing feedback loops' obtain (Palermos 2014: 33), and so on. Each of these aims to specify the conditions under which what would be 'ordinary' accounts be considered a part of the environment becomes part of a cognitive process (including those parts of the environment that are, or could be, part of the spatially internal skin-and-skull cognitive system.) We return to the question of what would count as an extracranial analogue of intracranial evidence-gathering in the next section.

According to process reliabilism (Goldman 2011), an epistemic subject *S*'s true belief that *p* qualifies as knowledge if and only if it is the product of a reliable belief-forming process. Compared with traditional internalist approaches to the problem of epistemic justification, process reliabilism greatly relaxes the demands on the epistemic subject. No longer does *S* need to have access to reflectively available reasons for belief in order for the belief that *p* to count as knowledge: as long as the belief is true and has, in fact, been reliably produced, *S* can be credited with knowledge that *p*. Whereas internalism requires *S* to shoulder the burden of proof, process reliabilism effectively delegates some of the justificatory work to processes that are external to what is reflectively accessible to *S*. In Pressure Gauge and Pressure Gauge*, Barton's evidence is equally reliable and yet whether or not he has evidence in the latter case is unclear, perhaps because things could so easily have been otherwise. We can easily imagine that the atmospheric pressure interfered with the gauge in an unreliable way. It seems lucky that the environment (the atmospheric pressure) intervened in the way that it did. In cases where an agent's epistemic status is in question, our intuition is that lucky changes in epistemically unfriendly environments can disqualify an agent from knowing the proposition. But in cases where an agent's evidential status is in question, our intuition is that lucky changes in similar *evidentially* unfriendly environments do not disqualify the agent from having the evidence. It would seem that reliability is not the determining factor over whether or not Barton has evidence.

Stephen Hetherington has recently argued that the contrast between internalism and externalism may be fruitfully characterized in terms of the attributability of knowledge to the person involved. Internalism essentially 'seeks to attribute knowledge *maximally* to a person': 'It has her self-consciously using good evidence, of which she is or could easily be aware,

to form a true belief; and with this sufficing, other things being equal, for the belief's being knowledge.' (Hetherington 2012: 212) But such an emphasis on maximal attributability quite obviously goes against the spirit of externalism, which seeks to disburden the person by delegating some of the epistemic work to her environment. Strict internalists might flatly deny that knowledge can be obtained by mere reliance on one's epistemic environment, on the grounds that whatever work the environment contributes is not fully attributable to the human cognizer—who, according to this view, is the only 'candidate knower'. Yet this way of 'resolving' the tension between, on the one hand, attributability as a precondition of ascribing knowledge and, on the other hand, our *de facto* dependence on our epistemic environment is unattractive, insofar as it renders unintelligible why we routinely take ourselves to know more—much more—than can be unequivocally attributed to our own epistemic efforts. A more attractive approach would maintain that knowledge can be acquired, but that it may not be solely attributable to the human cognizer. To whom, or what, *is* such knowledge attributable instead? Hetherington suggests that it can be 'for instance, *the person being reliable by consulting a thermometer*—the person-plus-the-thermometer': 'That unity knows; it is the agent of the knowing.' (Hetherington 2012: 213)

This suggestion has implications for a recent trend in reliabilist epistemology to regard knowledge not merely as the output of reliable belief-forming mechanisms, but as the successful exercise of a cognitive ability by an epistemic subject. By combining this *ability intuition* with process reliabilism's emphasis on the generation of reliable beliefs, one arrives at virtue reliabilism, which makes the acquisition of knowledge by an individual a matter of the successful deployment not only of her 'on-board' (Craig 1990: 11) cognitive faculties (such as memory and perception), but also of 'acquired methods of inquiry, including those involving highly specialized training or even advanced technology' (Greco 1999: 287). Provided such acquired methods of inquiry are appropriately integrated within *S*'s cognitive character, the true beliefs thus generated—i.e. the overall cognitive success of the extended system—remains 'to a significant degree creditable to her cognitive agency' (Pritchard 2010: 136–137, cited after Palermos and Pritchard 2013: 112; cf. Green 2012). In certain situations where many individuals (or even many individuals plus artefacts) are involved in gathering evidence and no single individual has overall responsibility for the evidence gathered, it seems more appropriate to credit the group as a whole with the achievement and to consider the reliability of the group-as-a-whole's evidence-gathering processes rather than to choose isolated cases. It is in this respect that virtue reliabilism may be of help.

Virtue reliabilism is sometimes presented as a defence of an individualist conception of knowledge—while at the same time acknowledging that 'the individual agent can be an advanced epistemic agent only within a given social structure necessary for supplying him with the reliable-belief

forming processes that he will later integrate within his cognitive character so as to come to know the truth of some proposition' (Palermos and Pritchard 2013: 115). Thus, Palermos and Pritchard argue that, even though 'the belief-forming process in virtue of which the subject formed his true belief is for the most part external to his organismic cognitive agency, it still counts as one of *his* cognitive abilities' (ibid.: 113; italics added). Yet, the trade-off identified earlier, between attributability (which forms the basis of our crediting specific human agents with knowledge) and our de facto dependence on our epistemic environment (which, in many cases, grounds the reliability of our beliefs) applies no less to virtue reliabilism. In extreme cases, although no human agent may be more deserving of being credited with knowledge than *S*, *S*'s overall contribution to the generation of the justification and knowledge in question may be quite insignificant. In such cases, one may well be tempted to credit the *non-human* elements involved with knowledge—as indeed Hetherington does when he writes that, in such cases, knowledge 'would be attributable to you and partly to aspects of the world beyond your awareness or perspective, such as when it is attributable wholly and only to you-*plus*-the-thermometer' (Hetherington 2012: 216).

Virtue epistemologists argue that knowledge involves cognitive success that is due to an exercise of cognitive ability. Conceiving of knowledge in this way appears to create fewer problems for extendedness than some other epistemological theories. (Cf. Kelp 2013; Pritchard 2010; Palermos and Pritchard 2013; Vaesen 2011) Similarly, we can think of evidence-gathering as a cognitive success that is due to an exercise of cognitive ability. As a result, when an agent gathers evidence, it is appropriate to credit her with a cognitive achievement. Like acquiring knowledge, gathering evidence is not usually a trivial matter. As we shall see in later sections, very often the bulk of the 'work' done in gathering evidence is not done by an individual but by groups of individuals, sometimes extended considerably across time or over technological networks. Virtue reliabilism, fortunately, can accommodate such cases very well since there is nothing in its principles that restricts credit and credit-worthiness to individual organisms. As Palermos and Pritchard have demonstrated, there is nothing in at least one central formulation of virtue reliabilism (due to Pritchard 2010: 136–7) that constrains knowledge-conducive cognitive abilities to those which take place intracranially (Palermos and Pritchard 2013: 7–8).

5. Parity Principle and Paradigmatic Forms of Evidence-Gathering: Perception vs. Extension

Can an individualistic account of evidence—which, as we have seen, is endorsed by 'traditional' and social epistemologists alike—fully capture evidence and evidence-gathering? Our task in this section will be to show that

there are several limitations to such an account which suggest that a socially-extended account better describes not just the norm but a necessary principle of evidence-gathering. Perception is often taken to be a paradigmatic case of evidence (e.g. Haack 2010). Several developments challenge this account. The first is automated computational systems that do the same job, functionally speaking, as individual perceivers and sometimes gather evidence that cannot be gathered by individuals. The following illustration comes from the world of petroleum engineering but it should be clear that many other illustrations could have been provided that would lead to similar conclusions.

Consider a computational system that retrieves data from inside an oil well. Measuring tools are sent down into the well, thousands of feet below sea level, far out of sight of ordinary unaided human perception. These tools record streams of data (e.g., data about temperature, pressure—from which is derived density—velocity, capacitance, etc.) about the fluids in the well. A computer at the surface retrieves this data and records it. This is then sent back to offices off-site and is gathered, filtered, and processed by software applications that then graphically present the data for human analysts to look at and interpret. Up until that point humans have been involved in producing the ‘end product’ only in a trivial way: they physically move the tool into position and release wires that send it down, they work with office computers running software and executing various commands. All these—as in the case of datamining in biology—we can imagine to be automated if it were practical.

Given such a scenario we are forced into two choices: either the end product of this process does not count as evidence or an individualistic, biological account of evidence-gathering is inappropriate for describing this case. It would be hard to maintain that the gathered, filtered, and processed data does not count as evidence: it serves as evidential support for claims about the properties of fluids and rock formations in just the same way that evidence gathered by human observers would. Before the development of this technology, evidence about the properties of fluids and surrounding rock formations was carried out by geologists examining ‘cuttings’—rock samples that are produced and ejected from the well during drilling. Evidence gathered by these geologists is perceptual—or observational—in a straightforward manner. They inspect the rocks and carry out some simple observational tests to draw conclusions about the potentially hydrocarbon-bearing formations below ground. In both cases, evidence-gatherers (whether biological or technological) are acquiring evidence for beliefs about certain empirical facts. The biological case (the geologists with their rock samples) is about gathering perceptual evidence; the technological case (the network of tools, instruments, and software) seems to be about gathering evidence that is sourced differently, but is nonetheless evidence.

The second development that challenges the individualistic account of perceptual evidence comes from HEC and research in cognitive science:

specifically, that it is not at all obvious that perception is necessarily individual i.e. some perception appears to be extended. Elisabeth Pacherie (1995) has argued that there are no stable criteria for distinguishing between the natural perceptual instruments in the skull and artefactual perceptual instruments such as those developed by microbiologists, petroleum engineers, and so on. Pacherie refers to instruments used in generating knowledge as 'epistemic artefacts' and compares their contribution to the one made by our naturally-endowed epistemic organs or systems. In particular, she focuses on microscopes and compares these with the visual organs and systems. She notes that all perception is indirect in that it involves acquired skills and learned knowledge beyond the innate abilities we are born with. Since both microscope and ordinary images 'carry information about the spatial properties of distal layouts' there are no stable, non-stipulative grounds for distinguishing between 'seeing with eyes' and 'seeing with microscopes' (Pacherie 1995: 182). Whether or not this argument is metaphysically sound, it seems plausible with regard to evidence-gathering. For ordinary unaided vision, reliability of evidence is guaranteed by the reliability of the organism's properly functioning perceptual faculties. For vision extended by epistemic artefacts, reliability of evidence is guaranteed by the reliability of properly functioning technical artefacts. It seems *bioprejudicial* to constrain reliability conditions only to the function of organic body parts and exclude artificial parts. The evidence itself gathered in both cases is the same and so, given that the process is analogous, we ought to conclude that no non-stipulative, non-bioprejudicial distinction can be drawn between internally and externally distributed perception.

To illustrate this claim, consider a typical case of visual perception: an individual agent sees that an object is blue. She does this without the aid of any technology. In another case the agent is still acting alone but wearing spectacles or contact lenses. It seems clear that this does not affect the epistemic status of her belief. In another case the agent is looking through a video recorder at the object. The time-lapse would seem to be the only thing that suggests that her epistemic status with respect to the colour of the object is in any significant sense compromised. In the last case, she uses a colourimeter to acquire knowledge of the object's colour. The connection between the agent and the artefact (spectacles, contact lenses, video recorders, colourimeters) and between the agent and the environment (thick lenses, cloudy days, atmospheric changes) are not typically taken to affect the agent's epistemic status—assuming they do not affect the reliability of said status—and so it looks like the same should be true of her evidential status. Her evidential status can be extended across artefacts and environments and remain intact.

There may be situations where a form of evidential parity is relevant. Consider again the Parity Principle as stated earlier:

Parity Principle

If, as we confront some task, a part of the world functions as a process which, *were it done in the head*, we would have no hesitation in recognizing as part of the cognitive process, then that part of the world *is* (so we claim) part of the cognitive process. (Clark and Chalmers 1998: 8)

There are some strong arguments against (and in favour of) metaphysical and epistemic parity. (See, e.g., Carter 2013.) However, evidential parity seems far more robust.

Evidential Parity Principle

If a process that, were it performed intracranially or by an individual epistemic agent, we would have no hesitation in recognizing as an evidence-gathering process, were performed extracranially or by an extended epistemic agent, then that process is an evidence-gathering process.

Suppose that within the offices of the engineering firm we spoke of earlier, evidence has been gathered by one of these networks of humans, artefacts and environments. The evidence has been collated by some central software that gathers, filters, and processes it, producing more evidence. Each separate set of evidence provided to the software was gathered and perceived by each individual human agent separately, but no single individual has gathered or perceived the output of the central software. What do we say about this evidence? Who do we credit with gathering it? It does not seem appropriate to credit it with the analyst who merely ‘logs on’ at the end of this process and collects the machine’s output. After all, it may be the case that this analyst has seen none of the previous data and may not even properly understand what they have collated and processed. They may be following simple, procedural rules. Suppose that it is possible, given time and resources, for an individual to gather exactly the same evidence that the extended epistemic agent gathered. This individual can take all the measurements, execute all the calculations, and derive all the required data that the EEA did. In such a case, evidential parity states that were we to give credit to the individual agent for such achievements, we should likewise give credit to the EEA when appropriate. Evidence, if you like, does not care if it is gathered by an individual, a group, or even a distributed computational system or extended human/non-human hybrid.

However, this only shows that an EEA can, in some cases, be credited with gathering evidence. There may even be a case to be made that

this is the norm in evidence-gathering (i.e. that most evidence is gathered by EEAs) although that empirical matter will have to be settled elsewhere. There is a stronger claim to be made that the Evidential Parity Principle is asymmetrical. Consider again the case of perception. Let us suppose that, typically, perception is internally distributed (i.e. it takes place within an individual organism). It remains the case that such perception *could* be externally distributed. That is, for any perceptual evidence that is gathered by an individual, such evidence could also have been gathered by an EEA.³ On the other hand, some evidence that is socially distributed cannot be internally distributed i.e. some evidence can be gathered only by EEAs.

In the engineering example above, it may be true that, in principle, an individual human agent could trawl through the data streams, performing the same algorithms and calculations, drawing the same graphs, filtering out the same outliers, and so on, that the computational system did. In principle, then, any evidence gathered could be reducible to the individuals that gathered it. It is just that, in practice, this does not happen. So perhaps all cases of evidence gathering are describable at the individual level but practical limitations mean that it is more efficient to use artificial systems? There are cases where evidence may be accessible in principle, but in practice it is so *impractical* that we might say that, for all intents and purposes, the practical/in-principle distinction is a spurious one. Consider the hypothetical 'deterministic machine' that has the computational power and all the data it needs to calculate every future event in the universe from currently available information about the earliest moments of the universe. We might suppose that a limitation of this machine is that any sufficiently sophisticated computer would take so long to calculate future events that such events were no longer in the future but had already occurred. Arguably then, the machine can, in principle, perform its function but the practical limitations are so limiting as to render the function redundant. In other words, some practical limitations hold *in principle*. If individual processing of the relevant data for a particular scientific question would require multiple—perhaps thousands of—human lifetimes, a 'merely' practical limitation quickly becomes a *de facto in-principle* limitation. If we are to make sense of scientific knowledge—and of how it is that we, as individuals, can legitimately credit ourselves with at least some such knowledge—we need to acknowledge that scientific evidence itself depends on the existence of reliable (external) networks of human reasoners and non-human actors (e.g. technical artefacts, measurement instruments, computational devices, information filters etc.). That is, it is unavoidable that certain forms of evidence-gathering need to be performed by and credited to EEAs, and descriptions at the externally-distributed level are not reducible in principle to descriptions at the internally-distributed level (a situation which Bedau calls 'in principle irreducibility in practice'; Bedau 1997: 449).

6. The Character of Extended Evidence

At the beginning of this paper, we suggested that the concept of ‘extendedness’ may be an apt lens through which to approach the problem of evidence—i.e. the disconnect that appears to exist between how philosophers have traditionally characterized the concept of evidence and how scientists employ the term. In particular, we expressed our hope that even those who, for various reasons, reject the concept of extendedness with respect to mind or knowledge may be able to accept the idea of extended evidence.

As an example, consider Ronald Giere’s critique of extendedness as applied to mental concepts. Giere (2007) invites the reader to consider, as an example of a distributed cognitive system, the Hubble Space Telescope. Not only does the successful operation of the Hubble telescope require coordination among a vast number of researchers and the successful functioning of a multiplicity of technical components, but some of the imaging techniques for distant objects also draw on other parts of nature. For example, some images ‘were produced by utilizing a cluster of galaxies, Abell 1689, as a gravitational lens’: ‘Abell 1689 is itself 2.2 billion light-years out into space, yet it was cleverly incorporated into the distributed cognitive system that produced the final images’ (Giere 2007: 317). This fact, Giere argues, causes substantial problems for attempts to credit the corresponding extended system with knowledge, consciousness, or agency:

‘If we treat the Hubble system as itself an epistemic agent with a mind of its own, it seems we would have to say that its mind extends from the Earth 2.2 billion light years out into space[.] Just how fast do intentions propagate? Do minds operate at the speed of light?’ (ibid.)

What these (rhetorical) questions convey is a certain unease regarding the very idea that concepts such as ‘knowledge’ and ‘intention’, which essentially depend on the more general concept of ‘mind’, should be extended to include objects and structures that are not in any obvious way capable of bearing mental states.

Yet Giere acknowledges that not all cases of extending our (human-centric) concepts to (non-human) structures (or hybrids consisting of human and non-human elements) are equally problematic: ‘There is no doubt that *some* extensions of concepts originating with humans beyond the bounds of biological agents are natural, even helpful.’ (Giere 2007: 317) A case in point, Giere argues, is memory: ‘Modern civilization, as well as modern science, would be impossible without various forms of record-keeping that are usefully characterized as external memory devices.’ (ibid.) Similarly unproblematic continuities can be established between some of our other organismic capacities, such as perception, and their technological extensions (e.g. detectors, measurement instruments, etc.). Significantly, the concrete

examples of extendedness that Giere endorses (record-keeping, storage devices, etc.) tend to coincide with the kinds of technologies scientists would consider repositories of data and evidence. In many cases, data only gets stored after significant automated filtering and cross-checking. For example, in experiments with particle accelerators, algorithms filter out background noise as well as frequent (but uninteresting) interactions between particles, focusing instead on 'events'—i.e. unusual occurrences that are deemed by the algorithm to be interesting and potentially novel.

Yet both traditional epistemology and most social-epistemological frameworks reject the thought that *evidence* may be 'out there'—i.e. may exist in the mode of extendedness—and instead defend an individualist conception of evidence and the warrant that derives from it. Thus, Haack argues that a theory of evidence 'must begin with the personal, and then move to the social, before it can get to grips with the impersonal sense in which we speak of a well-warranted theory or an ill-founded conjecture' (Haack 2001: 257). In this regard, she is indebted to a long tradition (mentioned in Section 3) of thinking of evidence as purely a matter of one's mental states or, more generally, as that which is cognitively accessible to a single human cognizer (See Conee and Feldman 2004, and references therein.) On this view, any talk of evidence that suggests an impersonal or extended dimension of evidence 'must be understood as an elliptical way of saying that it is well or poorly warranted by the evidence possessed *by some person or some group of people* at that time.' (ibid.: 271; italics added.) The basic idea of such evidential individualism seems to be that nothing that does not derive from, or has not passed through, a human mind qualifies as evidence. This contrasts with the Evidential Parity Principle defended in the previous section, according to which, loosely speaking, something may count as evidence even if it is never entertained by a single human mind—provided that, *were it to be instantiated inside a single mind*, we would have no difficulty recognizing it as a piece of evidence.

The examples we encountered earlier in such diverse areas as biology, petroleum engineering, and archaeology, are such that, if they were to be instantiated in a single mind, we would have no trouble recognizing them for what they are: processes of evidence-gathering. We know, of course, enough about these processes to realize that there is no realistic chance of replicating them 'intracranially', as it were. But the very same knowledge that leads us to recognize this impossibility also tells us that the way they operate is directly analogous, and functionally equivalent, to processes of observation, interpretation and assessment carried about by human cognizers. The data thus generated serves as evidential support for claims about various aspects of the world, in just the same way that evidence gathered directly by human observers would. Denying the deliverances of such extended, technologically mediated systems the status of evidence, thus, would at best reflect a

bioprejudicial attitude which unduly privileges the contingent way in which we process information about the environment ‘first-hand’.

What, in the examples discussed, is the contribution that human cognizers and their ‘on-board’ cognitive capacities make to the process of gathering evidence? To be sure, there remains a non-negligible element of human involvement—for example, when (human) curators classify incoming data, or when end users read off and interpret graphical representations of automatically processed data. But the fact remains that the bulk of the work done in such evidence-gathering systems is not carried out by individual human curators or end users, but by groups and hybrids of epistemic agents, artefacts, algorithms, technological infrastructures, etc. From a virtue-reliabilist perspective that wishes to do justice to the *actual* sources of reliability, credit should go where credit is due: to the extended evidence-gathering systems that guarantee the reliability of the information they generate through the proper functioning of the system. Where there is non-negligible human involvement, some credit will, of course, also go to individual human cognizers, but there is no reason to assume that their contribution is always going to be paramount.

One of the more radical implications of the idea of extendedness is precisely ‘the possibility of an epistemic agent in a particular case being a person-*plus-more-of-the-world*, or even (more radically) being a part-of-the-world-*minus-any-person*’. This might encourage us ‘to regard “It is known that *p*” (rather than “you know that *p*”, say) as at least sometimes the literally correct way of reporting an instance of knowledge’ (Hetherington 2012: 215–216). Critics of the extended knowledge hypothesis might reply that such a locution should at best be understood metaphorically, since attributing knowledge (or partial knowledge) to objects incapable of mental states would otherwise require rethinking the very foundations of our concept of knowledge. Rethinking the notion of *evidence* through the lens of ‘extendedness’, however, seems considerably less problematic. As we see it, it is entirely acceptable, and consonant with established usage, to speak—in an impersonal way—of there ‘being evidence that *p*’ (rather than someone in particular ‘having evidence that *p*’). In the examples discussed above, a particular dataset that has been gathered, filtered, and processed constitutes evidence, irrespective of whether any human cognizer has reflected on its evidential status or has used it to support a particular claim: the evidence-gathering infrastructure is simply set up in a way that ensures the data’s evidential significance (though, needless to say, does not guarantee its infallibility).

In at least some cases, then, evidence-gathering is extended in ways that render human involvement secondary or peripheral. However, in addition to this weaker claim, in what follows we wish to argue for two stronger claims. First, the extendedness of evidence-gathering processes often makes a positive contribution to the warrant of the corresponding claims. That is, it allows

for mechanisms which strengthen, rather than dilute, the evidential basis of various claims and which cannot easily be replicated by limited human cognizers. Second, extendedness is more descriptively adequate, in that it is the norm in science with respect to processes of generating and gathering evidence, and sometimes is the only way to describe certain cases of evidence. We argue that, as a corollary, it follows that whilst all 'internally-distributed' cases of evidence-gathering by individual cognizers are merely contingently internal, some cases of evidence-gathering are necessarily external.

Let us first turn to the ways in which extendedness might strengthen, rather than weaken, the evidential basis of a claim, hypothesis, or theory. In particular, we wish to discuss two such examples, each of which illustrates a more general point. The first example relates to those extended processes that gather evidence across a number of different domains (e.g., by drawing on methods and techniques from largely independent subdisciplines). Consider the case of archaeology, which we discussed earlier (Section 3) in connection with Wylie's account of evidence. At first sight, the fact that 'archaeological material is both so rich and persistently perplexing' (Wylie 2011: 375) might seem to add to its fragility as evidence, not least as the result of problems with interpreting and weighing different kinds of material against each other. However, assuming that evidence from different domains and sources is at least partially independent, any existing convergence between them will itself constitute evidence in support of their reliability (and, by extension, of the claims upon which the various types of evidence converge). The case is similar to that of corroborating testimony from independent eyewitnesses: as more and more eyewitness accounts agree on the point in question, it becomes increasingly unlikely that they have no basis in reality. Likewise, as more and more evidence from independent experimental methods and theoretical domains begins to support the same conclusions, it becomes less and less plausible to maintain that such agreement is merely the result of chance. As Wylie notes, 'the many different kinds of "theory" that play functionally different roles in archaeological inference are disjoint in epistemically consequential ways' (ibid.: 380), namely by lending support to those hypotheses with respect to which they converge. Convergence of evidence across different theoretical and experimental domains is all the more significant when the evidence is produced by different *independent* researchers and is generated using independent methodologies. In other words, the very fact that the gathering of convergent evidence is distributed over many (human and non-human) agents may itself constitute evidence. While those intent on defending an individualistic conception of evidence might argue that someone could individually come to know of the condition of distributedness (e.g., by being told), even an individual reasoner's appreciation of the significance of convergence rests on the recognition that there is additional—distributed—evidence further 'upstream', as it were.

The second example of how distributedness may contribute to the evidential status of claims is randomized controlled trials as routinely carried out in the biomedical sciences. As P. D. Magnus notes, such trials ‘rely on a network of agents each with partial information, so as to mitigate unconscious bias and placebo effects’ that might otherwise conflict with the task of following double-blind methodology: ‘The organizers of the trial would be rightly chastised if the process failed to implement the task.’ Yet, in virtue of how their task is *defined*—namely, as conducting a double-blind study—such trials ‘ipso facto cannot be carried out within a mind’ (Magnus 2007: 301). One might worry that this precludes the application of the Parity Principle and the characterization of drug trials as a form of *distributed cognition* (which is Magnus’s concern in the paper). However, as Magnus notes, whether or not a process or procedure is accurately described as a form of distributed cognition, depends on the specification of the task: ‘If the task is specified as conducting a double-blind trial, the drug trial is not d-cog; if the task is specified as determining how safe the drug is, then the trial is d-cog’ (Magnus 2007: 302). These considerations apply equally to the scientific evidence that is being produced by randomized double-blind trials. No single individual could implement a double-blind trial entirely on their own: given the definition of what double-blind methodology requires, this is simply a procedural impossibility. However, *if* an individual were to collate the various experimental results that are, in fact, being generated during such a trial, he or she could certainly come to appreciate their probative force—including the evidential significance of the fact that the various pieces of evidence were generated *independently*. But, in the case of an *actual* double-blind trial, it would be quite mistaken to credit an individual making such a final assessment with being the first to generate scientific evidence in the case at hand: such evidence as may be gleaned from double-blind trials is due to the various processes of administering the drugs, recording their effects, and transcribing the information, most of which are distributed across many individuals as well as across non-human agents.⁴

We have now seen various cases of evidence-gathering which are either internally or externally distributed i.e. which are accomplished either through the on-board faculties of an individual cognizer or through an extended network of individuals and epistemic artefacts. We have argued that a process need not be ‘personal’ (in Haack’s or Hardwig’s sense) in order for it to count as evidence-gathering. We have also argued for the appeal of evidential parity and for not privileging spatial or biological determinations in deciding what counts as evidence. A process that, if performed intracranially or by an individual epistemic agent, we would have no hesitation in recognizing as an evidence-gathering process, is an evidence-gathering process, even when performed extracranially or by an extended epistemic agent. Further, this parity is asymmetrical. Cases of evidence-gathering that are performed intracranially are only contingently internally distributed whereas

some extracranial processes of evidence-gathering are necessarily externally distributed. For any case of internally-distributed evidence-gathering we can conceive of an externally-distributed analogue. On the other hand, we cannot do the same for all externally-distributed evidence-gathering. Some cases may be practically irreducible to internal processes (such as evidence-gathering in petroleum engineering and geology), some may be 'in principle irreducible in practice' (such as evidence-gathering at CERN), and some may be irreducible in principle (such as evidence of convergence of evidence or double-blind trials).

Further, whilst the current paradigm of perceptual cases of evidence suggests that individual cases are the norm in evidence-gathering, we have suggested that the very opposite is most likely the case. The majority of scientific evidence-gathering depends upon irreducibly extended processes. Rather than intracranial processes, the norm of evidence-gathering appears to be some combination of these extracranial processes and so only the latter can hope to comprehensively characterize the nature of evidence and evidence-gathering in science and technology. In sum not only can some evidence *only* be gathered by EEAs, *most* evidence is. Given this, the burden of proof is on individualists to explain why descriptions at the individual, intracranial level are the most appropriate for characterizing evidence and evidence-gathering. Of course, our position also accounts for the descriptions of evidence-gathering at the individual level, but only as a special case of the much wider phenomenon of extended evidence-gathering. Hardwig, Haack, *et al.*, rightly argue for a social dimension to scientific knowledge but at the same time hold on to a quasi-perceptual concept of evidence that does not tally with the norm in scientific evidence-gathering. Here we have outlined an account that can preserve the perceptual case but goes beyond it to more comprehensively describe evidence and evidence-gathering.

Notes

1. One might even argue that the layperson's conception, according to which the knife simply *is* evidence, and the forensic scientist's perspective, according to which the knife *contains* evidence—which needs to be extracted using various forensic methods—have far more in common with each other than they do with the narrow legal conception, which holds that 'scientific evidence' consists of the—propositional—testimony of expert witnesses.
2. Examples include the debate about what Galileo could, or could not, see when looking through his telescope (Winkler and van Helden 1992) and the case of the spurious 'discovery' of N-rays—a putative new form of electromagnetic waves that were thought to manifest themselves in increases in the brightness of electric arcs, which led to a flurry of papers being published by numerous researchers between 1903 and 1906 (Nye 1980).

3. For example, we could hook up the individual perceiver to a machine that reads the perceptual data from the individual's brain and displays it on a screen or even directly into the brain of another perceiver.
4. Note that the point of this example is not to show that evidence from random controlled trials is always superior to 'personal' evidence—such as first-hand observation on the basis of clinical experience—but merely that it also constitutes a form of evidence; on the question of the meaning of 'evidence' in evidence-based medicine, see also Worrall 2002.

References

- Adler, J.E. 2002. *Belief's Own Ethics*. Cambridge, Mass.: The MIT Press.
- Bedau, M.A. 1997. Weak Emergence, *Philosophical Perspectives* 11 (s11): 375–399.
- Brown, M.J. 2009. Models and Perspectives on Stage: Remarks on Giere's Scientific Perspectivism, *Studies in History and Philosophy of Science Part A* 40(2): 213–220.
- Brown, M.J. 2011. Science as Socially Distributed Cognition: Bridging Philosophy and Sociology of Science, in K. François, B. Löwe, T. Müller, & B. van Kerkhove (Eds.), *Foundations of the Formal Sciences VII: Bringing Together Philosophy and Sociology of Science*. v. 32 of Studies in Logic. London: College Publications, 17–31.
- Carter, J.A. 2013. Extended Cognition and Epistemic Luck, *Synthese* 190(18): 4201–4214.
- Chisholm, R.M. 1989 [1966]. *Theory of Knowledge*, 3rd Edition [1st Edition published 1966], Englewood Cliffs: Prentice Hall.
- Clark, A. 1997. *Being There: Putting Mind, Body, and World Together Again*. Cambridge, Mass.: The MIT Press.
- Clark, A. 2006. Coupling, Constitution and the Cognitive Kind: A Reply to Adams and Aizawa, in R. Menery (Ed.) *The Extended Mind*. Cambridge, Mass.: The MIT Press, 81–100.
- Clark, A. and Chalmers, D. 1998. The Extended Mind, *Analysis* 58(1): 7–19.
- Clifford, W.K. 1879. *Lectures and Essays*. London: Macmillan.
- Conee, E. and Feldman, R. 2004. *Evidentialism*. Oxford: Oxford University Press.
- Conee, E. and Feldman, R. 2008. Evidence, in Q. Smith (Ed.) *Epistemology: New Essays*. New York: Oxford University Press, 83–104.
- Craig, E. 1990. *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford: Oxford University Press.
- Giere, R.N. 2002a. Discussion Note: Distributed Cognition and Epistemic Cultures, *Philosophy of Science* 69(4): 637–644.
- Giere, R.N. 2002b. Scientific Cognition as Distributed Cognition, in P. Carruthers, S. Stich, & M. Siegal (Eds.) *The Cognitive Basis of Science*. Cambridge: Cambridge University Press, 285–299.
- Giere, R.N. 2007. Distributed Cognition without Distributed Knowing, *Social Epistemology* 21(3): 313–320.
- Goldberg, S. 2010. *Relying on Others: An Essay in Epistemology*. Oxford: Oxford University Press.
- Goldman, A. 2011. Reliabilism, in E.N. Zalta (Ed.) *The Stanford Encyclopedia of Philosophy* (Spring 2011 Edition), URL=<http://plato.stanford.edu/archives/spr2011/entries/reliabilism/>
- Green, A. 2012. Extending the Credit Theory of Knowledge, *Philosophical Explorations* 15(2): 121–132.
- Haack, S. 2010. Clues to the Puzzle of Scientific Evidence, *Principia* 5(1–2): 253–281.
- Hacking, I. 1992. The Self-Vindication of the Laboratory Sciences, in A. Pickering (Ed.) *Science as Practice and Culture*. Chicago: The University of Chicago Press, 29–64.
- Hardwig, J. 1985. Epistemic Dependence, *The Journal of Philosophy* 82(7): 335–349.

- Hardwig, J. 1988. Evidence, Testimony, and the Problem of Individualism: A Response to Schmitt, *Social Epistemology* 2(4): 309–321.
- Hardwig, J. 1991. The Role of Trust in Knowledge, *The Journal of Philosophy* 88(12): 693–720.
- Hetherington, S. 2012. The Extended Knower, *Philosophical Explorations* 15(2): 207–218.
- Hutchins, E. 1995. *Cognition in the Wild*. Cambridge, Mass.: The MIT Press.
- Kelp, C. 2013. Extended Cognition and Robust Virtue Epistemology, *Erkenntnis* 78: 245–252.
- Kim, J. 1993. What is Naturalized Epistemology? In J. Kim, *Supervenience and Mind: Selected Philosophical Essays*. Cambridge: Cambridge University Press, 216–236.
- Kvanvig, J. 2004. Nozickian Epistemology and the Value of Knowledge, *Philosophical Issues* 14(1): 201–218.
- Leonelli, S. 2008. Circulating Evidence Across Research Contexts: The Locality of Data and Claims in Model Organism Research. *Working Papers on The Nature of Evidence: How Well Do "Facts" Travel?* No. 25/08. London: Department of Economic History, London School of Economics.
- Lewis, D. 1996. Elusive Knowledge, *Australasian Journal of Philosophy* 74(4): 549–567.
- Magnus, P.D. 2007. Distributed Cognition and the Task of Science, *Social Studies of Science* 37(2): 297–310.
- Nersessian, N., Kurz-Milcke, E., Newstetter, W.C., and Davies, J. 2004. Research Laboratories as Evolving Distributed Cognitive Systems. In R. Alterman and D. Kirsch (Eds.) *Proceedings of the 25th Annual Conference of the Cognitive Science Society* (July 31st–August 2nd, 2003). Boston: Taylor & Francis, 857–862.
- Nye, M.J. 1980. N-Rays: An Episode in the History and Psychology of Science, *Historical Studies in the Physical Sciences* 11(1): 125–156.
- Pacherie, E. 1995. Do We See With Microscopes, *The Monist* 78(2): 171–188.
- Palermos, S.O. 2014. Loops, Constitution, and Cognitive Extension, *Cognitive Systems Research* 24: 25–41.
- Palermos, S.O. and Pritchard, D. 2013. Extended Knowledge and Social Epistemology, *Social Epistemology Review and Reply Collective* 2(8): 105–120.
- Popper, K. 1972. *Objective Knowledge: An Evolutionary Approach*. Oxford: Clarendon Press.
- Pritchard, D. 2003. Virtue Epistemology and Epistemic Luck, *Metaphilosophy* 34(1–2): 106–130.
- Pritchard, D. 2010. Cognitive Ability and the Extended Cognition Thesis, *Synthese* 175(1): 133–151.
- Ross, A. 1986. Why Do We Believe What We Are Told? *Ratio* 28(1): 69–88.
- Vaesen, K. 2011. Knowledge Without Credit, Exhibit 4: Extended Cognition, *Synthese* 181(3): 515–529.
- Welbourne, M. 1986. *The Community of Knowledge*. Aberdeen: Aberdeen University Press.
- Williamson, T. 2000. *Knowledge and Its Limits*. New York: Oxford University Press.
- Winkler, M.G. and van Helden, A. 1992. Representing the Heavens: Galileo and Visual Astronomy, *Isis* 83(2): 195–217.
- Worrall, J. 2002. What Evidence in Evidence-Based Medicine? *Philosophy of Science* 69(S3): S316–S330.
- Wylie, A. 2011. Critical Distance: Stabilizing Evidential Claims in Archaeology, in P. Dawid, W. Twining, and M. Vasilaki (Eds.) *Evidence, Inference, and Enquiry*, Proceedings of the British Academy, Vol. 171. Oxford: Oxford University Press, 371–394.