epistemic pluralism

edited by
annalisa coliva
nikolaj j. l. l.
pedersen

Palgrave Innovations in Philosophy

Series Editors
Vincent Hendricks
Center for Information and Bubble Studies
University of Copenhagen
Copenhagen, Denmark

Duncan Pritchard
University of Edinburgh
Edinburgh, UK

Palgrave Innovations in Philosophy is a new series of monographs. Each book in the series will constitute the 'new wave' of philosophy, both in terms of its topic and the research profile of the author. The books will be concerned with exciting new research topics of particular contemporary interest, and will include topics at the intersection of Philosophy and other research areas. They will be written by up-and-coming young philosophers who have already established a strong research profile and who are clearly going to be leading researchers of the future. Each monograph in this series will provide an overview of the research area in question while at the same time significantly advancing the debate on this topic and giving the reader a sense of where this debate might be heading next. The books in the series would be of interest to researchers and advanced students within philosophy and its neighboring scientific environments.

More information about this series at
http://www.springer.com/series/14689

Annalisa Coliva
Nikolaj Jang Lee Linding Pedersen
Editors

Epistemic Pluralism

Editors
Annalisa Coliva
Department of Philosophy
University of California Irvine
Irvine, CA, USA

Nikolaj Jang Lee Linding Pedersen
Yonsey International Campus
Underwood International College
Incheon, Korea (Republic of)

Contents

Part III   Epistemic Pluralism and Epistemic Relativism

Part IV   Epistemic Pluralism: Some Applications

Editors and Contributors

About the Editors

Annalisa Coliva is a Professor of philosophy at the University of California at Irvine. She has written widely in epistemology, philosophy of mind and history of analytic philosophy. Her main books are *Moore and Wittgenstein. Scepticism, Certainty and Common Sense* (2010), *Extended Rationality. A Hinge Epistemology* (2015) and *The Varieties of Self-Knowledge*. A new book on *Relativism*, with Maria Baghramian, is forthcoming.

Nikolaj Jang Lee Linding Pedersen is an Associate Professor of philosophy and director of the Veritas Research Center at Underwood International College, Yonsei University. He works on truth, epistemology and metaphysics. He has published in journals such as *Noûs, Philosophical Issues, Philosophical Quarterly, Analysis, Synthese, Erkenntnis* and *The Monist*. He is a co-editor of *New Waves in Truth* (Palgrave Macmillan, 2010), *Truth and Pluralism: Current Debates* (Oxford University Press, 2013), *Epistemic Entitlement* (Oxford University, forthcoming) and *The Routledge Handbook of Social Epistemology* (Routledge, forthcoming).

Contributors

Delia Belleri is a fixed-term Assistant Professor at the University of Vienna. Previously, she held positions at the University of Hamburg, at the University of Barcelona and at UNAM (Mexico City). She is interested in a variety of issues in metaontology, with a special focus on the epistemology and semantics of ontological questions.

J. Adam Carter is a Lecturer in philosophy at the University of Glasgow, working mainly in epistemology. His work has appeared in (among other places) *Noûs*, *Philosophy and Phenomenological Research*, *Analysis*, *Philosophical Studies* and the *Australasian Journal of Philosophy*. Carter's first monograph, *Metaepistemology and Relativism*, was published in 2016 with Palgrave Macmillan.

Pascal Engel is director of studies at Ecole des hautes études en sciences sociales, Paris. He works primarily in epistemology, philosophy of logic and philosophy of mind. His recent interests are on belief and its varieties and epistemic norms. He has written a number of books, including *The Norm of Truth* (Simon and Schuster 1991) and *Ramsey, Truth and success* (Routledge, 2002).

Martin Kusch is a Professor of Philosophy of Science and Epistemology at the University of Vienna. Between 2014 and 2019, he is Principal Investigator of an ERC Advanced Grant, 'The Emergence of Relativism' (#339382). He is currently working on two books relating to this theme.

Robin McKenna is currently a postdoctoral researcher at the University of Vienna, where he works on the ERC-funded 'Emergence of Relativism' Project (Grant No. 339382). He completed his Ph.D. at the University of Edinburgh in 2013. Much of his published work is on topics at the intersections of epistemology, philosophy of language and metaethics such as contextualism and epistemic normativity and has been published in journals such as *Australasian Journal of Philosophy*, *Ethics* and *Synthese*. He is currently working on a book and a series of papers that argue that knowledge is a social phenomenon.

Anne Meylan is a SNSF Professor at the University of Basel. Since 2015, she is the head of the project 'Cognitive Irrationality' funded by the Swiss National Science Foundation. Her research profile lies at the intersection of epistemology and metaethics. Mainly, she is the author of numerous papers and books on the topic of the ethics of belief.

Jennifer Nado is an Assistant Professor at the University of Hong Kong. Her primary research area is metaphilosophy, with an emphasis on the philosophical implications of experimental philosophy.

Erik J. Olsson is a Professor and Chair in Theoretical Philosophy at Lund University, Sweden. His areas of research include epistemology, philosophical logic, pragmatism and, more recently, philosophy of the Internet. Olsson has published extensively on the subjects of epistemic coherence, reliabilism, social epistemology and logic of belief revision. His books include *Against Coherence: Truth, Probability, and Justification* (Oxford University Press, 2005), *Knowledge and Inquiry: Essays on the Pragmatism of Isaac Levi* (Cambridge University Press, 2006) and *Belief Revision Meets Philosophy of Science* (Springer, 2011).

Michele Palmira is a Juan de La Cierva Fellow at the University of Barcelona and a senior member of the research group LOGOS. He works in epistemology and the philosophies of language and mind. He has published articles on disagreement, relativism, moral deference and singular thoughts in journals such as *The Philosophical Quarterly*, *Synthese*, *Pacific Philosophical Quarterly* and *Ratio*.

Introduction

Annalisa Coliva and Nikolaj Jang Lee Linding Pedersen

1   Varieties of Epistemic Pluralism

Epistemic pluralism is an emerging area of research in epistemology with dramatic implications for the discipline. The aim of this edited collection is to draw out some of these implications, articulate and explore different versions of the view, consider its motivations and applications, and investigate its connection to other views in epistemology—in particular, epistemic relativism.

Given a schematic formulation epistemic pluralism takes the following form:

There are several (ways of being) *X*,

A. Coliva (✉)
Department of Philosophy, University of California,
Irvine, CA 92697, USA
e-mail: a.coliva@uci.edu

N.J.L.L. Pedersen
Underwood International College, Yonsei University,
Incheon
406-840, South Korea
e-mail: nikolajpedersen@gmail.com

1

where some epistemic notion or standing takes the place of "*X*". Depending on what is substituted for "*X*" different forms of epistemic pluralism ensue—for example:

(a)   There are several ways of being epistemically justified.
(b)   There are several ways of being epistemically warranted.
(c)   There are several ways of being epistemically rational.
(d)   There are several epistemic desiderata.
(e)   There are several epistemic principles.
(f)   There are several epistemic methods.
(g)   There are several epistemic goods.

Epistemic pluralism in these various guises is already present in the literature. Indeed, several prominent philosophers endorse (at least) one of these varieties of epistemic pluralism.

Alvin Goldman and Tyler Burge are pluralists about epistemic justification. Goldman thinks that there is such a thing as subjective justification and also objective justification—what he calls respectively "weak justification" and "strong justification". Members of epistemic communities that rely on a de facto unreliable method can be subjectively—or weakly—justified despite the unreliability of the relevant method. This is so if, roughly, there is no indication that the method is unreliable from the perspective of the subject and the method is widely relied on and regarded as reliable by members of the community. Objective justification—or strong justification—is reliabilist justification, i.e. it is (roughly) belief formed via a reliable belief-forming method. Goldman is known, of course, for his articulation and defence of reliabilism as a theory of justification, but he wants to make room for subjective—or weak—justification in order to capture mere blameless belief.[1]

Burge buys into internalist justification as well as externalist justification. According to Burge, the internalism/externalism distinction should be understood in terms of reasons. Internalist justification is reason-involving justification where there is an operative or relied-upon reason in the subject's psychology. Externalist justification is justification with no reason operative or relied-upon in the subject's psychology. Someone who believes that there are infinitely many primes on the

basis of a proof have this kind of justification, as the proof constitutes a reason that is operative in the person's psychology. Other justifications are not reason-involving. A young child's perceptual belief that there is a tree in front of her is justified, but there is no reason operative in the child's psychology that justifies the belief. An insistence that internalist justification is the only type of justification would deprive the beliefs of certain individuals—including young children—of enjoying a positive justificatory status. This, according to Burge, would amount to an over-intellectualization of epistemology.[2]

Crispin Wright is a pluralist about epistemic warrant: there is evidential warrant as well as non-evidential warrant—or, in his terminology, justification and entitlement. On his view, ordinary beliefs about the empirical world (e.g. the belief that there are trees) are warranted evidentially while beliefs in basic propositions (e.g. the belief that there is an external world) are warranted non-evidentially. The latter are warranted non-evidentially as a matter of principle, as the best attempts to acquire an evidential warrant to believe basic propositions are epistemically circular.[3]

Alvin Plantinga distinguishes between two kinds of rationality—internal and external rationality. Internal rationality is rationality "downstream from" phenomenal experience, while external rationality is rationality "upstream from" phenomenal experience. Internal rationality tracks whether a given belief is appropriately held in light of—or given—the subject's phenomenal experiences. External rationality, on the other hand, tracks the genesis or history of beliefs—whether the phenomenal experiences on which they are based came about as the result of properly functioning cognitive capacities or belief-forming processes.[4]

William Alston is known for arguing that the debate between advocates of different accounts of justification is futile. He adopts a pluralist attitude towards the debate: alternative accounts of justification need not be regarded as competitors—rather, they all track epistemically good-making features that individual beliefs or systems of beliefs may have. Alston drops the talk of justification and instead uses "epistemic desiderata" to refer to those good-making features. Epistemic desiderata include being true and, relatedly (since truth-conducive), being based

on adequate evidence, being formed via a reliable belief-forming process, and being formed via properly functioning cognitive capacities. Alston also includes three higher-order desiderata: having good cognitive access to evidence for a belief, knowing or having a good reason to believe that a given belief has a positive epistemic status and being able to carry out a successful defence of the probable truth of a belief. These are higher-order desiderata because they involve conceptualizing a belief as possessing a certain epistemic status, or assessing evidence or offering reasons in support of a belief.

The examples of Goldman, Burge, Wright, Plantinga, and Alston cover (a)–(d) above, i.e. pluralism about epistemic justification, warrant, rationality, and desiderata. To give an example of (e), (f), and (g) we can look at the debates on respectively epistemic relativism and the value problem.

According to a well-known brand of epistemic relativism—what we might call "systems-relativism"—justification is relativized to epistemic systems. Beliefs enjoy justification relative to this-or-that epistemic system. This kind of relativism can be found in Rorty and, at least according to some scholars, in Wittgenstein.[5] Systems-relativism involves a commitment to pluralism about epistemic principles or epistemic methods. This is so because epistemic systems are constituted by epistemic principles or epistemic methods, *and* because the systems-relativist additionally endorses the idea that there are several equally correct epistemic systems. A religious belief system might include reliance on the Bible as a justification-conferring method of belief acquisition while this kind of method is not included in scientific belief systems. Relative to the religious belief system the fact that the Bible says that *p* might justify a thinker in believing that *p*. However, relative to a scientific belief system that fact does not justify a thinker in believing that *p*.

Truth is required for knowledge. Additionally, suppose—as is widely done—that truth is the only epistemic good or value. The so-called value problem in epistemology is the problem of accounting for how or why knowledge is more valuable than mere true belief. It is difficult to see how this is possible if truth is the only epistemic good. For, in that case, whatever other components knowledge might have, it would seem that these other components cannot add any value to a belief on

top of the value accrued in virtue the belief's being true. Not wanting to abandon the idea that knowledge is more valuable than mere true belief, some take the value problem to show that truth cannot be the only epistemic good or value. There must be more than one. Michael DePaul is one advocate of this view, pluralism about epistemic goods or value.[6]

Epistemic pluralism is in a certain sense on the map already. As seen above prominent philosophers endorse epistemic pluralism in the various guises identified earlier. However, while epistemic pluralism is already on the philosophical map in this sense, it is still an emerging area of research. For, so far, when present or represented in philosophical work, epistemic pluralism has tended to be left in the background rather than being the main focus. This volume is the first edited collection dedicated exclusively to epistemic pluralism as a topic in its own right.

Much work remains to be done in terms of further investigating the varieties of epistemic pluralism presented above. The following questions are relevant to that task:

Q1.   How might different varieties of epistemic pluralism be articulated in greater detail?
Q2.   How might different brands of epistemic pluralism be connected?
Q3.   What compatibilities and incompatibilities might there exist between different forms of epistemic pluralism and monism?

Regarding Q1, one might think that it is relevant to draw a fundamental/non-fundamental distinction in understanding each of (a)–(g). It would seem relevant in at least some cases. One might think that there are several epistemic goods (say, reliability, evidence, reasons, and justification), but hold that there is only one fundamental good, truth, because those other goods derive their status from the connection they bear to truth. Someone who holds this type of view would count as a pluralist about epistemic goods if the pluralist thesis is understood in an unqualified sense or in terms of non-fundamental epistemic goods. However, if the pluralist thesis is understood in terms of fundamental goods, this type of view qualifies as a form of monism.

Regarding Q2, it may be that some kinds of epistemic pluralism imply other kinds. For example, if epistemic justification, warrant, rationality,

and desiderata qualify as non-fundamental epistemic goods, then it would seem that pluralism about any of these implies pluralism about non-fundamental epistemic goods. Regarding Q3, if epistemic justification, warrant, rationality, and desiderata qualify as non-fundamental epistemic goods, then pluralism about any of these is incompatible with monism about non-fundamental epistemic goods. However, it is worth noting that both pluralism about non-fundamental epistemic goods and pluralism about any of justification, warrant, rationality, and desiderata is compatible with monism about fundamental epistemic goods (and, in particular, with the idea that truth is the only fundamental epistemic good).

Work on epistemic pluralism will serve to reconfigure many debates in epistemology. For traditional projects in epistemology have been carried out under the aegis of an often-implicit assumption of epistemic monism. In particular, many philosophers have either implicitly or explicitly assumed monism about the nature of justification. The field has been dominated, for instance, by a plethora of proposals defending some brand of externalism or internalism. The only consensus reached is that the debate has not been adjudicated either way.

Now, such a result is only to be expected once one realizes that these different ideas about the nature and structure of justification depend on overarching different constraints imposed on the very notion of justification. These different constraints, in turn, depend on the kind of overarching epistemic projects theorists have pursued, varying from accounting for justification and knowledge reached through perception or within scientific enterprises, to confronting various kinds of skeptical challenges. There are reasons to think, however, that these projects, with their different attendant constraints on the notions of justification and knowledge, are equally legitimate.

Hence, exploring pluralism about justification will produce interesting first-order philosophical results, like paying closer attention to the various requisites one might want to impose on the notion of justification (and therefore of knowledge) and why. Moreover, it will produce also interesting second-order or meta-philosophical results. Namely, it will help us see why a long-lasting dispute in the domain

of epistemology about the nature of justification has reached a stalemate and how embracing an overarching different outlook would help go past it. Furthermore, it would help us better appreciate the relationship between the various epistemic projects epistemologists have been pursuing. The same points can be made in relation to other epistemic notions about which one might be a pluralist (rationality, warrant, etc.).

2 Motivations for Epistemic Pluralism

There are several *foci* of disagreement around which the debates on the nature of epistemic justification, warrant, rationality, and goods are articulated. A core debate, as we have seen, is the one between internalism and externalism. Despite its longevity and the amount of effort epistemologists have devoted to identifying the weaknesses and strengths of each position this debate is felt by many as intractable and no agreement or last word seems to be forthcoming.

A way to make some progress in the debate consists in abandoning epistemic monism, i.e. the claim that there is only one way for a belief to be justified (warranted, rational, good, etc.). As a consequence, although internalists and externalists take themselves to be offering incompatible accounts of the same target of investigation, if we reject epistemic monism and embrace pluralism instead, we are in a position to cast new light on this debate, to account for some striking facts that characterize it, and eventually to dissolve it.

The first motivation for pluralism hinges on the role that our intuitive judgments play in the theory of epistemic justification, warrant, rationality, and value. Although all the various theories explored in the debate are subject to counterexamples, each of them seems to perform well as far as it deals with certain cases of epistemic evaluation. This fact is nicely explained if it is acknowledged that there is a plurality of equally legitimate accounts of justification (warrant, rationality, value): internalist accounts will be in a position to capture some of our intuitions, whereas other intuitions will be best captured by externalist accounts. Thus, endorsing epistemic pluralism promises to dissolve

a debate that has so far proved to be intractable, and to account for our intuitions about justification, warrant, rationality, or value.

The second motivation stems from the realization that there appears to be a fundamental disagreement about the significance of the sceptical challenge for a theory of justification (warrant, rationality, or value): while externalists criticize internalism on the ground that it concedes too much to the sceptic, internalists object to externalism that it fails to take scepticism seriously. One might take this puzzling dialectic strongly to suggest that there are different equally legitimate philosophical projects that focus on different equally legitimate notions of justification (warrant, rationality, or value): one which aims to avoid scepticism by taking the sceptical challenge as seriously as possible, and another which downplays the importance of the sceptical challenge and focuses instead on conceptions of epistemic justification (warrant, rationality, or value) that promise to preserve the commonsensical core of our ordinary epistemic practices.

The third motivation, explored by Jennifer Nado (this volume, Chap. 6), connects with the debate over intuitions in philosophical methodology. Reliance on intuitions—e.g. ones concerning Gettier cases—have been criticized on various grounds. The common denominator of these criticisms is that intuitions do not meet a certain epistemic standard, where the relevant standard is associated with some concept (KNOWLEDGE, say), as it is used in philosophy. Nado counters this criticism by arguing that intuitions correlate with the ordinary concept of knowledge, and that this concept goes with lower standards than the standards for the philosophical concept(s) of knowledge. The upshot, according to Nado, is a form of epistemic pluralism: intuitions may generate evidence, justification, or knowledge by ordinary standards but not by philosophical standards—where the latter part of this claim could be expressed by saying that intuitions may fail to generate evidence$_P$, justification$_P$, or knowledge$_P$ Thus, there are several epistemic concepts—ordinary and philosophical ones—and they go with different epistemic states.

The fourth path towards epistemic pluralism is investigated by Erik Olsson (this volume, Chap. 2). He argues that by endorsing Carnap's method of explication we have an independent motivation for epistemic

pluralism. By explication, Carnap means "the transformation of an inexact, prescientific concept, the explicandum, into a new exact concept, the explicatum". The method is open to a pluralist view in two ways: by allowing the existence of more than one explicandum, and by making room for a plurality of equally legitimate ways of explicating the same explicandum.

3   Epistemic Pluralism: Further Articulations, Motivations, and Criticisms

As noted earlier, not much work has been done in order to understand what epistemic pluralism exactly amounts to and in what terms it is most interestingly formulated. This issue is confronted head on by Pedersen (Chap. 3). He discusses how to motivate epistemic pluralism, how it is most interestingly formulated, and how to defend it from two fundamental objections. His discussion draws inspirations from other philosophical debates in which pluralist theories have been recently pursued, especially the truth debate.[7] Pedersen shows that the epistemic pluralism to be found in the works of Tyler Burge, Alvin Goldman, and William Alston is motivated in a way structurally similar to the so-called alethic scope problem—the most widely cited motivation for pluralism about truth. The idea behind the alethic scope problem is that no single traditional account of truth has a scope sufficiently wide to accommodate all truth-apt discourse. However, by appealing to different accounts of truth for different domains the pluralist can accommodate all truth-apt discourse. In a similar fashion Burge, Goldman, and Alston want to accommodate a range of epistemic evaluations that no single traditional account of justification has a scope sufficiently wide to accommodate. The upshot: they adopt a pluralist approach, letting several forms of justification—or desiderata—work together to accommodate the relevant range of epistemic evaluations.[8]

Pedersen also suggests that most existing forms of epistemic pluralism—including those of Burge, Goldman, and Alston—are somewhat moderate in nature. This is because they are paired with the thesis that truth is there is only one non-derivative epistemic

good, truth. The various kinds of epistemic standings that Burge, Goldman, and Alston endorse—different forms of epistemic justification and desiderata—all derive their value from being truth-connected. In this sense, at a very fundamental level, these varieties of epistemic pluralism are very monistic in nature. In the light of this Pedersen proposes that epistemic pluralism is most interestingly—and strongly—formulated at the level of non-derivative epistemic goods, i.e. epistemic goods that do not derive their value from any other epistemic good. He refers to this form of pluralism as "pure epistemic pluralism" and attributes this kind of view to Michael DePaul and Jonathan Kvanvig. Lastly, drawing again on the truth debate, Pedersen formulates two "collapse arguments" against pure epistemic pluralism. These are arguments that purport to show that straightforward reasoning can be deployed to show that pure epistemic pluralism collapses into monism. Pedersen, a sympathizer of pluralism, offers lines of response to both collapse arguments.

Like Pedersen and Nado, Anne Meylan (Chap. 5) and Robin McKenna (Chap. 7) look at epistemic pluralism with sympathetic eyes. Meylan distinguishes between two senses of "justification": ordinary and technical. In order for uses of "justification" to qualify as ordinary they must have what Meylan calls the "trans-categorical feature"—the feature that "justification" is used in such a way that it is applicable to beliefs and actions alike (i.e. it applies across categories). According to reliabilism—the perhaps most well-known form of externalism—a given belief that p is justified in virtue of being formed via a reliable belief-forming process. According to accessibilism—or what might be taken to be a statement of internalism—a given belief that p is justified in virtue of the subject's having cognitive access to facts—possibly other beliefs—that support the truth of p. Meylan argues that "justification" does not have the trans-categorical feature when used in accordance with reliabilism and accessibilism. Accordingly, "justification" on the reliabilist and accessibilist pictures cannot be used in the ordinary sense. Instead "justification" is used in a technical sense. Meylan takes this to amount to a kind of *technical* epistemic pluralism.

McKenna articulates and supports a form of pluralism about knowledge that he calls "standards pluralism". Knowledge is usually taken to be a two-place relation between a subject and a proposition. Standards pluralism is the combination of two ideas: (i) knowledge is a three-place relation obtaining between a subject, proposition, and a standard for knowledge, and (ii) there is a plurality of standards for knowledge. McKenna makes a case for standards pluralism by arguing that it offers a prima facie plausible solution to several challenges. Standards pluralism can explain why we are inclined to assign different truth values to the same knowledge attribution in low- and high-stakes cases. Standards pluralism can likewise explain away the inconsistency between < I know that I have hands > , < If I know I have hands, I know that I'm not a brain-in-a-vat > , and < I don't know that I'm not a brain-in-a-vat > . For each of these propositions, there is no such thing as knowing simpliciter. Rather, a subject knows relative to a standard. So, it may be that the first and second propositions are true and the third false by ordinary standards while, by skeptical standards, the first proposition is false and the second and third true.

Standards pluralism is at least superficially similar to two prominent views: contextualism and relativism. However, McKenna explains why, upon closer scrutiny, standards pluralism is different from both of these views. Standards pluralism is different from relativism. For, while both views involve a relativization of an epistemic standing to some X (standards, systems, etc.), relativism involves the further idea that there is a plurality of equally correct Xs. Standards pluralism, on the other hand, does not involve commitment to the idea that the plurality of Xs cannot be ranked. Standards pluralism is different from contextualism. This is because contextualism is a *linguistic* thesis. It's a thesis about how a parameter—stakes—can vary across contexts and impact the semantic status of knowledge ascriptions. Standards pluralism, on the other hand, is a *metaphysical* thesis. It is a thesis about the nature of the knowledge relation: it is a three-place relation that obtains between a subject, a proposition, and a knowledge standard.

Most contributions to this volume aim to explore, support, or defend versions of epistemic pluralism. Pascal Engel's contribution (Chap. 4), however, is entirely devoted to criticizing the project. He critically

discusses whether a plausible version of the view is available. He considers five candidate views: the ambiguity view, dualism (attributed to Goldman and Sosa), the desiderata approach (Alston), functionalism (an epistemic analogue of Crispin Wright and Michael Lynch's form of truth pluralism), and disjunctivism (Pritchard). According to the first view—the ambiguity view—"justification" is ambiguous between an internalist concept of justification and an externalist one, with these two concepts being regarded as being in competition with one another and being mutually <u>in</u>compatible. According to the second view—dualism—"justification" is ambiguous between an internalist concept of justification and an externalist one, with these two concepts being compatible and both legitimate concepts or correct concepts of justification. According to the third view—the desiderata approach—there is a range of epistemic desiderata. These all mark epistemically good-making features of beliefs or systems of belief, as they are connected to the one epistemic goal: (significant) truth. According to the fourth view—functionalism—the concept of justification is characterized in a functionalist manner. That is, the concept is characterized by a set of principles that pin down its functional role and, through this role, serve to define the concept. The fifth view, disjunctivism, is restricted to perceptual belief. According to this view, the rational support enjoyed by a subject's belief in a good, veridical case of perception is different from the rational support enjoyed by the subject's belief in a bad, non-veridical case of perception. Good and bad cases may be indistinguishable to the subject. However, in the former case the subject has rational support for the belief that is both factive and reflectively accessible—namely, seeing that p. This kind of rational support is absent in bad, non-veridical cases.

Engel dismisses the ambiguity view and disjunctivism as genuine forms of pluralism, and so *they* cannot contribute towards the goal of putting a plausible form of *pluralism* on the market. Against the remaining three views Engel launches the same basic objection: all concepts of justification and all desiderata are related to truth and depend for their value on this relation. Thus, for the views considered Engel takes their *axiological unity* to undermine their candidacy for being genuinely pluralist views.

4   Epistemic Pluralism and Epistemic Relativism

In addition to exploring epistemic pluralism as a theme in its own right, this volume likewise explores potential connections and differences between relativism and pluralism. Martin Kusch (Chap. 8) and Adam J. Carter (Chap. 9) focus on relativism and pluralism, as they appear in the work of Paul Boghossian (2006). Epistemic relativism, according to Boghossian, is the view that results from combining the following three theses (2006, pp. 84–85):

(ER1)   There are no absolute facts about what belief a particular item of information justifies. (*Epistemic non-absolutism*)

(ER2)   If a person, *S*'s, epistemic judgments are to have any prospect of being true, we must not construe his utterances of the form '*E* justifies belief *B*' as expressing the claims *E justifies belief B* but rather as expressing the claim: *According to the epistemic system C (that I, S, accept) E justifies belief B.* (*Epistemic Relationism*)

(ER3)   There are many fundamentally different, genuinely alternative epistemic systems, but no facts by virtue of which one of these systems is more correct than any of the others. (*Epistemic Pluralism*)

Looking at this characterization we see that, as characterized by Boghossian, epistemic relativism carries a commitment to epistemic pluralism in the sense of *including* a form of epistemic pluralism.

Martin Kusch's contribution critically discusses Boghossian's (2006) objections to epistemic pluralism, that is, in this context, the view that there can be a plurality of equally legitimate epistemic systems, where two epistemic systems are different if they yield "*conflicting* verdicts on what it would be justified to believe under specified evidential conditions". Kush argues that Boghossian's objections are ineffective, and that epistemic pluralism so understood is a tenable position.

Carter explores how one might formulate a stable version of epistemic pluralism, and whether it is possible to formulate a non-relativism involving form of epistemic pluralism. Carter takes the first part of epistemic pluralism—the presence of genuinely alternative epistemic systems—to be uncontroversial. It is a mere expression of what most people take to be a fact: epistemic diversity. The second part, according to Carter, is much more controversial—i.e. the idea that these alternative systems are somehow on a par, there being no fact that makes one system more correct than any other. He argues that this idea can be accommodated within the framework of Wittgensteinian hinge epistemology because, within this framework, rational evaluation is essentially local. Carter likewise argues that the combination of epistemic pluralism and Wittgensteinian hinge epistemology is relativistic in nature and requires a commitment to epistemic incommensurability, i.e. the thesis that two subjects can have equally rationally justified beliefs without there being a rational basis for one subject to persuade the other to revise her view. Carter further explores hinge-epistemological proposals put forward by Duncan Pritchard. He ends by suggesting that a hinge epistemologist may be able to take on board epistemic pluralism and epistemic incommensurability but *not* commit to relativism, provided that relativism is understood along the lines of John MacFarlane's assessment-sensitive relativism.

5   Epistemic Pluralism: Some Applications

The last section of the book applies epistemic pluralism to different current debates in epistemology and metaphysics, to reconfigure their boundaries and to make room for much more compatibility between different extant positions than it is usually admitted under the auspices of monism.

In "How to be a pluralist about self-knowledge" (Chap. 10), Annalisa Coliva presents a pluralist account of self-knowledge—that is, knowledge of our own mental states—and connects it to discussions regarding epistemic pluralism in general. The first *existential* thesis, at the heart of this kind of pluralism, is that there is an *asymmetry*

between first- and third-personal cases of self-knowledge. The second *meta-epistemological* thesis is that a complete account of self-knowledge needs to address *both* first- *and* third-personal cases. The third thesis is an *axiological* one: both kinds of self-knowledge are equally epistemically interesting and existentially important. The fourth thesis is that pluralism about self-knowledge is philosophically committal: even though it makes room for the compatibility of different accounts of self-knowledge by reconfiguring their proper boundaries, it does not entail that all extant accounts of self-knowledge are in good standing. In particular, the proposed account endorses constitutivism with respect to first-personal self-knowledge and defends it from the *prima facie* devastating objection of rendering it a misnomer to call that kind of self-*knowledge* thus. Moreover, it embraces methodological pluralism with respect to third-personal self-knowledge and draws attention to the epistemologically different, so far largely unexplored ways in which we can gain it.

In "How to be a pluralist about disagreement" (Chap. 11), Michele Palmira assesses current debates on doxastic disagreement. He examines four different ways that doxastic disagreement can present itself: descriptive disagreement, conceptual disagreement, full disagreement, and credal disagreement. He argues that pluralism is one way to resolve issues concerning doxastic disagreement. One such pluralist account, developed out of John MacFarlane's work, is *disjunctive pluralism*. Palmira criticizes disjunctive pluralism and argues for an alternative pluralist theory of disagreement that he dubs *kinship pluralism*. He argues that kinship pluralism can be adequately extended to other varieties of disagreement, namely, *group* and *agnostic* disagreement.

Finally, Delia Belleri in "A pluralistic way out of epistemic deflationsim about ontological disputes" (Chap. 12) applies epistemic pluralism to ontological disputes. She counters epistemicism about the debate on material composition, according to which there are too little grounds to believe any of the competing theories—Nihilism, Universalism, and intermediate positions—with a form of epistemic pluralism that may be compatible with epistemic relativism. In particular, she argues that each party to the debate is justified

relative to the ranking of theoretical features and virtues that is "internal" to each position. This position, it is argued, is preferable for its greater charity towards the participants involved in the relevant ontological debate.

Notes

1. For Goldman on weak and strong justification, see his (1988). His comprehensive body of work on realiabilism include Goldman (1979) and (1986).
2. Burge (1993), (1998), (2003). A terminological remark: Burge himself refers to internalist justification as "justification" and externalist justification as "entitlement". This terminological divergence from Burge does not matter for present purposes. It should also be noted that there are significant similarities between Burge's distinction between two kinds of justification and Ernest Sosa's distinction between animal knowledge and reflective knowledge (see (2009a, b)). For this reason we count Sosa as an epistemic pluralist too.
3. Wright (2004). A terminological note: both Burge and Wright draws a distinction between justification and entitlement and take each of these two be species of the same genus, *warrant*. Although their terminology is the same, Burge and Wright's distinctions are orthogonal. For, as seen, the Burgean distinction is between internalist and externalist justification (or warrant) while the Wrightian distinction is between evidential and non-evidential warrant. Note also that, strictly speaking, Wright thinks that the attitude that goes with non-evidential warrant is trust—a species of acceptance distinct from belief. This does not matter for present purposes.
4. Plantinga (2000, pp. 110–112).
5. Rorty (1981), Wittgenstein (1969). See Boghossian (2006) and Kusch (forthcoming) for respectively criticisms and defences of epistemic relativism. For a critique of the idea that Wittgenstein was an epistemic relativist—a thesis that both Boghossian and Kusch endorse—see Coliva (2010).
6. DePaul (2001). What we refer to as the "value problem" is what Pritchard (2007) calls the "primary value problem". The secondary value problem is the challenge of explaining why knowledge is more valuable than any

proper subset of its parts (e.g. justified true belief, assuming that there is also an anti-Gettier condition). The tertiary value problem is the challenge of explaining why knowledge possesses a distinctive kind of value.

7. Pluralist approaches have proved to be fruitful ways of looking with new eyes at old issues in several philosophical debates. For alethic pluralism, see Lynch (2009) and Wright (2013); for logical pluralism, see Beall and Restall (2006); for ontological pluralism, see McDaniel (2009) and Turner (2010); for causal pluralism, see Godfrey-Smith (2009).

8. As remarked earlier, strictly speaking, Burge speaks in terms of warrant—reserving the expression "justification" for internalist or reason-involving warrant.

References

Beall, J.C., and Greg Restall. 2006. *Logical Pluralism*. Oxford: Oxford University Press.

Boghossian, Paul. 2006. *Fear of Knowledge*. Oxford: Oxford University Press.

Bonjour, Laurence. 2003. A Version of Internalist Foundationalism. In *Epistemic: Internalism Vs. Externalism, Foundations Vs. Virtues*, ed. Laurence Bonjour and Ernest Sosa. Malden: Blackwell Publishing.

Bonjour, Laurence. 2010. Recent Work on the Internalism-Externalism Controversy. In *A Companion to Epistemology*, ed. Jonathan Dancy, Ernest Sosa, and Matthias Steup, 33–43. Malden: Blackwell Publishing.

Burge, Tyler. 1993. Content Preservation. *Philosophical Review* 102 (4): 457–488.

———. 1998. Computer Proof, A Priori Knowledge, and Other Minds. *Philosophical Perspectives* 12: 1–37.

———. 2003. Perceptual Entitlement. *Philosophy and Phenomenological Research* 67 (3): 503–548.

Coliva, Annalisa. 2010. Was Wittgenstein an Epistemic Relativist? *Philosophical Investigations* 33 (1): 1–23.

DePaul, Michael. 2001. Value Monism in Epistemology. In *Knowledge, Truth, and Duty: Essays on Epistemic, Responsibility, and Virtue*, ed. Matthias Steup, 170–183. Oxford: Oxford University Press.

Foley, Richard. 2002. Conceptual Diversity in Epistemology. In *The Oxford Handbook of Epistemology*, ed. Paul K. Moser, 177–203. Oxford: Oxford University Press.

Fumerton, Richard A. 1995. *Metaepistemology and Skepticism*. Lanham: Rowman & Littlefield.

Gerken, Mikkel. 2013. Internalism and Externalism in the Epistemology of Testimony. *Philosophy and Phenomenological Research* 87 (3): 532–557.

Godfrey-Smith, Peter. 2009. Causal pluralism. In *The Oxford Handbook of Causation*, ed. Helen Beebee, Peter Menzies, and Christopher Hitchcock, 326–337. Oxford: Oxford University Press.

Goldman, Alvin I. 1979. What Is Justified Belief? In *Justification and Knowledge*, ed. George Pappas, 1–25. Dordrecht: D. Reidel Publishing Company.

———. 1986. *Epistemology and Cognition*. Cambridge: Harvard University Press.

———. 1988. Strong and Weak Justification. *Philosophical Perspectives* 2: 51–69.

Kelp, Christoph, and Nikolaj J.L.L. Pedersen. 2010. Second-Order Knowledge. In *The Routledge Companion to Epistemology*, ed. Sven Bernecker, and Duncan Pritchard, 586–596. New York: Routledge.

Kusch, Martin. Forthcoming. Epistemic Relativism, Scepticism, Pluralism. Forthcoming. *Synthese*. doi:10.1007/s11229-016-1041-0.

Lammenranta, Markus. 2004. Theories of Justification. In *Handbook of Epistemology*, eds. Ilkka Niiniluoto, Matti Sintonen, and Jan Woleński, 467–497. Dordrecht: Kluwer Academic Publishers.

Lynch, Michael P. 2009. *Truth as One and Many*. Oxford: Clarendon Press.

McDaniel, K. 2009. Ways of being. In *Metametaphysics: New Essays on the Foundations of Ontology*, eds. David Chalmers, David Manley, and Ryan Wasserman, 290–319. Oxford: Oxford University Press.

Plantinga, Alvin. 2000. *Warranted Christian Belief*. Oxford: Oxford University Press.

Pritchard, Duncan. 2007. Recent Work on Epistemic Value. *American Philosophical Quarterly* 44 (2): 85–110.

Rorty, Richard. 1981. *Philosophy and the Mirror of Nature*. Princeton: Princeton University Press.

Sosa, Ernest. 2009a. *A Virtue Epistemology: Apt Belief and Reflective*, vol. I. Oxford: Oxford University Press.

———. 2009b. *Reflective Knowledge: Apt Belief and Reflective Knowledge*, vol. II. Oxford: Oxford University Press.

Swinburne, Richard. 2001. *Epistemic Justification*. Oxford: Oxford University Press.

Turner, Jason. 2010. Ontological Pluralism. *Journal of Philosophy* 107 (1): 5–34.

Wittgenstein, Ludwig. 1969. *On Certainty*, ed. G.E.M. Anscombe and G.H. von Wright. Malden: Blackwell Publishing.

Wright, Crispin. 2004. Warrant for Nothing (and Foundations for Free)? *Aristotelian Society Supplementary Volume* 78 (1): 167–212.

———. 2013. A Plurality of Pluralisms. In *Truth and Pluralism: Current Debates*, ed. Nikolaj J.L.L. Pedersen and Cory Wright. Oxford: Oxford University Press.

Part I

Epistemic Pluralism: Methodological Issues

Explicationist Epistemology and Epistemic Pluralism

Erik J. Olsson

1   Introduction

By explication, Carnap explained, "we mean the transformation of an inexact, prescientific concept, the explicandum, into a new exact concept, the explicatum" (1950, p. 3), whereby the latter should satisfy the condition of being reasonably faithful to ordinary use, fruitful, exact, and simple. In the present article, I focus on the application of the method of explication in epistemology. I distill three senses in which explicationist epistemology is intrinsically pluralistic. It allows for there being a plurality of legitimate epistemological project corresponding to pursuing explication of different explicanda signified by the same ordinary term (e.g., "knowledge"); a plurality of reasonable and useful explicata corresponding to the same explicandum; and finally, a plurality of epistemological sub-methodologies for use in parts of the explication

E.J. Olsson (✉)
Department of Philosophy, Lund University, Lund, Sweden
e-mail: erik_j.olsson@fil.lu.se

process. Concerning the last point, I propose in a preliminary fashion that putative rivals such as intuition-based experimental and formal epistemology can be usefully viewed as representing complementary sub-disciplines within a general explicationist framework.

My first task, in Sect. 2, will be to explain the concept of explication, drawing on Carnap's 1950 account.[1] After that, I turn, in Sect. 3, to the question what it means to pursue epistemology in Carnap's sprit, and how this approach is different from other methodological proposals. I investigate the pluralist consequences of the present methodology in Sect. 4. Finally, I briefly comment on the relationship between the present methodology and the related but, as it turns out, not identical pluralistic epistemological picture advanced in William Alston's well-known article on what he calls epistemic desiderata (Alston 1993).

2   Carnap and the Method of Explication

The method of explication was introduced by Rudolf Carnap, most systematically in Carnap (1950), as a procedure for defining scientific concepts generally and philosophical concepts in particular. By explication, Carnap explains, "we mean the transformation of an inexact, prescientific concept, the *explicandum*, into a new exact concept, the *explicatum*" (1950, p. 3). Thus an explication can be seen as a function or mapping from an informal domain to a formal, exact domain. An important aspect of this fact is that there is no exact answer to the question whether a given explication is right or wrong. What we can meaningfully ask is whether it is fruitful, useful, simple, and so on.

The method proceeds in two steps: (1) the elucidation of the explicandum and (2) the specification (precise definition) of the explicatum. As for the first step, "[a]lthough the explicandum cannot be given in exact terms, it should be made as clear as possible by informal explanations and examples" (Carnap 1950, p. 3). Clarifying the explicandum serves the purposes of specifying, perhaps in relatively crude terms, what is to be included and what is to be excluded. Once this task has been accomplished we can meaningfully discuss possible explications of the concept or idea in question with the aim of finding a more exact

concept that is to replace the explicandum in certain specialized contexts. This is similar to how a scalpel is a more precise and useful a tool than the pocket knife in the operating room.

One of Carnap's examples of his method at work concerns the term "true." Suppose that we want to find a suitable explication of this term or concept. An elucidation of the explicandum would in this case involve stating that we do not intend the meaning "true" has in phrases like "a true democracy" or "a true friend," but rather the meaning it has in phrases like "this sentence is true," "what he just said is true," and so on. This does not yet mean that we have explicated the term "true"; we have only zoomed in on one particular meaning of the term. Explicating the concept of truth would involve specifying a formal or semi-formal theory of truth, e.g., in logical or set-theoretical terms, perhaps drawing on the work of Tarski and others.[2]

What requirements should be placed on a suitable explicatum once the explicandum has been sufficiently elucidated? The purpose of explication is to introduce a more or less vague or unclear intuitive concept into an exact framework. Thus, we wish to find a more exact concept that in some sense corresponds to the intuitive, everyday concept and that can do the job which the latter cannot do, or cannot do as well. But what does this relation of correspondence entail? It is obvious that we cannot hope for complete correspondence in meaning. The whole point of explication is, as it were, to diverge from the meaning of the intuitive concept by introducing a more exact correlate to the latter.

A natural proposal is that the explicatum should, nonetheless, be as close or similar to the explicandum as the latter's vagueness permits. Carnap finds, however, that this suggestion is undermined by scientific practice. One of his paradigm examples concerns the artificial scientific concept fish, defined in *Encyclopedia Britannica* as "any of more than 30,000 species of cold-blooded vertebrate animals (phylum Chordata) found in the fresh and salt waters of the world" and for which Carnap uses the Latin term *piscis*. This concept has come to replace the everyday concept of fish in scientific contexts. In fact, it has replaced fish even in everyday contexts. This happened, moreover, in spite of the fact that *piscis* is a narrower concept that excludes several kinds of animal that were previously subsumed under the concept fish, e.g., whales and seals.

Zoologists simply found that *piscis* is a more fruitful concept than fish. In general, Carnap explains, "[a] scientific concept is the more fruitful the more it can be brought into connection with other concepts on the basis of observed facts; in other words, the more it can be used for the formulation of laws." (1950, p. 6) Thus, *piscis*, unlike whales and seals, has a streamlined body for rapid swimming, extracts oxygen from water using gills or uses an accessory breathing organ to breathe atmospheric oxygen, has two sets of paired fins, usually one or two (rarely three) dorsal fins, an anal fin, and a tail fin, has jaws, has skin that is usually covered with scales, and lay eggs. Whales, by contrast, are taken to belong to the category of mammals because, like other mammals, they breathe air, are warm-blooded, nurse their young with milk from mammary glands, and have body hair.

By extrapolation from examples such as those discussed above, Carnap arrives at four general requirements on a suitable explicatum (1950, p. 7):

1. The explicatum [the thing that explicates] is to be *similar to the explicandum* [the thing that is explicated] in such a way that, in most cases in which the explicandum has so far been used, the explicatum can be used; however, close similarity is not required, and considerable differences are permitted.
2. The characterization of the explicatum, that is, the rules of its use (for instance, in the form of a definition), is to be given in an *exact* form, so as to introduce the explicatum into a well-connected system of scientific concepts.
3. The explicatum is to be a *fruitful* concept, that is, useful for the formulation of many universal statements (empirical laws in the case of a nonlogical concept and logical theorems in the case of a logical concept).
4. The explicatum should be as *simple* as possible; this means as simple as the more important requirements 1, 2 and 3 permit.

Of these requirements, the fourth—concerning simplicity—is the least important, Carnap thinks: "In general, simplicity comes into

consideration only in cases where there is a question of choice among several concepts which achieve about the same and seem to be equally fruitful: if these concepts show a marked difference in the degree of simplicity, the scientist will, as a rule, prefer the simplest of them." (1950, p. 7)

Details of Carnap's account can be questioned from a contemporary perspective. One example is the division between "logical" and "nonlogical" concepts, whereby, crucially, "nonlogical" is treated as synonymous with "empirical." This picture leaves out, among other things, the important categories of ethical and legal concepts, not to mention presumably evaluative epistemological expressions such as "good reason" or "justification." Behind Carnap's treatment lies presumably the thought that explication only concerns scientific concepts, and the further view, widely shared at the time, that ethical or legal concepts, to the extent that they cannot be reduced to logical or empirical concepts, do not belong to science, properly so-called.[3]

3   Explicationist Epistemology

From an epistemological perspective, Carnap's original account has the interesting consequence that evaluative epistemological terms like "justification," too, are either not explicable at all or they are explicable in either logical or empirical terms. If we rule out the former, for good reason it would seem, what remains is the latter. Thus there is an affinity between Carnap's method of explication, in its original formulation, and "naturalized epistemology" as championed by Quine (1969) and later by, among others, Goldman (1986) and Kornblith (2002).

It should come as no surprise that Quine was not only a "naturalist" regarding epistemology but also an advocate of Carnap's explicationist methodology. These two strands of Quine's thinking and their intimate interrelations become evident upon reading the final chapter of *Word and Object* (Quine 1960). Quine writes, in the closing paragraph:

> The philosopher's task differs from the others', then, in detail; but in no such drastic way as those suppose who imagine for the philosopher

a vantage point outside the conceptual scheme that he takes in charge. There is no such cosmic exile. He cannot study and revise the fundamental conceptual scheme of science and common sense without having some conceptual scheme, whether the same or another one less in need of philosophical scrutiny, in which to work. He can scrutinize and improve the system from within, appealing to coherence and simplicity; but this is the theoretician's method generally. He has recourse to semantic ascent, but so has the scientist. And if the theoretical scientist in his remote way is bound to save the eventual connections with non-verbal stimulation, the philosopher in his remoter way is bound to save them too. (Quine 1960, pp. 275–276)

If scientific definitions are explications, it follows that philosophical definitions are, or should be seen as, explications as well (cf. Quine 1960, pp. 257–262). In particular, definitions of knowledge, justification and other concepts of epistemological interest are best viewed as explications.

Even so, it is quite possible, in my view, to embrace the methodology of explication in philosophy and epistemology without being also committed to naturalized epistemology. A modern reader can happily dismiss the identification of the nonlogical with the empirical alluded to above as a relic of twentieth century positivism and proceed on the assumption that irreducibly epistemological, ethical, and legal concepts, if such there be and presumably there are, are just as amenable to explication as logical and empirical concepts are once the criteria of fruitfulness are correspondingly broadened.

I will use the term *explicationist philosophy* to refer to this view on the nature of philosophical definitions, in so far as they aim to explicate a concept occurring in natural language. We will henceforth refer to an epistemology based on this methodological position, correspondingly, as *explicationist epistemology*.[4]

Few epistemologists subscribe explicitly to Carnap's methodology. An exception is Keith Lehrer in his book *Theory of Knowledge* (1990). However, Lehrer introduces a subtle modification of Carnap's method, writing that "explication aims at producing concepts useful for articulating laws *and theories*" (1990, p. 6, our italics). Carnap, we recall,

explicitly writes that conceptual fruitfulness is a matter of being useful for the articulation of laws period. Moreover, the outcome of Lehrer's study is surprising given the proclamation that Carnap's method has been used. Thus, Leher's final definition of knowledge, arrived at on p. 147, is neither exact nor simple, and—unsurprisingly given his unorthodox interpretation of Carnap on this point—considerations of fruitfulness, in Carnap's strict sense, do not enter visibly in Lehrer's motivation of his definition. As for simplicity, the specification of Lehrer's explicatum requires no less than 13 definitorial clauses (ibid., pp. 147–149), prompting his remark that "[n]eedless to say, the attempt to analyze justification and undefeated justification in terms of acceptance, reasonableness, and truth has yielded a complicated analysis" (1990, p. 149).[5], although Lehrer thinks there is "underlying simplicity" (ibid.).

Alvin Goldman in his book *Knowledge in a Social World* (1999) is an example of an epistemologist who advances something very similar to explicationist epistemology, yet without reflecting on the relationship to Carnap. Early in the book (1999, p. 5), Goldman declares: "[v]eritistic epistemology (whether individual or social) is concerned with the production of knowledge, where knowledge is here understood in the 'weak' sense of *true belief*." Letting "S-knowledge" and "W-knowledge" stand strong and weak knowledge, respectively, Goldman writes a few pages later (1999, p. 24):

The present book, however, will have nothing to say about S-knowledge [i.e., knowledge as justified, true belief]. It is devoted entirely to the prospects for W-knowledge, which is simply *true belief*. One reason I focus on W-knowledge is to circumvent the intricate issues that surround the notion of S-knowledge. Addressing those issues would demand a major digression from the main thrust of the book. A second and more important reason is that people's dominant epistemic goal, I think, is to obtain true belief, plain, and simple. They want to be *informed* (have true belief) rather than *misinformed* or *uninformed*. The usual route to true belief, of course, is to obtain some kind of evidence that points to the true proposition and away from rivals. But the rationale for getting such evidence is to get true belief. Hence, the entire focus of this book is on W-knowledge.

Goldman proceeds to discuss the extent to which he has "invented" W-knowledge (p. 24) and indicates that he believes that he has not. There is, he thinks, an ordinary sense of "know" that corresponds to true belief. Yet he adds (Goldman 1999, p. 25): "If I am wrong about this, however, I am prepared to proceed cheerfully with weak 'knowledge' as a term of art (or technical term)." In the context of voter core knowledge (1999, p. 324), i.e., knowledge of the core question facing the voter, Goldman is explicit about the usefulness of the weak concept: "I shall show that certain significant consequences for democracy logically follow from widespread core knowledge even in the weak sense of 'knowledge', so we need not concern ourselves with core knowledge in the strong sense."

Thus, Goldman, in attending to "weak knowledge," seems to do something very similar to explicating one purported sense of "knowledge." First of all, his account is purpose-driven; he is interested in finding a fruitful account of knowledge mainly for use in social epistemology. Second, his account contains an "elucidation of the explicandum," i.e., the identification of a promising purported sense of knowledge in which the latter reduces to getting things right. He is then suggesting that this putative sense of knowledge is appropriately explicated as "true belief." In doing so, Goldman can be understood as being willing to sacrifice some similarity with ordinary use in order to gain the advantages of simplicity and usefulness. To the extent that this account is a correct reconstruction of Goldman's intentions, he is in his 1999 book essentially practicing explicationist epistemology. This becomes even clearer as Goldman proceeds to define the value of true belief, or "veritistic value," in exact terms utilizing the resources of probability theory.

The methodology of explication should be understood to imply that all four requirements on an explicatum be given positive weight. After all, Carnap himself refers to the conditions on an explication as "requirements." In particular, the requirement of exactness should be taken quite seriously. An explication is a mapping from an inexact to an exact conceptual domain, although Carnap acknowledged that exactness is not a matter of all or nothing, and that the precification of scientific concepts sometimes proceeds in stages. In practice, the explicatum need not be exact in the logical or mathematical sense, but it should at least be *more exact* than the explicandum.

One way to see more clearly what explicationist epistemology involves is to contrast it with other methodological approaches to epistemology. I will argue that other approaches can be thought of as "limiting cases" of explication in the sense that not all conditions on an explication are thought to be important. Thus interpreted, these other methodological accounts are explications only in a degenerative or improper sense.

If we assign zero weight to all requirements other than that of similarity to the explicandum, the result is an account close to that associated with ordinary language philosophers like J.L. Austin, Gilbert Ryle, and the later Wittgenstein. On this view, the main purpose of an account of a philosophical concept is to shed light, perhaps by means of examples, on the use of that concept or term in ordinary language. Few analytic philosophers nowadays officially subscribe to the ordinary language philosophy which dominated the Oxford scene in the mid-twentieth century. However, it is not unusual to find analytic philosophers who consider vagueness preservation to be a virtue of a philosophical analysis or definition. Faced with the objection that their preferred analysis of a given philosophical concept does not give a clear verdict as to how a particular example is to be classified, they will be inclined to argue, if possible, that this observation actually supports their analysis because the very same vagueness pertains to the explicandum.[6] A desire to preserve the vagueness of ordinary discourse is, for reasons noted, deeply at odds with the methodology of explication.

The ordinary language school has to some extent been revived recently in the form of experimental epistemology. Where the ordinary language philosophers sought to capture the meaning of epistemologically central terms through armchair reasoning and thought-experiments conducted by subjects perceived to be expert users (the philosophers themselves), the experimentalists are preoccupied with probing the intuitions of ordinary people in psychological experiments, e.g., by presenting philosophical laymen with questionnaires. The sole focus in both cases, however, is on the first Carnapian requirement: similarity to the explicandum. Where the two camps differ is in the methodology thought to be most appropriate for this common purpose.

On a traditional account, philosophical methodology amounts to conceptual analysis. A prominent advocate was G.E. Moore, who thought that "[a] thing becomes intelligible first when it is analysed into its constituent concepts" (1899/1970, p. 97). A prototypical example would be the conceptual decomposition of "bachelor" into "unmarried" and "man." Similarly, knowledge has been thought to be analyzable as "justified, true beliefs" and the like. To the extent that the advocate suggests that exploring meaning relations between terms in ordinary language is the sole aim of the philosopher's definitional activity, conceptual analysis is not compatible with the principles of explicationist epistemology but rather constitutes a variation on the theme of ordinary language epistemology.

The point just made is related to the *paradox of analysis*. Consider a proposed conceptual analysis of the form "A is C," e.g., "knowledge is justified, true belief," where A is the analyzandum (what is analyzed) and C the analyzans (what is offered as the analysis). Then either A or C has the same meaning, in which case the analysis is correct but expresses a trivial identity and hence is uninformative; or else A and C do not have the same meaning, in which the analysis is informative but incorrect. Hence, no conceptual analysis can be both correct and informative.

No corresponding puzzle arises for explicationist epistemology. An explication does not aim to deliver a conceptual analysis of the explicandum, in the sense of identifying the meaning constituents of the latter (if such there be). Rather, the aim of an explication is to identify, for a particular purpose, a more exact correlate of the explicandum such that the former satisfies the requirements of similarity to the latter while being in addition fruitful and simple.

A further striking advantage of Carnap's methodology over mainstream epistemological thinking in the conceptual analysis tradition is that the former, unlike the latter, is immune to the Gettier problem. Under what conditions would the Gettier problem be a threat to the claim that a given account of knowledge satisfies the first Carnapian desideratum, that of faithfulness to ordinary use? It would be if it would show that it is not true that, in most cases in which the ordinary concept of knowledge has so far been used, the proposed explicatum can be used in its stead. But there are good reasons to think that the Getter

problem does not show this: Gettier cases are not frequent enough to threaten the claim that a given account of knowledge can be substituted for ordinary knowledge in most cases. This basis for this contention is that Gettier cases involve the consecutive occurrence of two improbable events: a proposition (Brown owns a Ford) that is strongly supported by evidence turns out nonetheless to be false and yet by sheer luck (Brown is in Barcelona) the target proposition (Brown owns a Ford or is in Barcelona) comes out true anyway. Hence, the Gettier problem can never seriously undermine an explication (as opposed to a conceptual analysis) of knowledge.[7]

The methodologies contrasted with explicationalism so far can be reconstructed as being exclusively concerned with shedding light on the ordinary use of epistemologically central terms. These methodologies are limiting cases of explicationism in the sense that they assign all weight to the requirement that the outcome should be maximally similar to the ordinary language concept. At the other end of the spectrum, we find epistemologists who seem largely unoccupied with considerations of ordinary use. Jakko Hintikka's epistemic logic, as advanced in his famous book *Knowledge and belief* (1962) is a case in point. As Hendricks (2006, Chap. 6) explains, epistemic logic in Hintikka's sense was greatly influenced by the advances in modal logic. Specifically, standard systems of modal logic were given epistemic interpretations, and some main technical results of epistemic logic could then be extracted. The point of departure here was not an interest in the ordinary concept of knowledge as such, but a sense that something reminiscent of it could be formalized in the apparatus of modal logic. In standard modal logic, the necessity operator is interpreted as "being true in all possible worlds." More precisely, a proposition A is true in a world w if and only if A is true in all possible worlds accessible from w. In epistemic logic, a similar operator is introduced with the interpretation that A is true in all possible worlds compatible with what the agent knows. To the extent that epistemic logic is not interested at all in ordinary use, or interested only in a very qualified sense, it is merely a limiting case of explicationist epistemology and not a *bona fide* example thereof. The hard core epistemic logician is concerned with maximizing exactness, simplicity, and fruitfulness (in the formal sense that many interesting theorems can be derived).

Arguably, however, most practitioners of epistemic logic or formal epistemology, more generally, are interested in securing some relationship between the outcome of their activities and concepts in ordinary language. Whether a given practitioner of formal epistemology is pursuing explicationism depends on the extent to which he or she is giving sufficient emphasis on the goal of preserving the similarity to the explicandum in relation to Carnap's first requirement.

4   Explicationist Epistemology and Pluralism

As the reader can probably anticipate, explicationist epistemology represents a pluralist approach to epistemology. In fact, there are several senses in which this is true. First of all, explicationism allows for a *plurality of epistemological projects*. The initial step of elucidating the explicandum involves selecting the pre-systematic concept to be explicated among several such concepts. Following Bach (1985) and many others,[8] it is not unreasonable to think that there are two pre-systematic conceptions of epistemic justification—one externalist and the other internalist. If this is true, then it is completely legitimate from the present perspective to focus on one of these conceptions without also spending time and energy on the other. Similarly, if Goldman is right in thinking that there is a strong and a weak sense of knowledge, it would be rationally permissible to devote attention to one of them in the context of a particular epistemological investigation. There are many epistemological projects that can be fruitful and enlightening, depending on which pre-systematic concept (possible explicata) we decide to explicate.

A natural resource to turn to when explicating a given epistemological term is the map of meanings listed in an authoritative dictionary based on lexicology, which is the relevant empirical discipline for identifying word meanings in a systematic fashion. Such a map can be viewed as a list of possible explicanda, and the definitions given as "elucidations" of those explicanda. For example, Oxford Living Dictionaries (Oxford University Press) lists a number of "definitions of knowledge in English" that are useful to contemplate in this connection.[9] As a starting point is noted that the term can be used as a (mass) noun. As such, it

can refer to "[f]acts, information, and skills acquired through experience or education; the theoretical or practical understanding of a subject." Examples of this use would be "a thirst for knowledge" or "her considerable knowledge of antiquities." The term can also denote "the sum of what is known," as in the expression "the transmission of knowledge." A further sense is knowledge as "information held on a computer system." An example of this use is: "The server now has sufficient knowledge to honor a data transfer from the client." The dictionary also lists, under the same mass term heading, a particular use explicitly noted as belonging to philosophy: "True, justified belief; certain understanding, as opposed to opinion." One cited example is: "As a rationalist, he believed that the only path to true knowledge was through logic." According to the same source, knowledge can signify "[a]wareness or familiarity gained by experience of a fact or situation," as in "the programme had been developed without his knowledge" or "he denied all knowledge of the incidents."

Now any of these senses of "knowledge" could be the starting point of a potentially interesting epistemological project culminating in a more precise and (for a particular purpose) useful explicatum that also respects the goals of faithfulness to ordinary use and simplicity. If, for example, our aim is to explicate knowledge in the sense of "the theoretical or practical understanding of a subject," the resulting explicatum would naturally involve an account of the structure of a person's belief system and of how its elements fit together. The net result may very well be a coherentist account of systems of (true) beliefs tied together through inferential or explanatory connections. By contrast, knowledge as "the sum of what is known" would point more obviously in the direction of the "output aspect" of knowledge, i.e., something like Goldman's account in terms of "true belief." The computer sense of knowledge naturally suggests a Hinikka-style account in terms of excluded possibilities. The input aspect of belief, how the belief was obtained, would be salient if the project were one of explicating knowledge in the sense of "awareness or familiarity gained by experience of a fact or situation." An approach that takes into account the cognitive process and its relevant features would be a natural point of departure.

To continue the illustration, once a definite pre-systematic conception of an epistemological term has been selected for further explication, there is room for a *plurality of explicata*. The crucial observation here is that the goodness of an explication is a matter of satisfying all four desiderata, as a package, to as high a degree as possible. The desiderata may, however, work in opposite directions. For instance, a more fruitful explicatum may be less consonant with ordinary use (as the previous whale example illustrates in a scientific context). Carnap does not give any rule for how to weigh the different considerations against each other in cases where there are different plausible ways of explicating the same explicandum. His only advice, as we saw, is that simplicity should generally be the least important concern. It would be just as rational, from a Carnapian perspective, to favor fruitfulness over faithfulness to ordinary use, as it would be to entertain the opposite preference. This said, it is plausible to think that the relative weight of the desiderata will be, to some extent, guided by the context. Thus, one context may require a very exact explication of knowledge, e.g., for the purpose of AI programming. In another context, a relative unpolished explicatum may be sufficient if it proves to be empirically or otherwise fruitful.

Thus, even if we agree to attend to knowledge in the sense of "the theoretical or practical understanding of a subject" there are many ways in which this idea could be made more precise along Carnapian lines, even if we limit our attention to variants of coherentism. How we choose to proceed depends considerably on the purpose of the investigation. If our goal is a general philosophical one that focuses on the broader philosophical understanding of the concept, we might be satisfied with a rather rough characterization. For example, BonJour (1985) defines coherence in terms of the following aspects:

1. A system of beliefs is coherent only if it is logically consistent.
2. A system of beliefs is coherent in proportion to its degree of probabilistic consistency.
3. The coherence of a system of beliefs is increased by the presence of inferential connections between its component beliefs and increased in proportion to the number and strength of such connections.

4. The coherence of a system of beliefs is diminished to the extent to which it is divided into subsystems of beliefs which are relatively unconnected to each other by inferential connections.
5. The coherence of a system of beliefs is decreased in proportion to the presence of unexplained anomalies in the believed content of the system.

Yet BonJour has little to say about the details of these aspects and how they are interrelated. If system A contains more inferential connections than system B, but B is less anomalous than A, which system is more coherent, all things considered?[10] BonJour's relative silence on this matter and other similar intricacies reveals, on a charitable reading, that his theory is not intended to provide guidance for deciding concrete cases, but to shed light on the nature of belief systems and the human epistemic condition for general philosophical purposes, in particular the prospects of defending one's totality of beliefs against global challenges.

A strikingly different account of coherence has been advanced by Paul Thagard (e.g., Thagard 2000). His "principles of explanatory coherence" are listed below:

Principle E1 (Symmetry) Explanatory coherence is a symmetric relation, unlike, say, conditional probability. That is, two propositions A and B cohere with each other equally.

Principle E2 (Explanation) a. A hypothesis coheres with what it explains, which can either be evidence or another hypothesis. b. Hypotheses that together explain some other proposition cohere with each other. c. The more hypotheses it takes to explain something, the lower the degree of coherence.

Principle E3 (Analogy) Similar hypotheses that explain similar pieces of evidence cohere.

Principle E4 (Data Priority) Propositions that describe the results of observation have a degree of acceptability on their own.

Principle E5 (Contradiction) Contradictory propositions are incoherent with each other.

Principle E6 (Competition) If A and B both explain a proposition, and if A and B are not explanatorily connected, then A and B are incoherent with each other (A and B are explanatorily connected if one explains the other or if together they explain something).

Principle E7 (Acceptance) The acceptability of a proposition in a system of propositions depends on its coherence with them.

Thagard has achieved even greater precision by translating his principles into computer code, allowing a computer to answer questions like "Is system A more coherent than system B?" in relation to specific issues (e.g., Thagard 2000).

The bottom line is that BonJour and Thagard can be viewed as taking the same exlicandum as the starting point of their respective investigations and yet ending up with radically diverging accounts, whereby the theories differ perhaps most saliently in the level of exactness reached. On an explicationist reconstruction, the differences observed in the final results are plausibly due to variations in aims and purposes affecting the relative weight assigned to the Carnapian requirements.[11]

For another illustration of the plurality of explicata, consider reliabilism as an account of the sense of knowledge of "awareness or familiarity gained by experience of a fact or situation." Reliabilism states that knowledge, in much this sense, amounts to reliably acquired true belief (e.g., Goldman 1986). It is arguably easier to assess, or intersubjectively agree upon, whether someone possesses a reliably acquired true belief than it is to assess whether the person is in a state of "awareness or familiarity gained by experience of a fact or situation."[12] It is plausible, therefore, to view reliabilism as a more exact version of the latter concept. Yet different theories may still disagree about the precisification of "reliable." A first, rough answer is that a process of belief acquisition is reliable just in case it tends "to produce beliefs that are true rather than false" (Goldman 1979/1992, p. 113). However, this account could be understood either in a modal or purely probabilistic way depending on what class of beliefs we are considering as a reference. Do we mean all beliefs—past, present and future—ever acquired by the subject? Do we mean only actually acquired beliefs or should we also count possibly

acquired beliefs, i.e., the employment of the process not only in the actual but also in some possible worlds (for instance, the closest ones to the actual world)? Having reflected on these issues and concluded that the probabilistic account needs to be combined with a contextual choice of reference class, Peter Baumann (2009, p. 87) draws the following enlightening conclusions:

> Finally, what about the alternative between modal interpretations and probabilistic interpretations of "reliability"? Aren't they more or less on a par, at least with respect to the issues discussed here? I don't think so. I think there are clear advantages on the side of the probabilistic version. Let me quickly mention two. First, closeness rankings of possible worlds seem restricted to ordinal rankings while the apparatus of probability theory can capture more than that and represent relations between differences of probabilities. Second, probability theory is closer to home if you're a naturalist than modal logic. The natural sciences are happy to use probability theory but seem to have little use for modal notions. I would therefore propose three things (in the light of all of the above); stick with reliabilism, go for a probabilistic version of it, and accept the contextualist implications of all that.

If this is correct, then how we choose to make the reliabilist account more precise will in the end depend on the purposes of our investigation, thus leading to a possible manifold of different potentially interesting explicata.

In fact, I believe that there is a third sense in which explicationist epistemology is plausibly pluralist, although I will not argue this point at length here. The sense I have in mind is that explicationism allows for a *plurality of sub-methodologies*. The thought is that ordinary language philosophy, experimental philosophy etc. need not be seen as competing enterprises but rather as complementary parts of a larger explicationist picture. A resource such as the Oxford Living Dictionary will only get you so far in mapping out the meaning of central epistemological terms. Ordinary language and intuition-based epistemology, as well as experimental work, are obviously useful for spotting more fine-grained epistemological distinctions or proposing hypotheses in this direction.

They are thereby relevant for identifying the relevant explicandum with greater precision and for contrasting it with related, though distinct, concepts in a more extensive fashion than a general dictionary admits. These methodologies are also potentially useful when assessing how close a given proposed explicatum is to a given explicandum and thus relevant for evaluating the extent to which a given candidate explicatum satisfies the first Carnapian requirement. Formal epistemology, by contrast, has so far given rise to a rich catalogue of potential epistemological explicata (cf. Hendricks 2006)—a welcome resource when deciding how to explicate a given explicandum for a particular purpose requiring formal precision. Some researchers may be more drawn toward one submethodology rather than another. For example, some are interested in experimental work, others in thought-experiments and intuition-based methodology. All this is fine from the explicationist perspective. The results of these diverse activities, properly conducted, are useful in the explication of central epistemological concepts for various contexts and purposes.

5   Comparison with Alston's Theory of Epistemic Desiderata

Epistemic pluralism has been advanced by other researchers without reference to Carnap. One well-known case in point is William Alston's theory of "epistemic desiderata" (Alston 1993). It is interesting to compare Alston's approach with explicationalist epistemology, as outlined above.

Alston is, in his study, concerned not with the concept of knowledge but with that of (epistemic) justification. Traditionally, it has been held that there is one concept of justification, and that identifying and shedding light on that concept is a, or perhaps even *the*, central task of epistemology. Alston takes issues with this school of thought. His starting point is a survey of what epistemologists have been proposing to be necessary conditions of justification. Thus, some have suggested that the belief in question must be based on grounds of the appropriate sort. Others have emphasized truth-conducivity: "[t]he reason or its content

must be so related to the target belief and its content that, given the truth of the former, the latter is thereby likely to be true" (Alston 1993, p. 528). Still others have maintained that what justifies a belief must be accessible to the subject. There are also higher-level requirements of the sort that the subject knows, or is justified in believing, that the ground of her belief is an adequate one. Finally, having coherent beliefs has been advanced as a necessary condition for justification, as has satisfying certain intellectual obligations. Observing the diversity of these conditions, Alston concludes (1993, p. 534):

> If we take the full range of parties to the disputes we have been considering, some of whom have had their thinking about 'epistemic justification' nourished primarily by some of the roots just mentioned and others about others, there does not seem to be enough commonality in their pre-theoretical understanding of the nature of epistemic justification to warrant us in supposing that there is some uniquely identifiable item about which they hold different views. It seems, rather, that they are high-lighting, emphasizing, 'pushing' different concepts, all called 'justification'. It seems, to switch to the perspective of this paper, that they are selecting different epistemic desiderata, or packages thereof, as deserving of the honorific title 'justification'.

In place of the traditional epistemological enterprise Alston proposes that we should seek to "disentangle the various epistemic desiderata involved in these discussions, aim at a penetrating understanding of each and of their interrelations, and explore the implications of this for epistemology" (1993, p. 538). A few pages later (1993, p. 542), these new tasks are organized under four heading (original emphasis): "the *elucidation* of desiderata; their *viability*; their *importance*; their *interrelations.*" The elucidation of the desiderata involves "understanding the nature of each of the epistemic desiderata that have figured in that discussion" (ibid.). The viability of a desideratum concerns the possibility of actually satisfying it in practice. By importance, Alston understands the "relative importance or centrality of one or another desideratum" (ibid.) in comparison with other desiderata. Alston states that he has nothing to say in the article about the interrelations between the desiderata.

How does Alston's account compare with explicationalist methodology á la Carnap? Let us focus first on the communalities. First of all, Alston, like Carnap, does not assume that just because there is a term in ordinary language—be it "knowledge" or "justification"—this means that it signifies a unique concept or idea. Rather, both thinkers are open to the possibility that words in ordinary language can have a variety of meanings that are more or less interrelated. Second, Carnap would have welcomed Alston's list of "epistemic desiderata" as presenting, in his terminology, a collection of possible explicanda in relation to "justification," i.e., a set of pre-systematic meanings of the latter that can constitute the starting point for further conceptual endeavors. In this sense, both perspectives imply the existence of a plurality of legitimate and potentially interesting epistemological projects that can be pursued under the heading "epistemic justification," and I take it that Alston would agree that the same is true regarding other epistemological terms.[13]

Furthermore, Alston takes the new goal of epistemology vis-à-vis justification to be not only the attainment of a deeper understanding of the various meanings of "justification" and their interrelations, tasks that are naturally viewed as belonging, in Carnap's framework, to the initial step whereby the explicandum is elucidated, but he also thinks—as did Carnap—that there are further goals to pursue once this clarificatory step has been completed. Yet this is where the similarities between Alston and Carnap seem to end. For Alston, the additional aims involve shedding light on the viability and importance of the meanings, or desiderata, in question. Alston's own preliminary investigations into these matters regarding "justification" essentially consist in distinctly philosophical reflections on well-known distinctions in epistemology, as traditionally pursued, such as externalist vs. internalist justification and global vs. local skepticism. For Carnap, by contrast, what remains is the task of actually explicating the concept that has been singled out as the explicandum, i.e., transforming the latter into an exact concept that is adequate and fruitful for specialized purposes. Carnap's focus is on reconstructing and improving our pre-systematic concepts for use in circumstances that require greater precision. Carnap thought that serious philosophy presents just the kind of context requiring increased

exactness and conceptual development with respect to crucial terms, and that this was one way in which philosophy could advance beyond the stage of persistent disagreement to a scientific discipline. Alston, it seems, would disagree.

Notes

1. Another useful resource for understanding Carnap's method of explication is his debate with Strawson in the Schilpp volume dedicated to Carnap's work (Schilpp, 1963). See Strawson (1963) and Carnap (1963). I discuss various objections to Carnap's method, by Strawson and others, in Olsson (2015).
2. For the purposes of the present exposition, a sharp distinction between terms and concepts is unnecessary.
3. A second point of justified criticism concerns the way in which Carnap phrases his requirement of fruitfulness in terms of the number of universal statements in which the concept figures. Taken literary, Carnap is implying that any old universal statement will do. Yet his clarification in terms of "empirical laws" and "logical theorems" suggests that the statements in question must be plausibly true and also of a certain theoretical standing.
4. In Olsson (2015), I introduced the term "explicative epistemology" but I now prefer "explicationalist epistemology."
5. Lehrer proceeds to claim that there is "underlying simplicity" on the grounds that "[k]nowledge reduces to undefeated justification, a just reward for our arduous analytical efforts" (1990, p. 149). However, it remains true that Lehrer's definition of knowledge is hardly simply on a strict application of Carnap's criterion of definitional simplicity which requires that the form of the definition be simple (Carnap 1950, p. 7). Lehrer might respond that this shortcoming is offset by the simplicity of the laws for the formulation of which his explicatum is useful, in particular the supposed law that knowledge reduces to undefeated justification, appealing to the second aspect of simplicity on Carnap's account. However, it is doubtful whether this universal statement expresses a lawlike connection of the kind Carnap had in mind as it is a mere analytical consequence of the definitions of the concepts involved.

6. David Lewis is a well-known case. As Brian Weatherson (2016) notes, in his entry on Lewis in Stanford Encyclopedia, concerning the vagueness of Lewis's account of conventions: "Lewis, characteristically, thought this was a feature not a bug of the view. Our intuitive notion of a convention is vague, and any analysis of it should capture the vagueness. The idea that analyses of imprecise folk concepts should be imprecise recurs throughout Lewis's career."

7. This point is argued at length in Olsson (2015).

8. See Bach (1985), p. 248: "… there surely are two conceptions of justified belief involved in the debate, the internalist and the externalist conception. Laurence BonJour has contrasted them nicely. Internalism requires that a person have 'cognitive grasp' of whatever makes his belief justified. Being justified depends on how rational and 'epistemically responsible' (whatever this means more precisely) he is in coming to hold the belief. In contrast, the externalist (reliabilist) conception allows that the source of justification can be 'external to the person's subjective conception of the situation'." The BonJour reference here is to his 1980 article.

9. See https://en.oxforddictionaries.com/definition/knowledge, retrived on May 15, 2017. The sense in which the term applies to computer systems was added recently. When I first consulted the dictionary, in April of 2017, that sense was not yet listed.

10. See Olsson (2017) for a fuller discussion of BonJour's coherence theory. See also Olsson (2005), Chap. 4.

11. See Olsson (2017) for more details on Thagard and coherence theories in general. See also Olsson (2005), Chap. 9.

12. Conee and Feldman (1998) have forcefully questioned the possibility of reaching intersubjective agreement on matters of reliability except in special circumstances. The empirical study in Jönsson (2013) concluded, by contrast, that people often converge on judgements of reliability in normal cases. See also Olsson (2016) for a discussion of this aspect of the so-called generality problem for reliabilism.

13. See Peels (2010) for a contrasting assessment. Peels argues that Alston's theory of epistemic desiderata is not as pluralist as Alston claims it to be.

References

Alston, W.P. 1993. Epistemic Desiderata. *Philosophy and Phenomenological Research* 53 (3): 527–551.

Bach, K. 1985. A Rationale for Reliabilism. *The Monist* 68: 246–263.

Baumann, P. 2009. Reliabilism—Modal, Probabilistic or Contextualist. *Grazer Philosophische Studien* 79: 77–89.

BonJour, L. 1980. Externalist Theories of Empirical Knowledge. *Midwest Studies in Philosophy* 5: 53–73.

———. 1985. *The Structure of Empirical Knowledge*. Cambridge: Harvard University Press.

Carnap, R. 1950. *Logical Foundations of Probability*. Chicago: University of Chicago Press.

———. 1963. Replies and Systematic Expositions. In *The Philosophy of Rudolf Carnap*, ed. P. Schilpp, 859–1013. LaSalle, IL: Open Court.

Conee, E., and R. Feldman. 1998. The Generality Problem for Reliabilism. *Philosophical Studies* 89: 1–29.

Goldman, A.I. 1979/1992. What Is Justified Belief? In *Liaisons: Philosophy Meets the Cognitive and Social Sciences*, 105–126. Cambridge/MA & London: MIT Press. First published as pp. 1–23 in Pappas, G.S. ed. *Justification and Knowledge*. Dordrecht: D. Reidel, 1979.

———. 1986. *Epistemology and Cognition*. Cambridge: Harvard University Press.

———. 1999. *Knowledge in a Social World*. Oxford: Oxford University Press.

Hendricks, V.F. 2006. *Mainstream and Formal Epistemology*. New York: Cambridge University Press.

Jönsson, M. 2013. A Reliabilism Built on Cognitive Convergence: An Empirically Grounded Solution to the Generality Problem. *Episteme* 10: 241–268.

Kornblith, H. 2002. *Knowledge and Its Place in Nature*. Oxford: Clarendon Press.

Lehrer, K. 1990. *Theory of Knowledge*. Boulder: Westview.

Moore, G.E. 1899/1970. The Nature of Judgement. In *G E Moore: Essays in Retrospect*, Muirhead Library of Philosophy, eds. A. Ambrose and M. Lazerowitz, 89–101. London: Routledge. First published in *Mind*, New Series Vol. 8, no. 30 (April, 1899): 176–193.

Olsson, E.J. 2005. Against Coherence: Truth, Probability, and Justification. Oxford: Clarendon Press.

———. 2015. Gettier and the Method of Explication: A 60 Year Old Solution to a 50 Year Old Problem. *Philosophical Studies* 172: 57–72.

———. 2016. A Naturalistic Approach to the Generality Problem. In *Philosophers and Their Critics*, eds. B.P. McLaughlin and H. Kornblith, 178–199. Malden: Wiley-Blackwell.

———. 2017. Coherentist Theories of Epistemic Justification. *The Stanford Encyclopedia of Philosophy* (Spring 2017 Edition), ed. Edward N. Zalta. URL = <https://plato.stanford.edu/archives/spr2017/entries/justep-coherence/>.

Peels, R. 2010. Epistemic Desiderata and Epistemic Pluralism. *Journal of Philosophical Research* 35: 193–207.

Quine, W.V.O. 1960. *Word and Object*. Cambridge, MA: MIT Press.

Quine, W.V.O. 1969. Epistemology Naturalized. In *Ontological Relativity and Other Essays*, 60–90. New York: Columbia University Press.

Schilpp, P.A. 1963. *The Philosophy of Rudolf Carnap*. LaSalle, IL: Open Court.

Strawson, P.F. 1963. Carnap's Views on Constructed Systems Versus Natural Languages in Analytic Philosophy. In *The Philosophy of Rudolf Carnap*, ed. P. Schilpp, 503–518. LaSalle, IL: Open Court.

Thagard, P. 2000. *Coherence in Thought and Action*. Cambridge: The MIT Press.

Weatherson, B. 2016, David Lewis. *The Stanford Encyclopedia of Philosophy (Winter 2016 Edition)*, Zalta, E.N. (ed.).

Pure Epistemic Pluralism

Nikolaj Jang Lee Linding Pedersen

1   Pluralism

Here is a working definition of monism and pluralism (with respect to *X*):

N.J.L.L. Pedersen (✉)
Underwood International College, Yonsei International Campus,
Incheon 406-840, South Korea
e-mail: nikolajpedersen@gmail.com

Monism: there is exactly one way of being *X*.

Pluralism: there are several ways of being *X*.

Pluralist views have recently attracted considerable attention in different areas of philosophy. Truth and logic are cases in hand. According to truth pluralism, there are several ways of being true. Statements concerning the empirical world ('There are mountains'), legal discourse ('Speeding is illegal'), mathematics ('2 + 5 = 7'), and ethics ('Genocide is wrong') may be true in different ways. For example, 'There are mountains' might be true in virtue of corresponding with reality, while 'Speeding is illegal' might be true in virtue of cohering with the body of law.[1] According to logical pluralism, there are several ways of being valid. Validity is to be understood in terms of cases: an argument is valid if and only if in every case in which the premises are true and the conclusion is true. There are equally legitimate ways of construing 'case'. It can be construed as possible worlds (complete and consistent), constructions (potentially incomplete), or situations (potentially inconsistent). These different construals of cases go with different logics—respectively, classical logic, intuitionistic logic, and relevance logic. However, since the three construals of 'case' are equally legitimate, they capture equally legitimate ways of being valid. Hence, there are several ways of being valid.[2]

How about pluralism in epistemology? This chapter is dedicated to an investigation and reconfiguration of the epistemic pluralism debate. Once reconfigured, I explore the prospects of what I consider the most interesting form of epistemic pluralism: pluralism about non-derivative epistemic goods (i.e. epistemic goods that do not derive their value or goodness from some other epistemic good). I proceed as follows:

In Sect. 2, I look to the work of Tyler Burge, Alvin Goldman, and William Alston for examples of epistemic pluralism. Section 3 investigates what rationale can be given for epistemic pluralism. Drawing on the literature on truth pluralism, I suggest that one rationale for adopting a pluralist view in epistemology is that they put one in a position to accommodate a wider range of epistemic assessments. Epistemic pluralism has a wider scope than epistemic monism.

In Sect. 4, I do two things. First, I argue that the distinction between epistemic monism and epistemic pluralism is most interestingly drawn at the level of non-derivative epistemic goods. Second, I make the observation that, at a very fundamental level, the varieties of epistemic pluralism presented in Sect. 2 share a very significant feature: they are all combined with veritic unitarianism, a form of monism about fundamental epistemic goods. The view says that there are several epistemic goods but that truth is the only non-derivative epistemic good. The remainder of the chapter is dedicated to the task of investigating the prospects of pluralism about non-derivative epistemic goods—what I call 'pure epistemic pluralism'. I assume that truth is a non-derivative epistemic good. On this assumption, the identification of another fundamental epistemic good will point to a species of pure epistemic pluralism. What, other than truth, might qualify as a fundamental epistemic good?

In Sect. 5, I turn to the work of Michael DePaul and Jonathan Kvanvig. Each of them seems to be committed to pure epistemic pluralism. However, while DePaul seems to be committed to saying that there is *some* fundamental epistemic good distinct from truth, he does not give many pointers as to what that good might be. Kvanvig, on the other hand, does so. I suggest that he is committed to taking grasp of coherence–conferring relations to be a non-derivative epistemic good. In Sect. 6, I present two 'collapse arguments': arguments that are supposed to show that pluralism inherently unstable because it collapses into epistemic monism. I start by considering the collapse arguments, as formulated against truth pluralism. Having discussed the case of truth, I move on to the epistemic case. In Sect. 7, I tackle the two collapse arguments on behalf of the epistemic pluralist.

2   Some Varieties of Epistemic Pluralism

In Sect. 1, I noted that truth pluralism and logical pluralism have attracted a great deal of attention recently. However, pluralism has also made inroads into epistemology. Indeed, versions of epistemic pluralism can be found in the work of several prominent epistemologists.

In this section, I trace pluralist strands in the works of Tyler Burge, Alvin Goldman, and William Alston.

According to Burge, epistemic warrant is a genus that divides into two species: internalist warrant and externalist warrant. Burge presents epistemic justification as an internalist species of warrant and epistemic entitlement as an externalist one. (Below I allow myself to leave out the label 'epistemic'. Unless otherwise stated, I have in mind epistemic warrant, justification, rationality, etc.) Burge understands the distinction between justification and entitlement in terms of reasons. On his view 'A justification is a warrant that consists partly in the operation or possession of a reason. An individual is justified if and only if the reason is operative or relied upon in the individual's psychology'. Entitlement, on the other hand, requires reliability, but it does not consist in the individual's possessing or relying on a reason, not even partly.[3] Thus, a subject S may be entitled to believe that p although no reason is operative or relied on in S's psychology. Justification and entitlement are both species of warrant. However, they are radically different in nature. The former involves reasons, the latter does not. Hence, on Burge's view, there are several ways of being warranted. Burge is a pluralist about epistemic warrant.[4]

In a 1988 paper, Goldman says that he wants to propose a 'contrasting pair of conceptions of justification, and hold that *both* are defensible and legitimate'.[5] The distinction he wants to draw is between what he calls, respectively, 'strong justification' and 'weak justification'. A subject's belief is strongly justified provided that it is formed through a reliable process or method M, and the reliability of M is not undermined by the agent's cognitive state (e.g., by a belief to the effect that the process or method is unreliable). A subject S's belief is weakly justified provided that it is (i) formed through an unreliable process or method M, (ii) S does not believe M to be unreliable, (iii) there is no reliable process or method available to S which would classify M as being unreliable, (iv) there is no method or process that S believes to be reliable that, if used, would lead S to believe that M is unreliable, (v) S believes M to be reliable, and (vi) everyone in the community of which S is a member trusts and uses M, and S has good reason to trust other members of the community on many matters and has no decisive reason to distrust

their confidence in M.[6] Strong justification requires reliability. Matters are different for weak justification. Indeed, weak justification requires that the relevant method or process be <u>un</u>reliable. Strong justification is meant to capture a reliabilist notion of justification. In Goldman's words, weak justification is meant to capture '*ill-formed*-but- blameless belief'. This category covers cases of beliefs formed through unreliable processes or methods where it is beyond the intellectual scope of the subject to detect the unreliability of the relevant process or method. As noted, weak justification requires unreliability. However, since the basic idea is that the relevant method or process is reliable from the subject's perspective and its unreliability is beyond her epistemic reach, weak justification can still be regarded as being integrated into a broadly reliabilist framework.[7] Strong justification and weak justification are radically different in nature. Nonetheless, Goldman deems both of them 'defensible and legitimate'. He thus endorses pluralism about epistemic justification. He thinks that there are several ways of being justified.

William Alston arrives at epistemic pluralism through reflection on the history of epistemology. A prominent feature of this history is persistent disagreement concerning justification. According to Alston, the right reaction to this datum is to say that there is no unique property of beliefs that gets picked out by the word 'justified'. Instead, distinct notions of justification mark different epistemic desiderata that beliefs can enjoy.[8] Being true is one such desideratum. It is epistemically desirable for a belief to be true. Other related—but distinct—desiderata include having adequate evidence, being based on adequate evidence, being formed via a (sufficiently) reliable belief-forming process, and being Plantinga-warranted.[9] These desiderata are related to truth in the sense of being truth-conducive. They make it likely that the relevant belief is true. At the same time, these desiderata are distinct from the desideratum of being true because a belief's having them does not guarantee that the belief is true. Alston also endorses a cluster of desiderata that facilitate identification and formation of true beliefs: having high-grade cognitive access to the evidence for a belief, knowing or having a well-grounded belief to the effect that a given belief enjoys a certain positive epistemic status, and being able to carry out a successful defence of the probable truth of a belief.[10] These are higher-order

desiderata that, strictly speaking, apply to the subject. They involve the subject's conceptualizing some target belief as enjoying a certain epistemic status, the subject's accessing evidence, or offering reasons supporting the belief. In this respect, higher-order desiderata are different from the desideratum of being true and the three truth-conducive desiderata. None of these desiderata involves the subject's conceptualizing any belief, accessing evidence in its favour, or offering reasons to support it.

Being true, the three truth-conducive features and the three higher-order features are distinct. However, Alston thinks that they are all legitimate epistemic desiderata. While epistemologists have disagreed over whether these various features are necessary for justification, they agree that the features are desirable or good from an epistemic point of view. The features are all ways for a belief to be epistemically desirable. Hence, Alston is a pluralist about epistemic desiderata.[11]

3   The Epistemic Scope Problem

Burge, Goldman, and Alston are epistemic pluralists. However, one might wonder: why be an epistemic pluralist? I approach this question by considering the parallel question for truth pluralism. I present the motivation most commonly given by truth pluralists and then proceed to show that the same kind of motivation is what drives Burge, Goldman, and Alston.

Why think that truth has a plural nature? To answer this question, truth pluralists most commonly appeal to the so-called scope problem. In the abstract, the problem can be presented as follows:

> **Alethic scope problem**: propositions $p_1, \ldots, p_n$ should all be classified as being true, but no single theory of truth can plausibly account for the truth of all of $p_1, \ldots, p_n$.

Propositions $p_1, \ldots, p_n$ are drawn from the full range of truth-apt domains of discourse. The correspondence theory might plausibly be applied to a proper subset of $p_1, \ldots, p_n$—say, propositions about

concrete objects such as tables, chairs, and mountains. However, the theory does not seem like a plausible candidate for legal discourse. The coherence theory might plausibly be applied to the legal domain, but seems much less plausible as a candidate for discourse about concrete entities such as tables, chairs, and mountains. Similarly for other monist accounts of truth. In sum, the problem for monist accounts is that their scope is not sufficiently wide to plausibly account for the truth of *all* of $p_1, \ldots, p_n$.

Enter the truth pluralist. According to the truth pluralist, different accounts of truth should be regarded as complementary rather than as being in competition with one another. Several accounts should be brought together to do joint service. By taking on board several ways of being true, the scope of truth pluralism is sufficiently wide plausibly to accommodate the truth of all of $p_1, \ldots, p_n$. For example, the pluralist can say that the truth of 'There are mountains' is to be accounted for in terms of correspondence, while the truth of 'Speeding is illegal' is to be accounted for in terms of coherence with the body of law.[12]

Turn now to the epistemic case. Let F be a positive epistemic standing enjoyed by epistemic value bearers $x_1, \ldots, x_n$. This is the scope problem for F:

> **Epistemic scope problem**: epistemic value bearers $x_1, \ldots, x_n$ should all be classified as enjoying positive epistemic standing F, but no single theory of F can plausibly account for all of $x_1, \ldots, x_n$ being F.

Let us consider Burge, Goldman, and Alston in turn, substituting, respectively, 'warrant', 'justification', and 'epistemically desirable' for 'F'. I suggest that their endorsement of epistemic pluralism is motivated by an epistemic scope problem.

Burge writes:

> The claim that reason or justification is the only sort of epistemic warrant can be seen as a stipulative restriction on what "warrant" is to mean. But if it is a substantive claim, it hyper-intellectualizes epistemology. It focuses entirely on a kind of epistemic good that derives from the more intellectual aspects of the more intellectual representational systems. In so doing

it deprives epistemology of resources to account for more primitive, but nearly ubiquitous epistemic goods. Children and higher non-human animals do not have reasons for their perceptual beliefs. They lack concepts like reliable, normal condition, perceptual state, individuation, defeating condition, that are necessary for having such reasons. Yet they have perceptual beliefs. There is no sound basis for denying that epistemology can evaluate these beliefs with respect to norms governing their formation … There is no sound basis for denying that epistemology can evaluate their perceptual beliefs for epistemic warrant. There are legitimate questions about animals' and young children's entitlement to their perceptual beliefs.[13]

From this passage, it is clear that Burge thinks that it would be too restrictive to take the only kind of warrant to be justification, i.e. warrant involving a reason operative in the individual's psychology. Indeed, he explicitly says that justification 'cannot do the work of a full epistemology' and that it is 'not the only sort of epistemic warrant' (2003, p. 528). Taking a monist path, restricted to justification or internalist warrant, would hyper-intellectualize epistemology. It would turn it into a discipline focused only on an exclusive club of reasons-enriched subjects, leaving behind subjects without the requisite level of sophistication. According to Burge, this would be a mistake. Perceptual beliefs should lie within the realm of epistemic warrant, even when held by young children and higher animals—subjects incapable of possessing or relying on reasons. Burge's distinction between justification and entitlement is thus motivated by an epistemic version of the scope problem. In order to give warrant a scope sufficiently wide plausibly to cover all beliefs that should lie within the scope of warrant, both justification and entitlement are needed.

Turn to Goldman and his endorsement of strong and weak justification. To introduce the notion of weak justification, Goldman asks us to consider the following epistemic community:

Consider a scientifically benighted culture, of ancient or medieval vintage. This culture employs certain highly unreliable methods for forming beliefs about the future and the unobserved. Their methods appeal to the doctrine of signatures, to astrology, and to oracles. Members of the

culture have never thought about probability theory or statistics, never dreamt of anything that could be classed as 'experimental method'. Now, suppose that on a particular occasion a member of this culture forms a belief about the outcome of an impending battle by using one of the aforementioned methods, say, by consulting zodiacal signs in a culturally approved fashion. Call this method M. Is this person's belief justified …?[14]

Goldman himself answers the question concerning justification in the affirmative, providing the following explanation:

> Why … is some attraction felt toward a positive answer? This seems to stem from the cultural plight of our believer. He is situated in a certain spatio-historical environment. Everyone else in this environment uses and trusts method M. Moreover, our believer has good reasons to trust his cultural peers on many matters, and lacks decisive reason for distrusting their confidence astrology.
>
> … It is beyond his intellectual scope to find flaws in M. Thus, we can hardly *fault* him for using M, nor fault him therefore for believing what he does. The belief in question is epistemically *blameless*, and that seems to explain why we are tempted to call it *justified*.[15]

Goldman's distinction between strong and weak justification is motivated by an epistemic scope problem. Beliefs that are acquired through reliable belief-forming processes or methods qualify as justified. Visual perception and memory are reliable belief-forming processes, and so beliefs formed via these processes should count as justified. In Goldman's sense, perceptual beliefs and memory-based beliefs are strongly justified. Now, according to Goldman, the belief concerning the impending battle should also fall within the scope of justification. The same goes for other beliefs that are relevantly similar—that is, beliefs that have been acquired through unreliable belief-forming processes or methods, but blamelessly so. However, blameless beliefs formed via unreliable processes or methods cannot be strongly justified, as strong justification requires reliability. Thus, in order for the relevant

range of beliefs to qualify as justified, another species of justification is required. Goldman's suggestion is to take on board weak justification.

Finally, let us consider Alston. The epistemic scope problem lies at the very root of his epistemic desiderata approach. Recall that he proposes the desiderata approach against the background of a rejection of monism about justification. He takes the long-running debate between alternative accounts of justification to suggest that there is no uniquely correct concept of justification. Rather than being in direct conflict and competing for space, Alston thinks that different accounts of justification should all be accommodated. No single account captures exactly the conditions under which beliefs are justified. They are all on to something in the sense that they all track different epistemic desiderata.[16] Excluding certain desiderata would mean excluding certain features of beliefs that play a positive role in cognition or enquiry, thus leaving epistemology incapable of issuing a sufficiently wide range of positive epistemic assessments.

In the previous section, I traced pluralist strands in the works of Burge, Goldman, and Alston. In this section, I suggested that, in each case, it is the wider scope of pluralism with respect to some target epistemic standing that motivates the pluralist stance. In the next section, I further discuss the nature of the forms of pluralism represented by Burge, Goldman, and Alston. I suggest that they are rather moderate in nature and use this as a stepping stone for formulating a different, more radical form of pluralism.[17]

4   Pure Epistemic Pluralism

How pure are the kinds of epistemic pluralism discussed in Sects. 2–3? Not so pure. The goal of this section is to explain why and to formulate a purer kind of epistemic pluralism: pluralism about non-derivative epistemic goods. I call this kind of pluralism 'pure epistemic pluralism'.

To see why the forms of epistemic pluralism presented in Sects. 2–3 are not particularly pure, let us dwell on the connection between

positive epistemic standings and epistemic goods for a moment. Epistemic justification, warrant, rationality, and desiderata mark positive standings. If a belief is epistemically justified, warranted, rational, or desirable, it enjoys a positive epistemic standing. Similarly, if an action is morally justified, warranted, or rational, it enjoys a positive moral standing.

Positive standings do not float freely. Positive standings of a given kind qualify as such by being connected to goods of that kind. Specific standings are worthy of the label 'moral justification', 'moral warrant', or 'moral rationality' in virtue of conferring upon actions a positive standing vis-à-vis some moral good. In the same way, certain standings earn the label 'epistemic justification', 'epistemic warrant', 'epistemic rationality', or 'epistemic desideratum' in virtue of conferring a positive standing upon beliefs or other bearers of epistemic value with respect to some epistemic good. Put schematically:

> **Positive *X*-standings and *X*-goods**: a given standing counts as a positive standing of kind *X* in virtue of conferring upon its bearers a positive standing with respect to a good of kind *X*.

This thesis concerning positive standings and goods speaks of goods of kind *X* (or *X*-goods). But what might the space of *X*-goods look like? One possibility is what might aptly be called '*X*-unitarianism':

> **X-unitarianism**: there are several *X*-goods and a single non-derivative *X*-good. The former derive their value from the latter.

X-unitarianism is compatible with there being a plurality of *X*-goods. However, the space of *X*-goods comes with a certain structure. All derivative goods depend on the single non-derivative *X*-good for their value. As such, the space of *X*-goods is highly unified. Hence, the label 'unitarianism'. Formulated as an epistemic thesis unitarianism has it that there is a range of epistemic goods that all derive their value from a single non-derivative epistemic good.[18] The most prominent form of epistemic unitarianism places truth at the centre of the space of epistemic goods:

> **Veritic unitarianism:** there are several epistemic goods and a single non-derivative epistemic good, truth. The former derive their value from the latter.

Veritic unitarianism is a very widely held view in epistemology. Proponents include Burge, Goldman, and Alston.[19]

Why is veritic unitarianism such a widely held view? One plausible explanation is that it makes for a natural fit with the widely held view that truth is the goal of enquiry and cognition. One of the main tasks of epistemology is to evaluate beliefs, as they feature in enquiry and cognition. Since truth is the goal of enquiry and cognition, the relevant standard for epistemic evaluation of beliefs is how they do in relation to this goal. Different standings are thus evaluated positively from an epistemic point of view if they confer a positive standing upon beliefs with respect to truth.

I claimed earlier that the species of epistemic pluralism discussed in Sects. 2–3 are not particularly radical in nature. This can be explained by reference to veritic unitarianism. The varieties of epistemic pluralism presented earlier are not particularly radical in nature because they are all combined with veritic unitarianism. There is a plurality of positive epistemic standings (different species of epistemic warrant, justification, or desiderata). These are all as epistemic goods. If a belief enjoys any of these standings, there is something epistemically good about it. However, this is accounted for by the connection that each of these standings bears to truth, the goal of enquiry and cognition. Here, e.g., is Burge endorsing veritic unitarianism: 'Entitlement is a subspecies of epistemic warrant. Epistemic warrant, and hence entitlement, is an epistemic good. The epistemic good, warrant, is essentially associated with the fundamental representational good—truth. The notion of an epistemic good must be understood in relation to this fundamental representational good'.[20] As seen earlier, Burge endorses two distinct species of epistemic warrant, entitlement and justification. From the passage just quoted, it is clear that Burge regards both kinds of warrant as epistemic goods. However, it is likewise clear that he takes the epistemic goodness of both to derive from their connection to the same fundamental epistemic good, truth. For entitlement, it is straightforward to

account for why entitled beliefs enjoy a positive standing vis-à-vis the fundamental epistemic good of truth. They do so because reliability is a necessary condition for a capacity to generate entitled beliefs—where reliability is understood as the capacity's yielding a sufficiently high proportion of true beliefs. How about justification? Recall that Burge takes justification to involve a reason in the justified individual's psychology. Reasons to believe *p* support the truth of *p*. Hence, justified belief in Burge's sense enjoys a positive standing vis-à-vis the fundamental epistemic good of truth.

Goldman endorses veritic unitarianism.[21] His notions of strong justification and weak justification derive their value from being truth-connected, and so they fit into this axiological framework. We might say that strong justification is *objectively truth-connected,* while weak justification is *subjectively truth-connected.*[22] A belief's being strongly justified requires the belief-forming process or method to be reliable. Thus, strong justification confers on a belief a high objective probability of being true. In this sense, a belief's being strongly justified confers upon it an objective positive standing in relation to the fundamental epistemic good of truth. A belief's being weakly justified involves the subject's taking the relevant process or method to be reliable and doing so in a well-supported manner (i.e. everyone in the subject's community uses *M*, the subject has good reason to trust other members of the community on many matters, and the subject has no decisive reason to distrust other community members' confidence in *M*). Thus, while the relevant process or method is not in fact reliable, from the subject's perspective, there is reason to deem it reliable. Thus, a belief's being weakly justified confers upon it a subjective positive standing in relation to the fundamental epistemic good of truth.[23]

Alston, like Burge and Goldman, endorses veritic unitarianism:

> 'We evaluate something epistemically … when we judge it to be more or less good or bad from the epistemic point of view, that is, for the attainment of epistemic purposes. And what purposes are those?

> We can best approach this question by reminding ourselves that epistemology consists of a critical reflection on human cognition. And the

evaluative aspect of epistemology involves an attempt to identify ways in which the conduct and the products of our cognitive activities can be better or worse vis-a-vis the goals of cognition. And what are those goals? Along with many other epistemologists I suggest that the primary function of cognition in human life is to acquire true rather than false beliefs about matters that are of interest or importance to us'.[24]

Not surprisingly, as seen above, Alston takes being true to be an epistemic desideratum. To show that Alston's pluralism fits with veritic unitarianism, the other epistemic desiderata that feature in Alston's pluralism must be shown to derive their epistemic value from the fundamental epistemic good of truth. Let us first consider the desiderata of having adequate evidence, being based on adequate evidence, being formed via a (sufficiently) reliable belief-forming process, and being Plantinga-warranted.

I allow myself to leave reliability aside, as it has already been dealt with above in connection with Goldman. Alston takes a belief's being based on adequate evidence to be an epistemic desideratum because something counts as evidence for a belief in virtue of supporting the *truth* of that belief. Hence, a belief that is based on adequate evidence enjoys a positive standing vis-à-vis truth. Alston takes a subject's belief's being based on adequate evidence to be an actualization of the possibility given by the subject's having adequate evidence for the belief. On his view, a subject's having adequate evidence for a belief means that, if the subject were to base the belief on that evidence, the belief would be likely to be true. In this sense, having adequate evidence marks a truth-connection and confers a positive standing upon the relevant belief with respect to truth. However, Alston also considers having adequate evidence to be less of an epistemic good than being based on adequate evidence.[25]

As for Plantinga-warrant, Alston talks about beliefs being formed by properly functioning cognitive capacities. What he has in mind is warrant, as construed by Plantinga (hence the choice of label). According to Plantinga, a belief's being warranted requires (i) the belief to be produced by a properly functioning cognitive capacity C, (ii) C to be working in a cognitive environment appropriate for a capacity

of its kind, (iii) the segment of the design plan governing the formation of the belief to be aimed at the production of true beliefs, and (iv) given (i)–(iii), there has to be a high (statistical) probability that the belief is true.[26] Since the stated clauses are *necessary* for warrant, every instance of warranted belief will satisfy clauses (i)–(iii). However, since clause (iv) states that warranted belief must have a high statistical probability of being true given (i)–(iii), this means that any Plantinga-warranted belief has a high (statistical) probability of being true. Plantinga-warrant thus underwrites a truth-connection and confers upon beliefs a positive standing with respect to truth.[27]

Let us now consider the higher-order desiderata endorsed by Alston: (i) having high-grade cognitive access to the evidence for a belief, (ii) knowing or having a well-grounded belief to the effect that a given belief enjoys a certain positive epistemic status, and (iii) being able to carry out a successful defence of the probable truth of a belief. How might one argue for the truth-connectedness of features (i)–(iii)? There are several kinds of cases that one can make. First, the features seem to be at least subjectively truth-connected. If S has high-grade cognitive access to evidence for a given belief or is able to carry out a successful defence of the probable truth of the belief, then S's belief is subjectively truth-connected. From S's point of view, the belief is likely to be true, and so features (i) and (iii) confer on the belief a subjective positive status vis-à-vis truth. Feature (ii) would also seem to be subjectively truth-connected, assuming that any positive status is either objectively or subjectively truth-connected. If S knows or has a well-grounded belief to the effect that the belief that p enjoys a certain positive epistemic status that is truth-connected, this would seem to suffice for the subjective truth-connectedness of S's belief that p. From S's point of view the belief that p is likely to be true and, hence, feature (ii) confers on the belief a subjective positive status vis-à-vis truth.[28] Second, one can make a case for the three higher-order desiderata's being indirectly truth-connected because they help the subject manage her cognitive life in a way that favours acquiring and sustaining true beliefs. The desiderata can play this role because they involve the subject's recognizing or being able to defend the epistemically good standing of beliefs. What this tells us is that, when one of the three higher-order features is present, the

subject is well-positioned to acquire or sustain a true belief. That is, the higher-order features put the subject in a good position to acquire or sustain beliefs that are truth-connected. This makes them epistemically valuable.[29]

Earlier in this section I claimed that the versions of epistemic pluralism found in Burge, Goldman, and Alston are rather moderate in nature. I have defended this claim by arguing that each version of epistemic pluralism is combined with veritic unitarianism. Each species of epistemic warrant, justification, and desiderata embraced by respectively Burge, Goldman, and Alston derives its value from truth. Thus, the pluralist views of Burge, Goldman, and Alston are highly unified axiologically speaking.

Our reflections on veritic unitarianism and its subsumption of the three forms of epistemic pluralism discussed above suggest that the monism/pluralism distinction is most interestingly drawn at the level of non-derivative epistemic goods. If we want to see how radical or moderate a specific version of epistemic pluralism is, we should look at its commitments at the level of non-derivative epistemic goods. I thus suggest that epistemic pluralism—formulated at the most fundamental level—amounts to the following thesis:

> **Pure epistemic pluralism**: there are several non-derivative epistemic goods $G_1, \dots, G_n$.

Pure epistemic pluralism is a fundamental form of epistemic pluralism because non-derivative goods constitute the ultimate source of epistemic normativity. Positive epistemic standings qualify as such because they promote some non-derivative epistemic good. Endorsing pure epistemic pluralism radically widens the scope of epistemology. It does so because it widens the basis of epistemic goodness and, accordingly, the basis of epistemic normativity. Instead of connecting all positive epistemic standings to a single non-derivative epistemic good, pure epistemic pluralism ties positive epistemic standings to a plurality of non-derivative epistemic goods. Bearing in mind the constitutive connection between goods of a certain kind and positive standings of that kind, for each non-derivative epistemic good G_i ($1 \leq i \leq n$), there is a range of

positive epistemic standings: G_i-warrant, G_i-justification, G_i-rationality, and G_i-desideratum. Each of these is to be understood in terms of its connection to G_i. Within each G_i-cluster of positive epistemic standings, there is unity. It is provided by the non-derivative epistemic good of that cluster. However, there is no unity across the clusters in the sense of there being an overarching good from which each of $G_1, \ldots, G_n$ derives its value.

5   Pure Epistemic Pluralism, Anyone?

In the previous section, I introduced pure epistemic pluralism, the view that there are several non-derivative epistemic goods. The rest of the paper is dedicated to a preliminary investigation of this type of pluralism. One immediate issue is what reasons might be given in support of pure epistemic pluralism—and who, if anyone, holds the view. The present section addresses this issue. I proceed on the assumption that pure epistemic pluralists agree with advocates of veritic monism and veritic unitarianism that truth is a non-derivative epistemic good. Taking on board this assumption, pure epistemic pluralism can be supported by making a case that there is some non-derivative epistemic good distinct from truth. Drawing on the work of Michael DePaul and Jonathan Kvanvig, I present two ways that such a case might be made.

Let us first turn to DePaul. His take on the so-called value problem suggests that he is sympathetic to pluralism about non-derivative epistemic goods.

It is widely held that knowledge is more valuable than mere true belief. It is also widely held that truth is the only non-derivative epistemic good, as seen above. Truth is standardly taken to be required for knowledge. However, if truth was the only non-derivative epistemic good, it is difficult to see how other constituents of knowledge could contribute *additional* value and thereby push the overall value of knowledge beyond the value of mere true belief. This is the value problem. DePaul thinks that the value problem has bite. He takes it to show that there must be some epistemic good distinct from truth that does not derive its value from truth. Combining this line on the value problem

with the assumption that truth is a non-derivative epistemic good amounts to a case in favour of pure epistemic pluralism.[30]

It is worth noting that, if sound, the argument just given only delivers an unspecific version of pluralism about non-derivative epistemic goods. The argument at most delivers an existential conclusion: there is *some* non-derivative epistemic good distinct from truth. This naturally raises the question what goods other than truth might fall into the category of non-derivative epistemic goods. This question is our cue to move on to Kvanvig. Kvanvig, I suggest, is committed to taking grasp of coherence–conferring relations to fall into this category. In order to support this claim, it will be helpful first to review his account of what he calls 'objectual understanding'.[31]

Objectual understanding is understanding of a subject matter. It is marked grammatically by taking an object, as in:

Peter understands arithmetic.
Sophie understands physics.
Schweinsteiger understands the European Union.

Kvanvig takes objectual understanding to be factive. If S understands some subject matter, S's beliefs about that subject matter are true.[32] Additionally, objectual understanding involves subjective justification of a coherentist nature. Subjective justification applies when beliefs are formed or held on the basis of subjective standards for truth or falsity— that is, standards for truth or falsity employed by the subject who holds the beliefs in question. Subjective justification thus involves the subject's being guided by—and grasping—the marks of truth, as defined by her own standards. When a subject S understands a given subject matter, S's subjective justification is of a coherentist nature because the subject matter is characterized by the kinds of relations that coherentists have traditionally focused on in their account of justification (explanatory, probabilistic, logical, and conceptual relations), and because S has grasped these relations.[33] S's grasping the coherence-conferring features of the subject matter gives S subjective justification because it requires S to grasp the marks of truth in that subject matter, as fixed by S's own standards for truth.[34]

Let us now consider understanding from a value-theoretic perspective. Understanding is factive and, for this reason, objectively truth-connected. Given factivity, the subject's beliefs about the relevant subject matter realize the epistemic good of truth. Additionally, understanding involves subjective justification. Subjective justification is subjectively truth-connected since it requires the subject to grasp the marks of truth in the relevant subject matter, as fixed by the subject's standards for truth.[35]

The kinds of features or standings just dealt with derive their value from truth. The factivity of understanding is valuable because it is objectively truth-connected. It delivers the epistemic good of truth. Subjective justification is valuable because it is subjectively truth-connected. However, this is not all there is to say about the epistemic value of understanding. Recall that, in cases of understanding, Kvanvig takes subjective justification to involve the subject's grasping coherence–conferring relations in the relevant subject matter. Grasping such relations is valuable because it contributes to subjective justification—which, as just seen, is valuable due to its subjective truth-connectedness. However, grasp of coherence–conferring relations is also valuable because it organizes and systematizes the subject's thinking about the target subject matter. According to Kvanvig, this organization and systematization of grasped information is valuable because it involves true beliefs (about coherence–conferring relationships—e.g., about what follows from what). However, it also possesses epistemic value that does not derive from truth.[36]

As far as I know, Kvanvig does not offer any elaborate or detailed defence of the idea that grasped coherence–conferring relations possess non-derivative epistemic value. Let me try to do so. Suppose that *S* grasps coherence–conferring relations (i.e. explanatory, probabilistic, logical, and conceptual relations) characteristic of some body of information. *S*'s grasp of these relations organizes and systematizes *S*'s thinking about the body of information. This can happen independently of whether any of *S*'s beliefs are true. Consider, e.g., the following scenario:

Hologram Land:

Bob is in Hologram Land, unbeknownst to him. He has seen two bike holograms and holograms of respectively a boy and a girl. Let these four

holograms be respectively h_1, h_2, h_3, and h_4. Bob is not aware that he is Hologram Land, and so, he is not aware that he has seen bike holograms rather than real bikes and that he has seen holograms of a boy and a girl rather than real kids. Suppose that h_3 is coordinated with h_1 in such a way that it looks like there is a boy riding a bike, and that h_4 is coordinated with h_2 in such a way that it looks like there is a girl riding a bike. Lastly, suppose that Bob thinks that anything with two wheels is an animal; that any animal with a boy riding it is a herbivore; and that any animal with a girl riding it is a carnivore.

In the scenario just described, let us take Bob to believe the following propositions:

(1) h_1 has two wheels.
(2) h_2 has two wheels.
(3) h_3 is a boy.
(4) h_4 is a girl.
(5) h_3 is riding h_1.
(6) h_4 is riding h_2.
(7)   Anything with two wheels is an animal.
(8)   Anything that is an animal with a boy riding it is a herbivore.
(9)   Anything that is an animal with a girl riding it is a carnivore.
(10) h_1 is an animal.
(11) h_2 is an animal.
(12) h_1 is a herbivore.
(13) h_2 is a carnivore.
(14) h_3 is a boy, and h_4 is a girl.
(15) h_3 rides h_1, and h_4 rides h_2.
(16) h_1 is an animal, and h_2 is an animal.
(17) h_1 is a herbivore, and h_2 is a carnivore.

(1)–(17) are characterized by coherence–conferring relations. (10) follows logically from (1) and (7), and (11) follows logically from (2) and (7). Furthermore, (12) follows logically from (3), (5), (8), and (10), and (13) from (4), (6), (9), and (11). The conjunction in (14) follows from (3) and (4), and the conjunctions in (15), (16), and (17) follow from, respectively, (5) and (6), (10) and (11), and (12) and (13).

Given (7), (8), and (9), one might think that some of these logical relations track explanatory relations. (7) states a conceptual connection between **two wheels** and **animal**. In the light of this, one might take (1) and (7) to explain (10): h_1 is an animal *because* it has two wheels and anything with two wheels is an animal. For the same reason, one might take (2) and (7) to explain (11). (8) specifies a conceptual connection between **animal**, **boy rider**, and **herbivore**, while (7) specifies a conceptual connection between **animal**, **girl rider**, and **carnivore**. In the light of this, one might take (12) to be explained by (3), (5), (8), and (10): h_1 is a herbivore *because* it is an animal, a boy is riding it, and anything that is an animal and has a boy riding it is a herbivore. Similarly, one might take (13) to be explained by (4), (6), (9), and (11): h_2 is a carnivore *because* it is an animal, a girl is riding it, and anything that is an animal and has a girl riding it is a carnivore.[37]

Bob's beliefs in the Hologram Land scenario are all false. By design, the scenario is set up in such a way that truth is absent. Bob is looking at bike holograms and kid holograms. This accounts for the falsity of (1)–(6) and (10)–(17). (7)–(9) are false generalizations. Bikes falsify (7). A pig with a boy riding it falsifies (8), and a horse with a girl riding it falsifies (9). (Pigs are omnivores, not herbivores. Horses are herbivores, not carnivores.)

Despite being false, Bob's beliefs jointly exhibit epistemic goodness. This is because they jointly exemplify coherence–conferring relations. The beliefs stand in conceptual, logical, and explanatory relations, as seen above. The various coherence–conferring relations are epistemically valuable because they systematize and organize the propositions believed. Consider two sets of beliefs $B(p_1)$, … , $B(p_n)$ and $B(p_1^*)$, … , $B(p_n^*)$. Suppose that they are on a par in terms of their semantic status (all false, all true, or matching distribution of, respectively, false and true beliefs). However, assume that $B(p_1)$, … , $B(p_n)$ exhibit a great number of coherence–conferring relations, while $B(p_1^*)$, … , $B(p_n^*)$ are completely disconnected or unrelated beliefs. In that case, $B(p_1)$, … , $B(p_n)$ would seem epistemically better than $B(p_1^*)$, … , $B(p_n*)$. The former set of beliefs is organized and systematized by coherence–conferring relations, while, due to the absence of such relations, the latter set is not organized and not systematized.

What is more, the epistemic value of beliefs being organized or systematized by coherence–conferring relations does not derive from truth. Even if a set of beliefs involves just false beliefs, its members can stand in coherence–conferring relations. The Hologram Land scenario serves to illustrate this point. By design, all of Bob's beliefs are false, but they stand in coherence–conferring relations. They are organized and systematized by the conceptual, logical, and explanatory relations that obtain between them.

The considerations offered in the second half of this section support the idea that there is a non-derivative epistemic good distinct from truth. In presenting these considerations, I have taken my cue from Kvanvig's work on understanding. I believe that Kvanvig is committed to pure epistemic pluralism, i.e. that there are several non-derivative epistemic goods. He defends the idea that truth is intrinsically valuable, i.e. that truth is not valuable due to a relation it bears to anything else.[38] This would seem to imply that it does not derive its value from some other good, and so, that it is a non-derivative epistemic good. I also take Kvanvig to endorse the idea that grasp of coherence-conferring relations is non-derivatively epistemically valuable and have offered considerations in favour of this idea through an illustrative case.

One issue remains. Kvanvig speaks of *grasped* coherence–conferring relations rather than coherence–conferring relations themselves. Presumably this is because he wants to identify a good of subject's *beliefs*. Suppose that a subject S believes propositions $p_1, \ldots, p_n$, and that these propositions bear many coherence–conferring relations to each other. However, suppose that S does not grasp any of these relations. If so, the coherence–conferring relations have not been operative in S's cognitive life. In this sense, the relations do not serve to systematize and organize $\underline{S}$'s beliefs or thinking about $p_1, \ldots, p_n$. On the other hand, if S grasps the coherence–conferring relations, they do serve to organize and systematize S's beliefs or thinking about $p_1, \ldots, p_n$.

6   Collapse Arguments

Pure epistemic pluralism is a view worth exploring. It offers a conception of the structure of the domain of epistemic goods very different from the conceptions that underwrite monist and unitarian or moderate pluralist views. However, pure epistemic pluralism is not a widely held view. Epistemology is strongly dominated by veritic monism and unitarianism. In the next two sections, I do some preliminary ground clearing for a fuller exploration of pure epistemic pluralism. I do so by addressing two collapse arguments—what I call respectively the 'Instability Challenge' and 'Unity Challenge'. The Instability Challenge and Unity Challenge purport to show that epistemic pluralism is inherently unstable—that, upon reflection, it collapses into monism. Now, while I myself am sympathetic to pure epistemic pluralism, I keep an open mind as to whether it is ultimately tenable. However, I do not think that the collapse arguments give us grounds for scepticism.

Alethic Collapse Arguments

The Instability Challenge and Unity Challenge are prominent arguments in the literature on truth pluralism. I present them in the order mentioned.

The Instability Challenge:
According to truth pluralism, there is a plurality of ways of being true, T_1, … , T_n. Each of T_1, … , T_n is a way of being true in the sense that they are the way in which propositions from certain (but not all) domains are true. However, consider now the following disjunctive way of being true T_G (where D_i ($1 \leq i \leq n$) is the domain to which T_i applies):

$$(\text{T}_\text{G}) \quad (\forall p)(T_G(p) \leftrightarrow ((T_1(p) \wedge D_1(p)) \vee \ldots \vee (T_n(p) \wedge D_n)))$$

Since T_G is characterized by appeal to all of the ways of being true endorsed by the pluralist, T_G applies across all (truth-apt) domains. Furthermore, T_G is extensionally adequately: it applies exactly to the propositions that the pluralist countenances as true—and, given the

characterization of T_G, it does so necessarily. However, this means that T_G is a generic way of being true. There is no need to endorse a plurality of ways of being true. Being true is simply being T_G. Truth pluralism is thus an inherently unstable position. Straightforward reasoning shows that it collapses into monism.[39]

The Unity Challenge:
Assume with the truth pluralist that there are several ways of being true $T_1, \ldots, T_n$. Now ask what makes $T_1, \ldots, T_n$ ways of being *true*? Against the background of this question, the Unity Challenge can be presented in three steps.

Step 1—The unity requirement:

Properly classifying $T_1, \ldots, T_n$ as ways of being true—rather than ways of being something else—requires a significant degree of unity between $T_1, \ldots, T_n$. That is, $T_1, \ldots, T_n$ must share certain core features that unify them and make it the case that they are all ways of being *true*.

Step 2—Unity-underwriting features:

The core features that unify $T_1, \ldots, T_n$ are captured by certain core principles. These core principles might be taken to include (but not necessarily be limited to) the following:

Equivalence (E): for all p, it is true that p if and only if p.
  Contrast (C): for all p, it is possible for p to be justified but false, and for p to be true but not justified.
Preservation (P): for all p, the truth-aptitude of p (and any other truth-apt proposition) is preserved under basic logical operations (negation, disjunction, conjunction, etc.).[40]

Step 3—Collapse:

The unity-underwriting features captured by (E), (C), and (P) can be used to characterize a property T^*, as follows:

(T^*) T^* is the property such that, necessarily, for all p:

   (i) p is T^* if and only if p,

(ii)   it is possible for p to be justified but not T^* and for p to be T^* but not justified, and

(iii)   the T^*-aptitude of p (possibly in combination with other T^*-apt propositions) is preserved under basic logical operations (negation, disjunction, conjunction, etc.)

Since (T^*) captures the core features shared by the different ways of being true, (T^*) provides a specification of what really matters to truth. Indeed, (T^*) captures exactly what it means to be true. However, this means that T^*—the property characterized by (T^*)—simply <u>is</u> the property of being true. Truth pluralism thus collapses into truth monism.

Epistemic Collapse Arguments

The arguments just presented target a specific form of pluralism: pluralism about truth. However, the reasoning behind the arguments generalizes. The Instability Challenge relies on the fact that, given a plurality of properties, it is possible to characterize a disjunctive property that (necessarily) applies to exactly the things that those properties apply to. This applies in general, not just when we are considering supposed ways of being true. The Unity Challenge relies on the idea that unity is required in order for $F_1, \ldots, F_n$ properly to be labelled 'ways of being F'. But this should apply whether $F_1, \ldots, F_n$ are ways of being true or ways of being epistemically good. In the light of these two comments—and given the prominence of the Instability Challenge and Unity Challenge in the truth debate—I want to transpose the two challenges to epistemology and address them in that setting.[41]

The Instability Challenge:
According to epistemic pluralism, there is a plurality of epistemic goods, $G_1, \ldots, G_n$. Each of $G_1, \ldots, G_n$ is an epistemic good in the sense that it confers a positive epistemic standing on bearers of epistemic value ('ev-bearers'). However, consider now the following disjunctively characterized good G_G (where x ranges over ev-bearers):

$$(GG) \quad (\forall x)(G_G(x) \leftrightarrow (G_1(x) \lor \ldots \lor G_n(x)))$$

G_G is extensionally adequate from the perspective of the pluralist. It applies exactly to the ev-bearers that enjoy some positive standing by her lights—and it does so necessarily. This means that there is a single, extensionally adequate epistemic good. However, then there is no need to endorse a plurality of epistemic goods. Being epistemically good simply is being G_G. Epistemic pluralism is an inherently unstable position. Straightforward reasoning shows that it collapses into epistemic monism.

The Unity Challenge:
Assume epistemic pluralism, i.e. that there are several epistemic goods $G_1, \ldots, G_n$. What makes $G_1, \ldots, G_n$ epistemic goods? Against the background of this question, the epistemic Unity Challenge can be presented in three steps.

Step 1—The unity requirement:

Properly classifying $G_1, \ldots, G_n$ as epistemic goods—rather than some other kind of good—requires a significant degree of unity between $G, \ldots, G$. That is, $G_1, \ldots, G_n$ must share certain core features that unify them and make it the case that they are all *epistemic* goods.

Step 2—Unity-underwriting features:

Let the features that unify $G_1, \ldots, G_n$ be $F_1, \ldots, F_m$. $F_1, \ldots, F_m$ are captured by core principles $P_1, \ldots, P_m$. Each principle P_i $(1 \leq i \leq m)$ is of the form $(\forall x)\Phi(x)$ where x ranges over ev-bearers and $\Phi(x)$ specifies some necessary or sufficient condition for G-ness.

Step 3—Collapse:
Define G^* as follows:

(G^*) G^* is the property such that, necessarily, for all x: $F_1(x), \ldots, F_m(x)$,

where, as before, x ranges over ev-bearers, and '$F_1(x), \ldots F_m(x)$' says that x possesses the unity-underwriting features $F_1, \ldots, F_m$. (G^*) captures the unity-underwriting features of the pluralist's epistemic goods, $G_1, \ldots, G_n$. This means that (G^*) captures what really matters to being epistemically good. Indeed, it captures exactly what

it means to be epistemically good. However, this means that G^* <u>is</u> the property of being epistemically good. Epistemic pluralism collapses into epistemic monism.

7   Comments and Responses

In this section, I respond to the Instability Challenge and the Unity Challenge. However, before I do so, let me offer two comments that will inform or serve as background to the responses.

Comments

Comment 1: Impasse?
    In the epistemic case, the unity-underwriting features are treated entirely in the abstract. There is no specification of particular features that might be thought to unify the non-derivative epistemic goods G_1, ... , G_n. One proposal would be to say that G_1, ... , G_n are unified by all of them being truth-connected (where some $G_{1 \leq i \leq n}$ would be so in the sense of being truth itself).
    Many might be tempted by this answer, the reason being that the following conception of the boundaries of the epistemic realm is widely held:
    The Truth Conception of the Epistemic (TCE):

Truth is the sole epistemic good, and what it means for something—a species of warrant, rationality, reason, or desideratum, say—to be epistemic is for it to be truth-connected.

Many—or most—epistemologists either explicitly or implicitly buy into (TCE). This is reflected by the prominence of veritic unitarianism. Advocates of (TCE) tend categorically to dismiss the idea that epistemic standings or goods that are not truth-related can be epistemic in nature. Given (TCE) it is clear why G_1, ... , G_n must be unified by all being truth-connected. If they were not, they would not all be *epistemic* goods.

While (TCE) may be widely endorsed by epistemologists (e.g., by the kinds of moderate pluralists considered in Sects. 2 and 3), appealing to (TCE) in the present context would seem to beg the question against the pluralist. If appeal was made to (TCE), it might seem natural to say that $G_1, \ldots, G_n$ were derivatively epistemically good because they would derive their value from being connected to truth. However, this would run counter to the idea that all of $G_1, \ldots, G_n$ are *non-derivative* epistemic goods—which is precisely the distinguishing feature of pure epistemic pluralism.

It would thus seem that someone attracted by pure epistemic pluralism must reject (TCE). This raises the question whether there is a way to delineate the boundaries of the epistemic realm without begging any questions against the monist or the pluralist about non-derivative epistemic value. (I count unitarians as belonging to the monist camp.) If there is not, the discussion might well lead to an impasse. For present purposes, I believe that worries about a potential impasse can be set aside. To see this, recall that the task presently at hand is defensive in nature. The stability of epistemic pluralism has been questioned on the basis of the Instability Challenge and the Unity Challenge. The task at hand is to defend pluralism against these challenges. Since the charge is that pluralism is inherently unstable, it would seem legitimate simply to assume a pluralist perspective. Thus, if the pluralist rejects the idea that the non-derivative goods $G_1, \ldots, G_n$ are unified by all being truth-connected, they can simply help themselves to this assumption—even if the challenger holds a version of veritic monism or unitarianism.

Comment 2: Two Kinds of Pure Epistemic Pluralism

Earlier I suggested that pure epistemic pluralism is the purest form of epistemic pluralism. Strictly speaking, it is worth distinguishing between two ways of being a pure epistemic pluralist. We can get at these two ways of being a pure epistemic pluralist by considering two distinct kinds of response to the Unity Challenge. One kind of response is to insist that there is no genuinely unifying feature shared by all of the non-derivative epistemic goods $G_1, \ldots, G_n$. In calling them all 'epistemic goods' we are merely applying a single expression to a rather varied bunch. Another kind of response grants that the pluralist

range of non-derivative epistemic goods shares at least one unifying feature. However, this is paired with an argument to the effect that the presence of unity-underwriting features does not suffice to show that pure epistemic pluralism collapses into epistemic monism. The first line of response is much less concessive than the second. Accordingly, the kind of pure epistemic pluralism that the first response goes with is more radical than the kind of pure epistemic pluralism that goes with the second. In this section, I explore the second kind of response to the Unity Challenge. Between the two kinds of responses, this type of response possesses greater dialectical force because it grants the challenger a key component of her set-up.

Response to the Instability Challenge

Consider the following two claims:

(C1) $G_1, \ldots, G_n$ are more fundamental than G_G.
(C2) G_G is not a non-derivative epistemic good. Rather, it is a derivative good.

I am going to argue in favour of (C1) and (C2). This puts me in a position to defuse the Instability Challenge.

Let me start with (C1). Recall (GG):

$$(GG) \quad (\forall x)(G_G(x) \leftrightarrow (G_1(x) \vee \ldots \vee G_n(x)))$$

Here is an argument in favour of (C1): the bi-conditional in (GG) is to be read with priority from right to left. Instances of G_G are grounded by instances of $G_1, \ldots, G_n$. On the other hand, since grounding is asymmetric, we do *not* have any instance of G_G grounding any instance of any of $G_1, \ldots, G_n$.[42] The pluralist's non-derivative epistemic goods ground the supposed generic epistemic good, but not vice versa. Whenever G_G is instantiated, this is *because*—or *in virtue of*—some G_i's being instantiated (where $1 \leq i \leq n$). However, it is never the case that G_i is instantiated because—or in virtue of—G_G's being instantiated (again, $1 \leq i \leq n$). In this sense, $G_1, \ldots, G_n$ are more fundamental than G_G.

Let us now turn to (C2). The argument for (C1) gives us that G_1, ... , G_n are more fundamental than G_G in the sense of their grounding G_G and not being grounded by G_G. Now, return to (GG). It provides a crucial piece of extra information: G_G can *only* be grounded by G_1, ... , G_n. This supports the conclusion that G_G is nothing over and above G_1, ... , G_n. It is exactly the disjunction of these epistemic goods. Nothing more, nothing less. G_G is thus a derivative *property* at least in the sense that instances of G_G are *entirely metaphysically dependent* on G_1, ... , G_n. There is no way for G_G to be instantiated other than by one of G_1, ... , G_n's being instantiated.[43]

How does G_G's being a derivative property—in the metaphysical sense specified—impact what happens at the level of value? We seem to have at least this much: the epistemically good-making features of G_G do not get instantiated unless one of G_1, ... , G_n gets instantiated. Now, bear this in mind and also that G_1, ... , G_n are epistemic goods. Putting these two things together, the most natural candidate for an epistemically good-making feature of G_G would seem to be that its instantiation means that epistemically good-making features are instantiated *because* one of G_1, ... , G_n is instantiated. Taking on board this natural proposal gives G_G an epistemically good-making feature. The crucial thing to note, however, is that it also makes G_G a *derivative* epistemic good. It does so because the epistemic goodness of G_G derives from its instantiating something else that is an epistemic good.

I now am in a position to defuse the Instability Challenge. There are two points I would like to make. One point is metaphysical, the other axiological. The metaphysical point underwrites the axiological one.

First, since G_G is grounded by G_1, ... , G_n (and only by them), G_1, ... , G_n can reasonably be said to enjoy *metaphysical priority* over G_G. This metaphysical priority gives us at least that the supposed epistemic goodness of G_G is only ever realized because one of G_1, ... , G_n is instantiated. In this sense, the pluralist's epistemic goods serve as a metaphysical basis of the supposed generic epistemic good. I take this point to help defuse the Instability Challenge. It makes it clear that G_1, ... , G_n are indispensable. Even if G_G cannot be resisted, it must be granted that G_1, ... , G_n are important. They are needed as the metaphysical basis of

G_G. The only way to maintain that G_G has instances is to grant that G_1, … , G_n likewise have instances.[44]

Second, given the first point about metaphysical priority, our above considerations show that G_1, … , G_n enjoy *axiological priority* over G_G. G_1, … , G_n are not only more fundamental than G_G in a metaphysical sense. They are likewise more fundamental than the supposed generic good G_G in an axiological sense. This is because G_G is a derivative epistemic good. The epistemic goodness of any instance of G_G derives from the epistemic goodness of G_1, …, Gn. This defuses the Instability Challenge. For, the question we are considering is whether the challenge undermines pure epistemic pluralism. Pure epistemic pluralism is the view that there are several *non-derivative* goods. However, G_G—given its status as a *derivative* good—does not pose a threat to this thesis.

Response to the Unity Challenge

Let us now turn to the Unity Challenge. Recall that the supposed problematic epistemic good is characterized as follows:

(G*)   G* is the property such that, necessarily, for all x: $F_1(x)$, … , $F_m(x)$,

where x ranges over ev-bearers, and $F_1, \ldots, F_m$ are the features that are meant to unify the pluralist's epistemic goods G_1, … , G_n.

My goal in this section is to defuse the Unity Challenge. I do so by investigating the following basic question:

(Q)   What goods satisfy (G*)?

The Unity Challenge is mostly silent on this matter. I say 'mostly silent' because, while the Unity Challenge comes with no specification of a good satisfying (G*), the statement of the challenge does reveal a crucial presumption. The advocate of the Unity Challenge uses the definite article—speaking of *the* good that necessarily possesses the relevant features. Proper use of the definite article requires uniqueness. However,

this is a problematic presumption. Further reflection on (G^*) undermines it.

Let us distinguish between two candidate kinds of goods for satisfying (G^*):

(a)   Goods that necessarily have features $F_1, \ldots, F_m$ (as required by (G^*)) *and* are necessarily co-extensional with $\cup\{\text{Ext}(G_1), \ldots \text{Ext}(G_n)\}$, i.e. the union of the extensions of the pluralist goods $G_1, \ldots G_n$.

(b)   Goods that necessarily have features $F_1, \ldots, F_n$ (as required by (G^*)) but need not be necessarily co-extensional with $\cup\{\text{Ext}(G_1), \ldots \text{Ext}(G_n)\}$, i.e. the union of the extensions of the pluralist goods $G_1, \ldots G_n$.

I argue that, for both (a) and (b), the presumption of uniqueness is undermined or more work must be done to support it.

Let us consider (a). Goods that conform to (a) have the features that the 'unity challenger' takes to capture the true characteristics of epistemic goods. Dialectically one might think that a single, unique good that conformed to (a) would have some force. Given the necessary co-extensionality requirement, any good of type (a) will be extensionally adequate from the point of view of the pure epistemic pluralist. Necessarily, any such good applies *exactly* to the things that the pure epistemic pluralist takes to possess some (non-derivative) epistemic good. However, what are candidate goods of type (a)? One good immediately comes to mind: G_G, the disjunctive good of being G_1 or … or G_n. Given its characterization, this good is necessarily co-extensional with $\cup\{\text{Ext}(G_1), \ldots \text{Ext}(G_n)\}$. Now, does the disjunctive good possess the unity-underwriting features $F_1, \ldots, F_m$ as a matter of necessity? I believe that it does and make a case for this in the Appendix. (The material contained in the Appendix is of independent interest. It appears as an appendix rather than as a part of the main text in order to prevent us from getting caught up in details that would blur what is really at stake with G_G.) The question is whether G_G poses a threat to pure epistemic pluralism, assuming that it conforms to (a).

There is a serious question as to whether uniqueness can be granted—whether, indeed, G_G can be shown to be the only type (a) good. Let us leave this issue aside for now and grant uniqueness, i.e. suppose that G_G is the only type (a) good. Even so, pure epistemic pluralism is not in trouble. This is because of the nature of G_G, the supposed unique type (a) good. Recall Sect. 7.2. In that section, it was argued that it is implausible to think that G_G is a non-derivative good. Instances of G_G are good because they are instances of $G_1, \ldots, G_n$. For this reason, G_G is a derivative rather than a non-derivative epistemic good. Since pure epistemic pluralism is the thesis that there are several non-derivative epistemic goods, G_G does not pose a threat to the view.

Might there be other candidates for type (a) goods? In order to pose a threat to pure epistemic pluralism, there would have to be just one such good, *and* it would have to be a non-derivative good. Furthermore, in order for this kind of good to pose a threat to pure epistemic pluralism, the unity challenger would have to make certain assumptions about the metaphysics of properties. In particular, it would have to be assumed that properties are individuated at least in part intensionally. If they were individuated purely extensionally, any goods satisfying (a) would be identical to one another. This is so because of the requirement of necessary co-extensionality. Now, assuming that goods are at least partly intensionally individuated, it would have to be shown that there is a unique good that is necessarily co-extensional with $\cup\{\mathrm{Ext}(G_1), \ldots, \mathrm{Ext}(G_n)\}$, necessarily has the unity-underwriting features $F_1, \ldots, F_m$, and, unlike G_G, is a non-derivative good. In this context, I take it that it is fair to say that the onus is on the unity challenger to show that there is an epistemic good that meets all of these conditions. After all, it is the unity challenger who claims that pure epistemic pluralism is inherently unstable.

Let us now move on to consider (b). Going along with (b) the unity challenger must show that there is a unique good that necessarily has the unity-underwriting features $F_1, \ldots, F_m$ and qualifies as a non-derivative good. There is no requirement to the effect that this good must be co-extensional with $\cup\{\mathrm{Ext}(G_1), \ldots, \mathrm{Ext}(G_n)\}$. However, in that case, $G_1, \ldots, G_n$ would all seem to be candidates. They certainly all have features $F_1, \ldots, F_m$, as the very idea behind the Unity Challenge

is to focus on $F_1, \ldots, F_m$—the *shared* features that unify $G_1, \ldots, G_n$. Does each of $G_1, \ldots, G_n$ have these features necessarily? Arguably so. Otherwise $F_1, \ldots, F_m$ would not be truly unifying. The features would not unify instances of $G_1, \ldots, G_n$ within and across possible worlds unless $G_1, \ldots, G_n$ possessed $F_1, \ldots, F_m$ necessarily. In sum, if the necessary co-extensionality requirement is dropped, the unity challenger would have to endorse all of $G_1, \ldots, G_n$ as goods. However, this commitment does not disagree with pure epistemic pluralism—indeed, it <u>is</u> pure epistemic pluralism!

I conclude that the Unity Challenge—whether understood against the background of (a) or (b)—does not pose a threat to pure epistemic pluralism.

8   Concluding Remarks

I hope to have accomplished six things in this chapter. First, I hope to have shown that epistemic pluralism is endorsed by several prominent epistemologists. In particular, I traced a pluralist strand in the works of several prominent epistemologists. Burge is a pluralist about epistemic warrant, Goldman about epistemic justification, and Alston about epistemic desiderata. Second, I suggested that the principal motivation for the pluralism embraced by these three prominent epistemologists is the same. Each form of pluralism is motivated by appeal to the scope problem. Unless epistemic pluralism is taken on board, it is not possible to accommodate all cases that seem to be cases of warrant (Burge), justification (Goldman), or epistemic desiderata (Alston). Third, while attributing a species of pluralism to each of Burge, Goldman, and Alston, I suggested that it is somewhat moderate in nature. This is because they are all combined with veritic unitarianism. There is a plurality of epistemic goods (different species of warrant, etc.), but they are all unified by a single, overarching good—truth. I used this observation to formulate what, in my view, is a much purer form of epistemic pluralism: pluralism about non-derivative epistemic goods.

Fourth, I presented DePaul and Kvanvig as two examples of epistemologists who seem to be committed to pure epistemic pluralism. In connection with Kvanvig, I spelled out in some detail why one might think that grasp of coherence–conferring relations is non-derivatively epistemically good. Fifth, I presented two arguments targeting epistemic pluralism: the Instability Challenge and the Unity Challenge. I called them 'collapse arguments' because they both purport to show that epistemic pluralism collapses into epistemic monism. Sixth, some ground-clearing work was done on behalf of pure epistemic pluralists. I am sympathetic to the view and think it deserves to be explored in detail. If either of the collapse arguments was compelling, the prospects of pure epistemic pluralism would seem rather bleak. The view would not even get off the ground. I hope to have cleared the ground for further investigation of the view by responding to the two arguments.

I hope that this chapter will help advance the discussion of epistemic pluralism. In particular, I hope that it offers what might be considered the first steps towards a comprehensive investigation of pure epistemic pluralism. In framing the chapter, I have drawn parallels between the epistemological literature and the literature on truth pluralism. I suggested that epistemic pluralists motivate their view in a way that is structurally similar to the way in which truth pluralists motivate their view. Both kinds of pluralists appeal to the superior scope of pluralism compared to monist rivals. I also suggested that certain fundamental challenges faced by truth pluralists can be transposed to the epistemic realm. However, fortunately the challenges can be defused—or at least so I argued. Hopefully, there will be cross-fertilization between pluralism debates from different areas of philosophy—truth, logic, epistemology, ontology, morals, and other areas. In order to achieve cross-fertilization, it is important to focus on differences as well as similarities. Only by doing this, will it become clear what models of pluralism are workable within different areas of philosophy. This, in turn, will impact the kinds of motivations and challenges pluralists face within a given area of philosophy.[45]

Notes

1. For different kinds of truth pluralism, see Cotnoir (2013); Edwards (2011, 2012, 2013); Lynch (2001, 2004, 2009, 2013); Pedersen (2006, 2010, 2012a, b); Wright (1992, 2013).
2. For this kind of logical pluralism, see Beall and Restall (2006). For a different kind of logical pluralism—one tied to different kinds of structures—see Shapiro (2014).
3. Burge (2013, pp. 3–4).
4. The distinction between justification and entitlement—and the pluralism that goes with it—is a core component of Burge's epistemological work. Sample references include Burge (1993, 1998, 2003, forthcoming).
5. Goldman (1988, p. 51).
6. Goldman (1988, Sect. II–III).
7. See Goldman (1988, p. 56) for a characterization of weak justification. Let me emphasize that clauses (v) and (vi) above are *my* additions. Goldman's own official characterization does not include these clauses. However, they ought to be included. It is pretty clear that clause (vi) is satisfied in the kind of case to which Goldman wants weak justification to apply (see p. 52). It would also seem that clause (v) should be included if, as on Goldman's view, reliability is to be the core epistemological notion. Both clauses (v) and (vi) do crucial work to support the idea that the subject's belief in the target range of cases should enjoy a standing worthy of the label 'justification'. Clause (i) states that, as a matter of fact, M is an unreliable process or method. This clearly cannot confer a positive epistemic standing on M, not even from the subject's own perspective. Clauses (ii)–(iv) are all negative clauses, requiring the subject not to have any reason to think M unreliable (which would amount to the subject's having a defeater for thinking M reliable). None of clauses (ii)–(iv) can confer a positive standing upon M either. However, clauses (v) and (vi) achieve this. Clause (v) specifies that the subject believes M to be reliable, while clause (vi) ensures that, from the subject's perspective, the belief in the reliability of M is well supported. Again, I take it to be clear that clause (v) is satisfied in the kind of case that Goldman uses to motivate the introduction of weak justification, and that clause (vi) should be incorporated because weak justification is meant to be embedded within a broadly reliabilist

framework. Also, as just suggested, clauses (v) and (vi) do crucial work in terms of account for why, within this kind of framework, weak justification is worthy of the label 'justification'.

8. Alston (1993, p. 531), (2005, p. 22). Some are desiderata discussed by Alston apply to individual beliefs, others for systems of beliefs. Here, I restrict attention to the former kind.

9. Alston (2005, pp. 43–44). In Sect. 4, I spell out what is meant by 'Plantiga-warranted.'

10. Alston (2005, pp. 43–44).

11. Alston (1993, p. 531), (2005, p. 22).

12. Lynch (2004), (2009); Edwards (2011, 2012); Pedersen (2006, 2010); Wright (2013).

13. Burge (2003, pp. 528–529).

14. Goldman (1988, pp. 51–52).

15. Goldman (1988, p. 52).

16. Alston (2005, pp. 53–57).

17. Not much work has been published on comparisons between truth pluralism and epistemic pluralism. However, in thinking about similarities between pluralist views within these two areas, I have benefitted from discussion with a number of people, including Luca Zanetti, Michael Lynch, Filippo Ferrari, Sebastiano Moruzzi, and Annalisa Coliva. There is a consensus on the point made in this section, i.e. that epistemic pluralism can be (and is) motivated by appeal to the scope problem.

18. My use of the term 'unitarianism' is inspired by Goldman (2001). However, let me note that Goldman uses 'unitarianism' to denote the epistemic virtue analogue of unitarianism as defined here.

19. Alston (1989, 1993, 2005); Burge (2003, forthcoming); Goldman (2001). Many others endorse the view, either explicitly or implicitly. See, e.g., BonJour (1985, pp. 7–8); Moser (1985, p. 4); and Pritchard (2014). For discussion of various problems with veritic unitarianism (and monism)—and suggested solutions—see David (2001, 2005).

20. Burge (2003, pp. 505–506).

21. Goldman (2001), (1999: Chap. 1–3).

22. See Goldman (2001, Sect. 7) for a similar point.

23. I do not take a belief's being subjectively truth-connected to involve the subject's explicitly believing or otherwise conceptualizing the belief as such (or as having some other feature that ensures a truth-connection). It suffices that the subject acts as if the belief is true

and would deem it likely to be true (or enjoy some other truth-connected status) if asked explicitly.

24. Alston (2005, p. 29).
25. Alston (2005, pp. 89–90).
26. Plantinga (1993, pp. 46–47).
27. Alston (2005, pp. 148–150). Alston also considers deontological features and intellectual virtues as candidate epistemic desiderata, but ultimately includes no feature of either kind in the list of epistemic desiderata. See Chap. 4 and Alston (2005, pp. 151–162) for details.
28. Again, as suggested in connection with Goldman, I do not think that subjective truth-connectedness should be thought to require, on the part of the subject, an explicit belief or other conceptualization of the relevant belief as being truth-connected. It suffices that the subject acts as if the belief is true and would deem it likely to be true (or enjoy some other truth-connected status) if asked explicitly.
29. The second path is the one taken by Alston himself. See Alston (2005, pp. 43–45, 50). See Peels (2010) for a discussion of Alston's epistemic pluralism and the various desiderata considered by Alston.
30. DePaul (2001). See DePaul (2004) and Zagzebski (2004) for arguments against specific forms of veritic monism. DePaul (2004) targets Goldman's veritic epistemic consequentialism. Zagzebski (2004) targets veritic monism as it features in Sosa's virtue reliabilism.
31. Kvanvig distinguishes objectual understanding from propositional understanding which, grammatically, is marked by a that-clause. Below, I focus exclusively on objectual understanding and, for this reason, allow myself to leave out the marker 'objectual'.
32. Strictly speaking, this needs to be qualified. Understanding a subject matter is compatible with having *some* false beliefs about it. See Kvanvig (2003, pp. 201–202) for discussion.
33. Kvanvig (2003, pp. 192–193, 197–198, 200–202).
34. Kvanvig (2003, p. 202).
35. Kvanvig takes subjective justification to involve the adoption of 'intentional means to the goal of truth'. According to him, intentional means to a goal possess a kind of value different from instrumental value. Adopting intentional means to a given goal does not necessarily increase the likelihood of reaching the goal. However, in order for a means to be instrumentally valuable, it needs to raise the likelihood of reaching the relevant goal. See Kvanvig (2003, pp. 63–64).

36. Kvanvig ([2003], p. 202).
37. Some might think that the occurrence of singular terms for the holograms in the content of beliefs is implausible or odd. If so, the case can be rewritten in terms of demonstratives.
38. Kvanvig ([2003], p. 42).
39. This argument is given by Tappolet ([2000]). See Pedersen ([2006]) and ([2010]) for pluralist responses.
40. (E), (C), and (P) are core principles concerning truth favoured by Wright ([1992]). The full list of core principles in Wright is longer than the three-membered list given above (and Wright uses the label 'Embedding' rather than 'Preservation'). Also, other authors give different lists of core principles (e.g., Lynch ([2009])). I leave this issue aside for present purposes. What does the crucial work in the Unity Challenge is the assumption that the different ways of being true are unified by certain features.
41. The literature on pluralism is not as comprehensive in epistemology as in truth studies. However, for what it is worth, in discussion the idea of epistemic pluralism often prompts someone to present a version of the Instability Challenge or Unity Challenge.
42. A relation R is asymmetric if and only if, for all x and y, if $R(x, y)$, then $\neg R(y, x)$.
43. The claim that G_G is nothing over and above $G_1, \dots, G_n$ is made within the dialectical context that the Instability Challenge is formulated. Recall that the idea behind the Instability Challenge is that pure epistemic pluralism is inherently unstable because straightforward reasoning undermines it. Here, 'straightforward reasoning' means using a disjunctive characterization to introduce G_G—a move that it seems difficult to resist. Surely, one can introduce a new notion as the disjunction of existing ones. However, the reason why this move seems so difficult to resist is precisely because the newly characterized notion is not supposed to be anything over and above the disjunction of the existing ones. One might claim, say, that G_G must possess some feature that unifies all of $G, \dots, G_{1n}$. However, if this claim is advanced, we move past the point of 'straightforward reasoning' by introducing a new notion by characterizing it merely as the disjunction of existing ones. Indeed, we move into the territory of the Unity Challenge—which I respond to in due course.
44. See Pedersen ([2010]) for this kind of argument in the case of truth.

45. Again, I wish to thank Luca Zanetti for fruitful and illuminating conversations about epistemic pluralism, including differences and similarities with the truth pluralism. Zanetti (2014) contains much material of interest.

46. $\left(\text{GE}^*_{\text{G1}}\right)$ and $\left(\text{GE}^*_{\text{G2}}\right)$ have to be conditionals rather than bi-conditionals. By (GE_{G1}) and (GE_{G2}), both G_1 and G_2 are goals of cognition. So, making $\left(\text{GE}^*_{\text{G1}}\right)$ a bi-conditional would incorrectly predict that instances of G_2 are not goals of enquiry, while making $\left(\text{GE}^*_{\text{G2}}\right)$ a bi-conditional would do the same for instances of G_1.

Appendix: G_G

$G_1, \ldots, G_n$ are the non-derivative goods endorsed by the pure epistemic pluralist. $F_1, \ldots, F_m$ are the features that unify $G_1, \ldots, G_n$ according to the Unity Challenge. G_G is the property of being G_1 or … or G_n. This appendix argues that G_G possesses the features $F_1, \ldots, F_m$ as a matter of necessity.

In giving my argument, I assume that the pluralist endorses two epistemic goods, G_1 and G_2, and that these epistemic goods are unified by (and only by) the feature of being a goal of enquiry:

(GE_{G1}) G_1 is a goal of enquiry.
(GE_{G2}) G_2 is a goal of enquiry.

I make the assumptions just stated for expository purposes only. Nothing in my argument hangs on them.

Now, (GE_{G1}) and (GE_{G2}) capture a feature of epistemic goods rather than ev-bearers. However, the 'goal feature' can be transposed to the level of ev-bearers as follows:

$\left(\text{GE}^*_{G1}\right)$ For all x, if $G_1(x)$, then x is a goal of enquiry.
$\left(\text{GE}^*_{G2}\right)$ For all x, if $G_2(x)$, then x is a goal of enquiry.

where x ranges over ev-bearers.[46] Combining (G^*) (from Sect. 6) with the idea that the unifying feature of epistemic goods is to be a goal of enquiry yields:

(G^{*GE}) G^* is the property such that, necessarily, for all x: if $G^*(x)$, then x is a goal of enquiry.

Here, G^* is the supposed unique epistemic good. Above, I claimed that G_G—the property of being G_1 or … or G_n—has $F_1, \ldots, F_n$ necessarily. Assuming (G^{*GE}) supporting this claim amounts to showing the following:

(G^{*GE-GG}) Necessarily, for all x: if x is G_1 or x is G_2, then x is a goal of enquiry.

Let offer support for (G^{*GE-GG}). In order to do so, I rely on strengthened versions $\left(GE^*_{G1}\right)$ and $\left(GE^*_{G2}\right)$:

$\left(GE^*_{G1\square}\right)$ Necessarily, for all x, if $G_1(x)$, then x is a goal of enquiry.
$\left(GE^*_{G2\square}\right)$ Necessarily, for all x, if $G_2(x)$, then x is a goal of enquiry.

(G^{*GE-GG}) is a consequence of $\left(GE^*_{G1\square}\right)$ and $\left(GE^*_{G2\square}\right)$.

(GE^*_{G1}) and (GE^*_{G2}) can be strengthened to $\left(GE^*_{G1\square}\right)$ and $\left(GE^*_{G2\square}\right)$ because the status of G_1 and G_2 as goals of enquiry is not a mere coincidence. It is modally robust. This modal robustness is explained by the modal robustness of the connection between the status of G_1 and G_2 as goods and their status as goals of enquiry together with the modal robustness of G_1 and G_2's status as epistemic goods. Turn first to the modal robustness of the good–goal link. G_1 and G_2 are goals of enquiry *because* they are epistemic goods. This means that, in any possible case, if something is G_1, then it is a goal of enquiry. Similarly for G_2. So we have:

$(G\text{-}GE_\square)$ Necessarily, for all x, if x is epistemically good, then x is a goal of enquiry.

Consider now the status of G_1 and G_2 as epistemic goods. They do not enjoy this status as a mere coincidence. The status of G_1 and G_2 as epistemic goods is modally robust—a matter of necessity—because they are non-derivative epistemic goods. The epistemic goodness of G_1 and

G_2 is not due to their bearing some relation to some other good. The goodness of G_1 and G_2 resides solely in G_1 and G_2 themselves. Thus, whenever G_1 is instantiated, this by itself—or on its own—is epistemically good. Similarly for G_2. Here, I mean to use 'whenever' to indicate generality along two dimensions. The first dimension is intra-worldly: within a given possible world, the instantiation of G_1 and G_2 is always epistemically good. The other dimension is inter-worldly: the instantiation of G_1 and G_2 is epistemically good across all possible worlds. Taking on board these two dimensions of generality, we get:

(G_1-G$^\square$)   Necessarily, for all x, if $G_1(x)$, then x is epistemically good.
(G_2-G$^\square$)   Necessarily, for all x, if $G_2(x)$, then x is epistemically good.

   (G_1-G$^\square$) together with (G-GE$_\square$) delivers $\left(\text{GE}^*_{G1\square}\right)$, while ($G_2$-G$^\square$) together with (G-GE$_\square$) delivers $\left(\text{GE}^*_{G2\square}\right)$.
   Claim: (G_1-G$^\square$) and (G-GE$^\square$) entails $\left(\text{GE}^*_{G1\square}\right)$. Argument: assume ($G_1$-G$^\square$) and (G-GE)—that is, necessarily, for all x, if $G_1(x)$, then x is epistemically good and also, necessarily, for all x, if x is epistemically good, then x is a goal of enquiry. Now consider an arbitrary possible world w and an arbitrary ev-bearer x. Suppose $G_1(x)$ in w. Then, by (G_1-G$^\square$), x is epistemically good in w. By (G-GE$^\square$), x is a goal of enquiry in w. Hence, in w, if $G_1(x)$, then x is a goal of enquiry. Since w and x were both arbitrary, we have: necessarily, for all x, if $G_1(x)$, then x is a goal of enquiry. A completely analogous argument shows that (G_2-G$^\square$) and (G-GE$^\square$) entails $\left(\text{GE}^*_{G2\square}\right)$.
   Having established $\left(\text{GE}^*_{G1\square}\right)$ and $\left(\text{GE}^*_{G2\square}\right)$, let us return to the task of showing that ($G^{*GE\text{-}GG}$) is a consequence of $\left(\text{GE}^*_{G1\square}\right)$ and $\left(\text{GE}^*_{G2\square}\right)$. Here is the argument: assume (GE$_{G1}$) and (GE$_{G2}$)—that is, necessarily for all x, if $G_1(x)$, then x is a goal of enquiry, and also, necessarily, for all x, if $G_2(x)$, then x is a goal of enquiry. Now, for an arbitrary world w and an arbitrary ev-bearer x, assume $G_1(x)$ or $G_2(x)$. Suppose that $G_1(x)$. By $\left(\text{GE}^*_{G1\square}\right)$, x is a goal of enquiry. Suppose that $G_2(x)$. By $\left(\text{GE}^*_{G2\square}\right)$, x is a goal of enquiry. Thus, whether $G_1(x)$ or $G_2(x)$, x is a goal of enquiry. Hence, x is a goal of enquiry. So, if $G_1(x)$ or $G_2(x)$, then x is a goal of enquiry. Since x was arbitrary, we can generalize on x: for all x, if $G_1(x)$ or $G_2(x)$, then x is a goal of enquiry. Since w was

arbitrary as well, we can necessitate: necessarily, for all x, if $G_1(x)$ or $G_2(x)$, then x is a goal of enquiry. This shows that $(G^{*GE\text{-}GG})$ follows from $\left(\mathbf{GE}^*_{G1\square}\right)$ and $\left(\mathbf{GE}^*_{G2\square}\right)$, i.e. that if G_1 and G_2 are both goals of enquiry as a matter of necessity, then so, too, is the disjunctive property of being G_1 or G_2.

Let me restate what I take the significance of this result to be. Being a goal of enquiry is the feature that, in the argument just given, is assumed to be the unifying of the pluralist's epistemic goods—and so the feature in terms of which the generic epistemic good G^* is characterized. It has been shown that the property of being G_1 and G_2 has this feature—i.e. it is a goal of enquiry—as a matter of necessity. I thus take it that, by the lights of (G^*), being G_1 or G_2 is a generic epistemic good. Now, if G^* is identified with being G_1 and G_2, we have a somewhat familiar good on our hands—namely, a specific case of G_G. As argued earlier, the pure epistemic pluralist does not have to worry about this specific epistemic good.

Now, the argument just given proceeded on the assumption that there are two epistemic goods and precisely one feature that unifies these two goods. Can the argument be generalized? Yes. As said, the assumptions of two goods and one unifying feature were made for the sake of exposition. Nothing in the argument seems to hang on either of them, and so it would seem that the argument can be generalized in at least two ways. First, the argument can be generalized to cover an arbitrary number of non-derivative epistemic goods, meaning in this particular context an arbitrary number of disjuncts. Second, the argument can be generalized to cover an arbitrary number of unifying features. Given non-derivative epistemic goods $G_1, \ldots, G_n$ what we have, then, is the conclusion that the corresponding disjunctive property will satisfy (G^*), i.e. it will possess the unifying features as a matter of necessity:

$\left(G^*_{Fm\text{-}GG}\right)$ Necessarily, for all x: if x is G_1 or … or x is G_n, then $F_1(x), \ldots, F_m(x)$.

The generalized argument supports the conclusion that the epistemic good characterized as the disjunction of the pluralist's non-derivative goods, however many, is going to possess the unifying features as a matter of necessity.

References

Alston, W. 1989. *Epistemic Justification*. Ithica, NY: Cornell University Press.

Alston, W. 1993. Epistemic Desiderata. *Philosophy Phenomenological Research* 53: 527–551.

Alston, W. 2005. *Beyond 'Justification': Dimensions of Epistemic Evaluation*. Ithica: Cornell University Press.

Beall, Jc, and G. Restall. 2006. *Logical Pluralism*. Oxford: Oxford University Press.

BonJour, L. 1985. *The Structure of Empirical Knowledge*. Cambridge: Harvard University Press.

Burge, T. 1993. Content Preservation. *Philosophical Review* 102: 457–488.

Burge, T. 2013. *Cognition Through Understanding: Self-Knowledge, Interlocution, Reasoning, Reflection*. Oxford: Oxford University Press.

———. 1998. Computer Proof, A Priori Knowledge, and Other Minds. *Philosophical Perspectives* 12: 1–37.

———. 2003. Perceptual Entitlement. *Philosophy and Phenomenological Research* 67: 503–548.

———. Forthcoming. Entitlement: The Basis for Empirical Warrant. In *Epistemic Entitlement*, ed. P. Graham and N.J.L.L. Pedersen. Oxford: Oxford University Press.

Cotnoir, A. 2013. Validity for Strong Pluralists. *Philosophy and Phenomenological Research* 83: 563–579.

David, M. 2001. Truth as the Epistemic goal. In *Knowledge Truth, and Duty: Essays on Epistemic, Responsibility, and Virtue*, ed. M. Steup, 151–169. Oxford: Oxford University Press.

———. 2005. Truth as the Primary Epistemic Goal: A Working Hypothesis. In *Contemporary Debates in Epistemology*, ed. M. Steup and E. Sosa, 296–312. Oxford: Blackwell.

DePaul, M. 2001. Value Monism in Epistemology. In *Knowledge, Truth, and Duty: Essays on Epistemic, Responsibility, and Virtue*, ed. M. Steup, 170–183. Oxford: Oxford University Press.

———. 2004. Truth Consequentialism, Withholding and Proportioning Belief to the Evidence. *Philosophical Issues* 14: 91–112.

Edwards, D. 2011. Simplifying Alethic Pluralism. *Southern Journal of Philosophy* 49: 28–48.

———. 2012. On Alethic Disjunctivism. *Dialectica* 66: 200–214.

———. 2013. Truth, Winning, and Simple Determination Pluralism. In *Truth and Pluralism: Current Debates*, ed. N.J.L.L. Pedersen and C.D. Wright, 113–122. New York: Oxford University Press.

Goldman, A. 1988. Strong and Weak Justification. *Philosophical Perspectives* 2: 51–69.

Goldman, A. 1999. *Knowledge in a Social World*. Oxford: Oxford University Press.

———. 2001. The Unity of the Epistemic Virtues. In *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*, ed. A. Fairweather and L. Zagzebski, 30–49. Oxford: Oxford University Press.

Kvanvig, J. 2003. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.

Lynch, M. 2001. A Functionalist Theory of Truth. In *The Nature of Truth: Classic and Contemporary Perspectives*, ed. M. Lynch, 723–749. Cambridge: MIT Press.

Lynch, M. 2013. Three questions for alethic pluralism. In *Truth and Pluralism: Current Debates*, eds. N.J.L.L. Pedersen and C.D. Wright, 21–41. Oxford: Oxford University Press.

———. 2004. Truth and Multiple Realizability. *Australasian Journal of Philosophy* 82: 384–408.

———. 2009. *Truth as One and Many*. Oxford: Oxford University Press.

Moser, P. 1985. *Empirical Justification*. Dordrecht: D. Reidel.

Pedersen, N.J.L.L. 2006. What Can the Problem of Mixed Inferences Teach Us About Alethic Pluralism? *The Monist* 89: 103–117.

———. 2010. Stabilizing Alethic Pluralism. *Philosophical Quarterly* 60: 92–108.

———. 2012a. Recent Work on Alethic Pluralism. *Analysis* 72: 588–607.

———. 2012b. True Alethic Functionalism? *International Journal of Philosophical Studies* 20: 125–133.

Peels, R. 2010. Epistemic Desiderata and Epistemic Pluralism. *Journal of Philosophical Research* 35: 193–207.

Plantinga, A. 1993. *Warrant and Proper Function*. New York: Oxford University Press.

Pritchard, D. 2014. Truth as the Fundamental Epistemic Good. In *The Ethics of Belief*, ed. J. Matheson, and R. Vitz, 112–129. Oxford: Oxford University Press.

Shapiro, S. 2014. *Varieties of Logic*. Oxford: Oxford University Press.

Tappolet, C. 2000. Truth Pluralism and Many-Valued Logics: A Reply to Beall. *Philosophical Quarterly* 50: 382–385.

Wright, C. 1992. *Truth and Objectivity*. Cambridge: Harvard University Press.

———. 2013. A Plurality of Pluralisms. In *Truth and Pluralism: Current Debates*, ed. N.J.L.L. Pedersen and C.D. Wright, 123–153. Oxford: Oxford University Press.

Zagzebski, L. 2004. Epistemic Value Monism. In *Sosa and His Critics*, ed. J. Greco, 190–198. Oxford: Blackwell.

Zanetti, L. 2014. Epistemic Pluralism: As One and Many. Master's thesis, University of Modena e Reggio-Emilia, Modena, Italy.

A Plea for Epistemic Monism

Pascal Engel

1 Introduction

There are at least three senses in which one can talk of epistemic pluralism. The first is a version of epistemic relativism, which one may formulate thus: "There are many fundamentally different, genuinely alternative epistemic systems, but no facts by virtue of which one of these systems is more correct than any of the others" (Boghossian 2005, p. 73). This is the idea that different communities, often distant in space and in time but possibly also close to our present day world, can have different, and incommensurable, conceptions of the epistemic properties of their beliefs, and in particular of what it takes for a belief to be justified, to be true or to be known. Such views are illustrated by historical analyses of the "social history of truth", by constructivist views in the sociology of knowledge, or by various claims in feminist

P. Engel (✉)

Centre de Recherches sur les Arts et le Langage, Ecole Des Hautes études
En Sciences Sociales, 96 Boulevard Raspail, 75006 Paris, France
e-mail: pascal.engel@ehess.fr

epistemology or in experimental philosophy, which aim at deriving from historical, anthropological, psychological and sociological "data" claims about the fundamental diversity of epistemic evaluations relative to culture, background and cultural frameworks, and to argue against the "absolutist" view according to which there can be but one cross-cultural and trans-temporal system of evaluation. The second is a version of epistemic contextualism, the view that the meanings of epistemic weighty terms such as "true", knows", "is justified", or "is warranted" vary across contexts and circumstances, depending on who attributes such properties and on the epistemic situation of the subjects. The third is the meta-epistemological view emerging from the discussion about the nature of justification and knowledge, according to which there cannot be a unified conception of these notions, which are necessarily diverse. It is this controversy which interests me here. Although it has indeed important ties with the other two debates, it does not bear on the conceptions of justification, of truth and of knowledge which are implicit in common sense and in different cultures, but on the explicit conceptualizations of these notions by philosophers. The issue is whether epistemology ought to proceed with one and only one notion of justification and of knowledge, or it ought to accept that there are several such notions, which do not serve the same purposes but are equally valid. Epistemic pluralism in this sense is a widespread assumption in contemporary epistemology. Many epistemologists agree that our *philosophical* notions of knowledge and justification do not always answer the same criteria and that our epistemic evaluations may vary depending on our practices, informational needs and the circumstances in which they are issued. This is reinforced by studies, purportedly "experimental", on our ordinary knowledge ascriptions and their trappings. Epistemic pluralism in this sense is often associated with epistemic contextualism, but it does not entail it. Pluralism about theories of justification does not by itself entail that the ordinary notion of knowledge is context dependent. My question here is whether epistemic pluralism in this theoretical sense is correct: Do we have to renounce the idea that there is but one dimension of epistemic evaluation, to which the variety of our practices are related? I shall argue that the pluralism is only apparent, that has to be only one dimension of epistemic evaluation and that there are no good reasons to espouse a form of epistemic pluralism.

In the first section, I shall give some *prima facie* reasons to adopt epistemic pluralism. In the second section, I shall try to characterize epistemic pluralism in this meta-epistemological sense, by distinguishing different versions of this view. Each of them starts from the familiar idea that there is an ambiguity in the notion of justification and in other related notions, in particular between an internalist and an externalist sense. From there one can defend three kinds of views. The first consists in the claim that the various senses of justification are actually compatible and can be integrated, if they are located at the appropriate level. This is a version of what is often called "epistemic compatibilism". The second is the claim that the various senses of the notion justification are incompatible, and serve different functions corresponding to distinct "epistemic *desiderata*". The third is a version of functionalism about justification: justification is a property which is realized in different ways according to the kind of domain to which it applies.

In the third section, I present my objection to these three versions. It is that they are pluralist only *prima facie*. They in fact presuppose that there is but one fundamental dimension of evaluation, which is centred on truth, evidence and knowledge, which do not vary in meaning and do not bifurcate into various notions. So I shall defend, against each form of epistemic pluralism, a form of epistemic monism.

In the fourth section, I examine another possible version of epistemic pluralism, Duncan Prichard's epistemological disjunctivism, and argue that it fails to be pluralist in the required sense. In the last section, I consider two other illustrations of the primacy and centrality of one kind of epistemic status: pragmatic encroachment and epistemic injustice. Each of these purported phenomena presupposes the very concept of evidence rather than they provide reasons for pluralizing it. Epistemic monism should not, however, deny the existence of various dimensions of epistemic assessment. It should rather put them in their place, which is secondary to the central dimension.

Within the limits of this essay, it is possible only to give the sketch of a defence of epistemic monism. A full argument that the notions of truth, of justification and of knowledge cannot be pluralized would require much more, for in a sense the issue is all what the epistemological debate is about. I hope nevertheless to give some reasons to resist some of the pluralistic moves which are so influent in contemporary epistemology.

2   Motivations for Epistemic Pluralism

One of the most common assumptions of contemporary epistemology is that although there are several ways of evaluating beliefs and of sorting out those which are susceptible to count as knowledge, these evaluations form a very tight conceptual circle. On this view there is but one dimension of epistemic evaluation: beliefs are true or false, justified or not, rational or not, conform or not to evidence, come to be knowledge or not in one sense only, and there is but one notion of truth, of justification, of rationality and of evidence. Indeed, these notions are distinct and do not entail each other: it is one thing for a belief to be true, another thing for it to be justified or rational, yet another thing for it to suit evidence. Belief can be justified without being true, rational without answering evidence. But these properties are closely connected: evidence and justification are evidence and justification *for* truth, and rationality is in large part a source of justification for one's beliefs, hence for their potential truth (without being identical to it). The fact one has reasons to believe that p does not entail that p is true, but the fact that p can be a reason for believing that p. Truth is generally taken to be the goal of inquiry. Let us call this assumption of the unity of epistemic evaluation *epistemic monism*. It is implicit in most of contemporary epistemology. It receives an explicit formulation in a passage from Laurence Bonjour:

> The distinguishing characteristic of epistemic justification is thus its essential or internal relation to the cognitive goal of truth. It follows that one's cognitive endeavors are epistemically justified only and to the extent that they are aimed at this goal, which means very roughly that one accepts all and only those beliefs which one has good reason to think are true. To accept a belief in the absence of such a reason … is to neglect the pursuit of truth; such acceptance is, one might say, epistemically irresponsible. My contention here is that the idea of avoiding such irresponsibility, of being epistemically responsible in one's believings, is the core of the notion of epistemic justification. (Bonjour 1985, p. 8)

Bonjour clearly relates the goal of truth to epistemic justification and to the internalist requirement that one ought to be responsible for one's

beliefs. Now epistemic monism—in this sense of there being only one epistemic goal, namely truth—is perfectly compatible with the existence of various theories of what justification ought to be. Some philosophers understand it, as Bonjour, in an internalist sense—any justification or reason for a belief has to be such that one has access to it—but others deny that this requirement is mandatory. There are diverging views of what justification is, and externalists about justification, who hold the opposite view, reject this requirement, but no one in this debate denies that there has one single property to which justification corresponds. The fact that some philosophers define justification in different and often conflicting senses and characterize it in deontological terms (what one ought to believe or is permitted to believe), whereas others define it in terms of good evidence, reliability, proper function, reason, virtue or other concepts, is perfectly compatible with epistemic monism: one of these definitions might be the right one, while the others are incorrect. Either internalism or externalism is the right view, or either foundationalism or coherentism is correct, but they cannot be all correct together. Epistemic pluralism becomes an option when one suspects that all of these senses apply equally and legitimately to the notion of justification, without there being only one notion which is *the* correct one. A first motivation for epistemic pluralism arises with the observation that there is not a unique dimension along which we evaluate our beliefs and our claims to knowledge. When we ask deeply theoretical questions about our knowledge of the world, in particular about the validity of our scientific theories, or when we raise deep metaphysical issues about the threat of scepticism, we inquire along one single and very demanding dimension of evaluation, but when we leave the theoretical space of the scientific journals or of the epistemology seminar room, we use much more relaxed standards. This is a contextualist commonplace: our epistemic evaluations depend on the perspective of attribution. For instance, in Harman's (1973) well-known case of the person who reads about the assassination of the President in an early edition of a newspaper, but where there is a media conspiracy later to replace it with a modified edition that retracts the story, the agent's belief is justified, but it does not constitute knowledge, because of the new information that is available in her practical or social environment. Contextualists exploit such situations to argue that epistemic evaluation

is highly sensitive to circumstances and points of view. Moreover, it may not be the same concept of justification and of knowledge which is at work here. The former seems to be tied to individual reflective evaluation of our own beliefs, the latter to a socially determined context which is utterly different. Why should we expect that the criteria which apply to individual knowledge should apply to social knowledge?

A second motivation for epistemic pluralism comes from the fact that there are different kinds of knowledge, related to diverse sources, which seem not to call for the same kind of justification: thus perceptual knowledge, testimonial knowledge, memory knowledge, knowledge of other minds, socially extended knowledge through artefacts, a priori knowledge, inferential self-knowledge do not seem to obey the same criteria of justification. Empirical knowledge and mathematical knowledge do not appear to be the same ilk. Neither perceptual knowledge nor knowledge of our own minds. Knowledge is also relative to biological species, and within one species to different sources. We share a large part of our cognitive equipment with other mammals, and our "animal knowledge" is rather different from the "reflective knowledge" which we may be the only species to enjoy. To each kind of knowledge seems to be attached a specific kind of justification and a specific level of justification. Why should perceptual knowledge obey the kind of criteria which apply to a priori knowledge or to testimonial knowledge? Self-knowledge, inferential knowledge and a priori knowledge seem to enjoy a level of certainty which other kinds of knowledge do not have. Moreover, within one domain of knowledge, for instance self-knowledge, there may be different varieties of it (Coliva 2016). This has led a number of philosophers to defend the view that these species of knowledge should be granted a different kind of epistemic status from the ordinary, reflective, notion of justification: some kind of *prima facie* kind of justification, or some form of entitlement, the nature of which has to be determined.

A third motivation concerns the varieties of the entities which are subject to epistemic evaluation. Most of the time, the candidates are beliefs' contents. But there is a difference between on the one hand the evaluation of the proposition which a believer is susceptible to entertain on the basis of a given piece of evidence ("propositional justification") and on the other hand the actual possession of the appropriate

belief ("doxastic justification"). Sometimes the objects of evaluation are not the contents of beliefs, but the believings, that is the processes of belief formation, their reliability and their safety. Some theories, notably those based on virtues, dispositions and skills, make the agent the object of epistemic evaluation. Conceptions of epistemic virtues are themselves diverse, depending on whether one insists on their voluntary acquisition or not, and on the respective emphasis on moral and intellectual virtue. And indeed when epistemology deals with forms of social knowledge, the agents whose beliefs are evaluated are collective. So one may suspect that justification and knowledge may not have the same meaning depending on the entities to which they are assigned. Moreover, there is also a conflict between epistemic and ethical evaluations. These are, arguably, quite different, and ethics is not epistemology. But there are cases where these evaluations overlap.

A fourth motivation for epistemic pluralism consists in the difficulty to conciliate the various kinds of epistemic evaluations proposed by epistemologists. Internalism requires that agents have access to the properties which justify their beliefs, whereas externalism shuns away such a requirement. Internalist justification is typically tied to the mental states of the agent and to their reasons for believing, whereas externalist justification is typically tied to the world or environment in which beliefs operate, and to their truth. The same sort of duality of internal vs external face affects the very notion of reason. On one reading, reasons for belief are internal states of the agent, seen from his point of view or perspective. On another reading, reasons are facts in the world, which agents may not be aware of. The very concept of reason is also subject to another ambiguity: some reasons are said to be "normative", i.e. those that an agent ought to have whether or not he is aware of the facts, whereas other reasons are "motivating" and relative to the agent's evidence when acting or believing. The concept of evidence may refer on the one hand to the non-factive mental states which support a subject's beliefs and on the other hand to their factive mental states.

Given this diversity of epistemic notions, one argument for epistemic pluralism might be that there are too many competing and incompatible accounts of justification for it being possible that only one property corresponds to the notion or concept of justification. But one might also argue

that these accounts are actually compatible and that the choice between them depends on one's purposes (Alston 2005). Why are there so strong divergences of intuitions about the epistemologists' favorite cases such as the story of Samantha the clairvoyant who "knows" that the President is in New York or of Barney who "knows" that this is a barn within a landscape full of fake barns? The existence of strong disagreements among epistemologists on the notion of justification is but one argument, which is actually not very good: Does, for instance, the disagreement between philosophers about the notion of truth entail that there are different notions of truth? A further argument has been advanced by "experimental philosophers": lay subjects, when presented with various test cases, such as Gettier cases, about how they understand terms like "knowledge", reply differently, depending on their cultural background or circumstances of evaluation. These variations in the ordinary attributions of knowledge, which parallel the variations in our epistemic intuitions about "cases", seem to suggest that they do not pick out a single property.

This diversity is also reflected in the variety of terminologies which philosophers have used in discussing epistemic evaluation. Some stick to the traditional terminology of "justification", but others prefer to advance other notions, in order to remove the internalist sounding: thus "warrant" has been used to designate the property which corresponds to externalist justification (Plantinga 1992). Some philosophers prefer to use "justification" only in the internalist sense and reserve the externalist requirements to knowledge (Audi 1988). This line is indeed tempting when one reflects on the Gettier problem: Gettierised subjects are justified but their true beliefs do not amount to knowledge. Some philosophers prefer to recast these debates in terms of "reasons" (Skorupski 2011). Those who defend a "knowledge-first" conception of knowledge prefer to get rid of the very notion of justification. These different labels for designating positive epistemic status often reflect strong theoretical oppositions, but they may also signal purely terminological differences with no substantive import.

Last but not least, there has been, in particular under the influence of E.J. Craig's (1991) conception of knowledge and due to the notorious failure to solve the Gettier problem, a tendency among epistemologists

to consider the *function* which the concept of knowledge plays in various practices of epistemic evaluation, rather than trying to define this concept through necessary and sufficient conditions. Many thinkers have drawn the conclusion that there might be more than one such function, hence that epistemic evaluation is inevitably plural.

To such a variety of dimensions of epistemic evaluation, one can react in two different ways. One can take the classical line and say that behind this diversity of vocabularies there must be a unity: there cannot be several stories about epistemic evaluation, and one theory has to account for its central features. Alternatively, one can go pluralist and accept that there are several equally legitimate concepts of epistemic evaluation. Such a pluralism needs not be a form of relativism, which would say that all kinds of epistemic evaluation are worthy of the name. Epistemic pluralism is most of the time the view according to which our epistemic evaluations serve different purposes depending on the aims and the circumstances. To use a familiar metaphor, just as our concept of knowledge might well be like a "Swiss Army knife", operating in different ways in different contexts (Weiner 2005), our other epistemic concepts might be necessarily diverse and purpose oriented.

I reject this idea and take the classical line. But before giving my reasons for this conservative position, it is necessary to formulate epistemic pluralism more precisely.

3   Three Kinds of Epistemic Pluralism

There are at least three kinds of views that one might propose under the label of epistemic pluralism. The first one I shall call the *ambiguity or dualistic view*, the second the *epistemic desiderata approach* and the third the *functionalist* view.

The *ambiguity view* says that there are two main senses of "justification", the internalist and the externalist, which are *prima facie* incompatible and which do not designate the same property, just as "bank" can designate a financial institution or the border of a river. When epistemologists discuss the notion of justification, and when one camp understands it in the internalist sense according to which justification

entails that one has access to one's reasons for believing and when the other camp defends an externalist and reliabilist conception of justification, they equivocate and talk past each other. They refer to two distinct conceptions of justification. This verdict can be reached from many angles. For instance in Gettier examples, one can concentrate on the internal side of justification, take in consideration how the Gettierised agent reaches his belief (e.g. that Mr. Nogot owns a Ford), and accept that one can be justified in believing something that is false, or, alternatively, one can consider the situation from the outside, and accept that the agent can have a lucky true belief which nevertheless has grounds, unbeknownst to him (Fogelin 1994). The same verdict of ambiguity is present in the examples which internalists have opposed to externalists, such as the case of the clairvoyant, who is "justified" in the reliabilist sense in believing that the President is in New York, but does not have a clue about why she is justified in believing this (Bonjour 1988), or in the case of Lehrer's Truetemp (Lehrer 1990) who has correct beliefs about temperature without having the slightest idea about how they come to him. The same verdict can be issued in all the cases where people have false but epistemically blameless beliefs. In one sense, the Ancient astronomers ought to have believed that the Earth revolves round the Sun; in another sense, they were justified in believing the contrary. Still, the ambiguity view is not a form of pluralism. If epistemologists talk past each other, this does not mean that their respective concepts of justification are correct. It is open to the internalist to object, as they actually did, that externalists do not talk of justification at all, when they accept a reliabilist view. Pluralism emerges when one accepts that *both* the internalist and the externalist are correct and describe justification from two perspectives which one takes to be compatible. This view, which is a form of dualism, can be argued for in several ways. I shall consider only two examples.

The first is Goldman's (1991) complex view of justification which aims at combining the different senses of this notion and which he presents as "a marriage of internalism and externalism" (1991, p. 139). First Goldman distinguishes justification at the level of reliable processes from justification at the level of methods. To be strongly justified, one's beliefs must be based on reliable processes and on adequate methods.

But one can form a belief by way of adequate process but inadequate method, when one has a blameless or non-culpable belief. In this weak sense of justification, one can form a belief through an unreliable method, although one does not believe that one's method is unreliable. A belief is strongly justified if and only if (i) it is produced by a reliable process, and (ii) the process is not undermined by the agent's cognitive state. A belief is weakly justified if and only if (i) it is produced by an unreliable process (ii) the agent does not believe the process to be unreliable, and (iii) the agent has no evidence that the process is unreliable. Thus, weak justification is internalist. And strong justification is externalist.

The second example of a double-level theory is Sosa's (2009) distinction between animal and human knowledge. Animal knowledge is only a matter of arriving at true beliefs by the employment of apt faculties, whereas reflective (or human) knowledge requires, in addition, that the subject be internalistically justified in thinking that his belief be grounded in a reliable cognitive faculty.

Such proposals can be interpreted in two ways. In the first place, they can be interpreted in the previous sense as saying that notions like "knowledge" and "justification" are ambiguous and that strong justification and animal knowledge on the one hand, based on such notions as reliability of processes and aptness, give us one sense of "knowledge" and "justification", while weak justification and reflective knowledge on the other hand, give us another sense. On this view, "knowledge and justification" can be understood *either* in the first—animal and nonreflective—sense *or* in the second—reflective and internalist—sense. This is a dual aspect epistemology, just as there is a dual aspect theory of the mind ("property dualism"). In the second place, Goldman and Sosa's respective proposals can be interpreted as saying that, although the two kinds of knowledge and justification are distinct, they can be combined. On Sosa's view in particular, animal knowledge is indeed knowledge, but one knows "full well" only if one combines this animal level with the reflective level. Both obey different criteria, but the former is a precondition for the latter. In terms of Sosa's metaphor of the archer:

On one level, how apt the shot is depends on the degree of competence manifest by its success. But, on another level, the full aptness of the shot depends also on the meta-competence manifest by its aptness and by its success. A performance is fully apt only if its first-order aptness derives sufficiently from the agent's assessment, albeit implicit, of his chances of success (and, correlatively, of the risk of failure). (Sosa 2011, p. 11)

On such a view, the two kinds of knowledge and justification have to be *combined* in order to reach a better form of justification and knowledge than the animal one. Each one contributes to the total result, which is full knowledge, and full knowledge is an exclusive status.[1] This is a form of epistemic pluralism only by courtesy, for such a view entails that the two kinds of knowledge or justification are equally entitled to receive the name of knowledge and to perform this role, whereas animal knowledge and reflective knowledge are not by themselves sufficient for knowledge. On Sosa's view, both are necessary, but they are sufficient only jointly for full knowledge. This is better described as a kind of epistemic dualism, or, as. Goldman calls his own view, a "duplex" conception of justification.

A second and more likely candidate for epistemic pluralism is the view that there are several notions of justification, depending on the purpose which they are meant to serve, and which all are suited for having a justificatory status. This is, in particular, Alston's position in *Beyond Justification* (Alston 2005). Alston first argues that the property of justification can serve a number of objectives, which are all equally important and respectable. He lists a number of candidate conditions for epistemic justification. These are:

(i)   cognitive accessibility to one's reasons
(ii)   higher-level knowledge of the low-level conditions of one's knowledge
(iii)   capacity to defend one's epistemic status
(iv)   reliability of the process of belief production
(v)   coherence between one's beliefs
(vi)   intellectual virtue and proper epistemic credit
(vii)   capacity to reflect on one's evidence.

None of these conditions are sufficient, says Alston, but they are all *prima facie* sufficient and compatible with several concepts of justification: internalist, externalist, evidentialist, reliabilist, deontological or virtue theoretic. This means, according to Alston, that each of these properties satisfies equally well the requirements for justification, and consequently that *there is no unique property* to which they answer:

> There isn't any unique, epistemically crucial property of beliefs picked out by 'justified'. Epistemologists who suppose the contrary have been chasing a will-of-the-wisp. What has really been happening is this. Different epistemologists have been emphasizing, concentrating on, "pushing" different epistemic desiderata, different features of belief that are positively valuable from the standpoint of the aims of cognition. These include the features we have been listing above in surveying views as to the nature and conditions of justified belief. They include such features as a belief's being permitted by relevant rules or norms, a belief's being based on adequate grounds, a belief's being formed in a reliable way, a belief's fitting coherently in a coherent system, and so on. Somehow the practice has spread of taking one's attachment to a certain epistemic desideratum as deriving from its being part of what it is for a belief to be "justified" or what is required for that. But the supposed connection with "justification" has nothing to do with what makes a desideratum epistemically desirable. There is no substance to that connection; it is an honorific title that carries no remuneration, perks, or further implications along with it. It is not as if one needs to show that, for example, reliability of formation or evidential support or coherence is what matters for justification in order to validate its epistemic credentials. There is no such reality as epistemic justification to perform that function. All we have is the plurality of features of belief that are of positive value for the cognitive enterprise. They need no validation from a connection with a supposed master epistemic desideratum picked out by 'justified'. There isn't any such. A belief's being justified has no more objective reality than ether or ghosts. (Alston 2005, p. 22)

Alston's view is pluralist in the sense that (a) it rejects the idea that there could be a unique property of "justification" the nature of which could be revealed by our epistemological inquiries, (b) considers that in

spite of the fact that there is but one epistemological goal, significant truth, there are various dimensions along which this goal can be attained and (c) our beliefs can be evaluated along several of the dimensions (i)–(vii) above, or along one of them only, depending on the kind of purposes one has. No single notion of epistemic justification, internalist or externalist, deontic or reliabilitist, evidentialist or not, is the correct one. They are all suitable, and equally valuable, according to one's emphasis and objectives. This entails clearly not only that justification as such does not exist, but also that it is relative to our aims. There is then, according to Alston, no single story to be told on epistemic justification.

The third possible view which may qualify as a form of epistemic pluralism is a form of *functionalism*. As far as I know, the view has never been defended in print, but we may try to formulate it on the model of the kind of alethic functionalism which has been defended about truth by Crispin Wright (2001) and Michael Lynch (2009), which their proponents sometimes call a form of pluralism. In a nutshell, alethic functionalism says that truth is a high-order functional property characterized by a set of platitudes, such as disquotation, objectivity, normative role, correspondence role, which are realised in different ways depending on the domains (mathematics, physics, ordinary objects, ethics, aesthetics) to which it applies. There is only one truth role, but there are distinct truth properties at the level of their realization within a domain of discourse. Thus, in mathematics, truth might be a form of assertibility, in physics a form of correspondence, in ethics a form of coherence, perhaps something else when we consider aesthetic discourse, etc. It is not decided a priori which property will be instantiated in each domain. Similarly, we might construct *justification functionalism* as the view according to which justification is a high-level property characterized by a set of very general properties, realised differently in different domains. The list of the properties in question would strongly overlap Alston's the list of *desiderata*. Along these lines we may suggest that there is a core concept of justification, constituted by the following platitudes:

(a)   *reliability*: a belief is justified if it is caused appropriately by reliable processes.

(b)   *access*: potential access internalism: a belief is justified if the agent has potential access to his reasons.
(c)   *critical assessment*: a belief is justified if the agent is able to critically assess it.[2]
(d)   *basis*: a belief is justified if it is properly based on evidence.
(e)   *coherence*: a belief is justified it potentially coheres with the set of an agent's belief.
(f)   *normativity*: a belief being justified entail that it has a certain normative status.

Justification functionalism would be the view that these platitudes, taken together, implicitly define the property of justification, which can be realized in different ways in different domains.[3] Thus, one may suggest that these properties will be realized differently, and to varying degrees, for perceptual knowledge, inferential knowledge, testimonial knowledge, or memory knowledge. For example, perceptual knowledge would obey essentially the reliability requirement and conform only marginally to the requirement of coherence. We can also envisage the functionalist idea with respect to the kind of question for which the issue of justification arises. When the issue concerns scepticism, one may suggest that the reliability requirement is not enough, as the "new evil demon intuition" suggests: it is possible for someone's beliefs to be justified even when the processes that produced those beliefs are not reliable. But when it comes to testimony, we may suggest that a weaker form of justification is enough. This may coincide with a weaker notion of justification, such as that of *prima facie* justification or of entitlement, which has been invoked for perceptual and testimonial knowledge.

The main difficulty, if one tries to formulate justification functionalism, is that the various platitudes do not really capture a unique minimal property of justification and may conflict, unlike those which, according to alethic functionalism, implicitly define truth. Thus, the reliability feature potentially conflicts with the access internalist feature and with the normativity feature. All the problem is: How can they *both* be correct justifications? So the functional property of justification can hardly be a "thin" property in the manner of truth according to alethic

functionalism. It must be a property which combines externalist and internalist features, in the manner of the dualist views described above.

Functionalism about justification, thus sketched, seems a bit cheap. It gives us pluralism about justification without really telling us how justification in each case is achieved. But it is a possible way of formulating a pluralistic idea.

4   Why Epistemic Pluralism Is Not Really Pluralist

A conception of epistemic justification, in order to count as pluralist, must meet certain conditions. In the first place, it must be such that it concerns *epistemic* justification and not other kinds of justification. Practical judgments about actions, moral judgments, aesthetic judgments can all said to be justified, in various senses, when the respective evaluations are correct. But the fact that the notion of justification applies to diverse domains does not count as displaying a variety within *epistemic justification*. An action, if justified, is not justified in the same sense in which a belief is said to be justified, even though there are similarities between the two. Reasons for action and not reasons for belief. Although there are important similarities and although the questions often arise of the relationships between practical and epistemic evaluation, this does not entail any identity between the two, unless one subscribes to a strong form of pragmatism.[4] Neither can the kind of dualist view or two levels defended by thinkers like Goldman and Sosa.

As we have seen, for Goldman strong and weak justification, and for Sosa competence-based justification and reflectively based justification, do not compete: they can be applied to one and the same belief and both contribute to its justification. Actually, in Sosa's view (2011), although there are two kinds of knowledge and two kinds of justification, the two strands must compose. Although animal knowledge is actual knowledge because it can be apt (competent), accurate (reach truth successfully) and adroit (accurate because of the exercise of competence), full knowledge is achieved only when the agent acquires a

reflective grasp of his aptness, accuracy and adroitness (the "AAA" scheme). The unity of the whole "AAA" scheme is, in Sosa's recent presentations of his views, secured by the idea that believing and knowing are certain kinds of performances, which involve agents aiming at truth and acquiring certain dispositions, and then skills and intellectual virtues, in order to be able to assess their risks and guide rationally their apt animal knowledge, thus coming to know full well. This picture is well summarized by Sosa:

> Fully apt performance goes beyond the merely successful, the competent, and even the reflectively apt. And it is the human, rational animal that can most deeply and extensively guide his performances based on the risk involved, in the light of the competence at his disposal. That is why reason must lord it over the passions, both the appetitive and the emotional. (Sosa 2015, p. 87)

This is by no means a pluralistic view of justification: it is a very unified one. Sosa is not claiming that there are two kinds of justification, one based on a lower-level competence, the other on a higher-level competence, which would obey different criteria. On the contrary, the same scheme applies to both, and they are meant to be articulated.[5]

For a theory to be genuinely pluralist, it has to be the case that more than two distinct notions of justification can be applied. In this respect, Alston's conception is probably the one which fits best the pluralistic objective. However, I want to claim that a genuine epistemic pluralism in this sense is hard to come by. For a simple reason: whatever concept of epistemic justification we consider—be it reliability, evidence, reason, epistemic responsibility or epistemic virtuousness—they are all related to, and indeed dependent upon, one and only one dimension of epistemic evaluation, namely *truth*. Evidence is evidence for truth, justification is justification *for* truth and reliability is truth *conduciveness*. In so far as they are epistemic, all these forms of justification are supposed to be related to truth as the unique goal of cognition. Sometimes Alston lists truth as one among others of his epistemic *desiderata*. But this cannot be right: truth is not one dimension of justification among others, it is the essential one. Truth is what we want justified belief *for*.[6]

Pluralism can only be vindicated if it epistemic goals were plural. It has been argued that understanding is a possible candidate for being an epistemic goal. There is no doubt that one of the goals of cognition is to understand what beliefs and theories are, and to grasp their meaning, in addition to being able to grasp their truth (Kvanvig 2003). But it is very dubious that there can be understanding without knowledge of the truth-conditions of belief (Elgin 2009). If this is correct, truth is not one among other epistemic *desiderata*, it is the main *desideratum*. The same can be said about knowledge. In any theory of epistemic justification, beliefs are supposed to be justified in order to become knowledge. Indeed, one kind of argument one can give for epistemic monism is that justification amounts to knowledge and that one is justified to believe that p if and only if one knows that p (Sutton 2007). Knowledge-first epistemology provides one kind of strong argument for epistemic monism. On Williamson (2000) version of this view, not only being justified entails knowing, but knowledge is evidence, and obeys also a condition of safety, which is typically of the externalist kind. If the knowledge-first program is vindicated, epistemic pluralism has to be rejected at the outset. But actually there is no need to commit oneself to the knowledge-first program in order to reject epistemic pluralism. It is enough to argue that one cannot defend the diversity of our concepts of epistemic justification without presupposing that there is *one* central dimension which dominates and actually unifies the others. And that dimension is clearly truth or knowledge. Take the previous dimensions. Reliability of belief is, I said above, the capacity to produce as many *true* beliefs as possible. Access to one's beliefs is clearly the capacity either to have direct awareness of *the truth* of one's beliefs (perhaps in the "transparency" sense (Shah 2003)) or to have second-order beliefs about the truth one's first-order beliefs. One way or another, thinking that one believes that p entails thinking that one takes one's belief that p to be *true.* Evidence is evidence *for truth*. Critical assessment is assessment with respect *to truth*. In particular, the main and most fatal objection to a belief is that it is *false*. And there are good reasons to argue that the aim, and in this sense the norm, of belief is *truth*.[7]

Deflationists and minimalists about truth, relativists of all sorts, Rortyan pragmatists and a number of thinkers who doubt the centrality

of truth in our epistemic evaluations will balk at such affirmations. They will claim that the aim of belief is not truth, but honest justification. They will argue that the very independence of justification from truth is what allows us to rely on a variety of conceptions of justification, depending on our specific purposes. Indeed, the thesis of epistemic pluralism often flows naturally from the deflationist view that truth has no essence and is not the goal of our inquiries: since justification is not oriented towards truth, it is necessarily oriented towards other goals, and these goals are diverse. Their argument, well summarized by Rorty is that:

> The need to justify our beliefs to ourselves and our fellow agents subjects us to norms, and obedience to these norms produces a behavioural pattern that we must detect in others before confidently attributing beliefs to them. But there seems to be no occasion to look for obedience to an additional norm – the commandment to seek the truth. For – to return to the pragmatist doubt with which I began – obedience to that commandment will produce no behaviour not produced by the need to offer justification (Rorty 1995, p. 26).

In fact, the Rortyan argument is more a *reductio* of the idea that justification is independent from truth than a vindication of it. For there is an answer to this kind of argument, which has been given by Wright (1992) and by Price (1998,2003) there is indeed a difference that truth makes in practice for the nature of assertions, since assertions which would not be governed by a norm of truth would not be assertions at all.[8] There would be no difference between objective assertions and merely opinionated assertions. The argument carries over to beliefs. Wright and Price do not conclude from this that truth as to be correspondence in the realist sense, but that any theory of justification has to make room for a contrast between justification and truth: even if one does not accept a strong realist conception of truth, one has to accept that some assertions and beliefs are more justified than others and that there is some objective standard which they answer. Wright calls this standard *superassertibility*, the capacity of our assertions to resist our best efforts to refute them. Whatever the form which the norm of

belief and of assertion can take, there has to be an objective standard, which plays the role of truth. Justification cannot be thought independently from *some* objective concept of truth. This is why the kind of justification functionalism envisaged in the previous paragraph cannot really make sense unless one adds to the five platitudes (a)–(e) above a sixth one:

(f) *truth*: justification aims at truth.

In other words, truth is always presupposed by justification. It is the property which makes justification, evidence, reliability and understanding and meaningful. Otherwise, we are not dealing with *epistemic* evaluation.

One may object to this that the fact that truth cannot be dissociated from justification does not by itself show that epistemic pluralism is impossible. There could well be several concepts of justification which were nevertheless truth oriented. But would this possibility entail a genuine epistemic pluralism?

It is difficult to examine this possibility without discussing specific proposals. We have seen that Goldman and Sosa's dualistic views, which are supposed to render compatible two concepts of justification, one internalist and the other externalist, are not actually pluralist: they involve in fact two distinct levels of justification, both truth driven, and both contributing to the aim of knowledge. In order to have a genuine epistemic pluralism, it is necessary to show that distinct forms of justification can concur to the same goal of knowledge. But is this possible?

5   Epistemological Disjunctivism Is Not an Epistemic Pluralism

In order to try to answer this question, let us consider a different style of epistemic pluralism, which aims at combining internalist and externalist justification, Pritchard's *epistemological disjunctivism* (Pritchard 2012, 2015). Epistemological disjunctivism is a thesis according to which:

> In paradigm cases of perceptual knowledge the knowledge in question enjoys a rational support that is both factive and reflectively accessible. In particular, it is the view that when one has perceptual knowledge in such cases, the reflectively accessible rational support one has for one's knowledge that p is that one sees that p. Note, however, that seeing that p is factive, just like knowing that p. That is, if one knows that p then it follows that p must be true; likewise if one sees that p, then p must be true. (Pritchard 2015, p. 124; Pritchard 2012, p. 11)

Epistemological disjunctivism is meant to apply to perceptual knowledge only. It says that in ordinary perceptual cases, our justification ("rational support") for knowing that p is that one sees that p. This is internalist justification—based on one's direct reflective access to the contents of one's perception. But unlike internalist justification, which does not entail truth, this perceptual justification is factive—one's seeing that p entails that one knows that p. But this condition is realised only when one is in "good" case, where one actually sees that p, that is when one is not in a case of hallucination or in an Evil Demon sceptical scenario. When one is in such a "bad" case, not only one's justification is not factive—one does not know—but it is different from one's justification in the good case. Epistemic disjunctivism does not say, as the "new evil demon intuition" suggests, that our justification is internal and that it is the same in the good and in the bad case. On the contrary in the bad case, we are not justified and do not know. This is what makes the view disjunctivist. It entails that the justification ("rational support") in the two cases is fundamentally different. For in the good case, the reflective accessibility is factive, and in the bad case, the reflective accessibility must be fundamentally different. But if this is so, the justification that one has in the good case—the accessibility that one has to the content of one's belief—cannot be the same in the bad case, for in the good case, one knows, whereas in the bad case one does not know. So we have one justification—internally accessible—which is not truth conducive and another—factive—which is truth conducive. This disjunction seems, by our previous criteria, entail a genuine dualism of justificatory status. Pritchard highlights the difference when he notes that in the factive case, the basis of one's knowing perceptually that p

is that one sees that *p*, which is identical to knowing that *p*. This second sort of justification entails knowledge, whereas the first—the one which answers internalist criteria—does not. Note first that truth is the standard, not justification. This confirms our previous diagnosis that any notion of justification has to presuppose the possibility—however remote—of accessing truth. But for epistemological disjunctivism to make sense, and to be true, it cannot be the case that there is a double standard for justification. The view presupposes that the only kind of justification which can yield knowledge is the one which is factive. It is not *in virtue* of one's internalist rational support (the direct access to one's perceptual belief) that one is really justified, hence that one knows. So, contrary to our hypothesis, epistemological disjunctivism is not really pluralistic (or dualist). The only justification which counts is the one which warrants the epistemic position of knowing. This is the distinctive thesis of epistemological disjunctivism: in perception one's rational support can be both accessible by reflection alone and factive. This is quite opposed to the externalist conception of justification, according to which one can be justified—in the sense which entails knowledge—in believing that *p* without one's having access to one's reasons or justification. This is also very distinct from a commitment to a merely internalist conception of justification, since such a conception does not say that perceptual beliefs are factive. So one way or another epistemological disjunctivism cannot be pluralist about justification.

One of the main motivations Pritchard's epistemological disjunctivism is an attempt to answer the sceptical challenge, in several of its versions. In the most common one, the "closure" version, the sceptic's reasoning is the following:

(1)   I know that I'm not a brain in a vat (BIV) only if I have reflectively accessible evidence that justifies believing that I'm not a BIV.
(2)   I don't have reflectively accessible evidence that justifies believing that I'm not a BIV.
(3)   Therefore, I don't know that I'm not a BIV.

Externalists about justification do not accept (2). But internalists about justification cannot accept (2) either. The epistemological disjunctivist

on the contrary accepts that one can have reflective evidence that one is not a BIV which is factive. But this claim rests neither on the externalist nor on the internalist conception of justification. It must rest on another conception of justification—which one may call "disjunctivist"—which is such that direct access and factive belief *are* compatible. Then again this is a far cry from a view which would be epistemically pluralist. For such a view would have to say that one can *either* have internalist *or* externalist justification. This is not the kind of disjunctivism that Pritchard aims at, for what he wants to say is that epistemological disjunctivism is neither internalist nor externalist. But this is far from what Prichard advertises, when he claims that epistemological disjunctivism may achieve the "Holy Grail" of epistemology by reconciling the internalist view that justification is reflectively accessible with the externalist view that justification is objectively truth conducive. And, as I have tried here to suggest, such a view, whether or not it is true, is not pluralist at all.

Epistemological disjunctivism is only one strategy among those which Pritchard (2012) uses to rebut the sceptical challenge. The strategy that we just considered against the closure-based sceptical argument is actually a strategy which, according to Pritchard, aims at "overriding" the sceptical paradox, and which is committed to the thesis that all our beliefs are subject to rational evaluation ("the universality of reasons"). But there is, according to Pritchard, another strategy against radical scepticism, which does not consist in trying to reject the sceptic's argument by overriding it, but by trying to *undercut* the sceptic's stance. This is the Wittgensteinian strategy: rather than addressing the sceptic's denial that we have knowledge because our beliefs fail to be rationally grounded, it points out that our whole belief system rests upon hinge commitments towards certain kinds of contents which are ungrounded, hence that we cannot rebut. These are contents of the form "I have hands", "the Earth has not been created five minutes ago", "My name is NN", etc. These are neither beliefs, nor candidates to knowledge, but rather presuppositions of our whole world picture, which are taken for granted and never the object of our epistemic inquiries. No rational support or justification can be given for these, because there is no such support to be had.

The undercutting strategy is quite different from the overriding strategy: it does not rest on the idea that we can have *knowledge* of perceptual contents which can be both *factive* and *rationally supported*, but on the idea that our hinge commitments can be factive, reflectively accessible, without being rationally grounded. But they are not unjustified either. They enjoy a kind of status which is better described with terms like assumptions, constitutive rules or entitlements which govern our whole belief system.[9] Whatever name we give them, the point is that they are not in the game of justification and of knowledge. As Pritchard insists, they pertain only to *local*, but not to universal ways of giving reasons to our beliefs. If this is so, can these hinge commitments be considered as alternative ways of justifying our beliefs, and can a view which admits them be characterized as a form of epistemic pluralism? Hardly so. For if epistemic pluralism is, as I have characterized it here, the view that there are a number of alternative forms of justification, hinge commitments cannot fit the bill, since they are not epistemic *justifications*. They are neither susceptible of being true, nor of being knowledge. They are neither beliefs nor propositions, and cannot provide any "rational support". They do not play any role in our epistemological inquiries or in the "fixation of belief". They do not call for evidence, reliability or an of the ordinary notions associated to epistemic appraisal. They are, Pritchard's phrase, "arational" and one could add: anepistemic. Pritchard, when he describes this kind of answer to or "epistemic angst", talks of a "biscopic vision", engendered by his acceptance of two ways of answer the sceptic: his form of epistemological disjunctivism one the one hand and the Wittgensteinian strategy on the other (Pritchard 2015, pp. 173–179). He claims that these are not incompatible and that they "support each other". But I do not see how this is possible. If epistemological disjunctivism can be associated with an overriding strategy against the sceptic, it has to participate to the project of a rational and universal answer to the sceptical challenge. It cannot be put on a par with the local game of giving reasons. If this is a form of vision, it is more cyclopic than biscopic. This is not a form of epistemic pluralism, and it is hard to see how epistemic pluralism could be defended on such grounds.[10]

6 Two Corollaries: Pragmatic Encroachment and Epistemic Injustice

As a corollary to the previous discussion, I would like to examine, albeit briefly, two alleged phenomena which seem potentially to pose a threat to epistemic monism, pragmatic encroachment and "epistemic injustice".

I have argued that epistemic pluralism is not a real option. But this is not to deny that there can be other dimensions of evaluation of beliefs than epistemic ones. There are indeed practical reasons to believe, and various kinds of evaluation of our beliefs, ethical, aesthetic, or otherwise, which can often compete or coexist with our epistemic evaluations. We can also have reasons *not* to consider some reasons that we have, be they epistemic, practical or other.

But can we count these among our *epistemic* evaluations? Truth, evidence, reliability, justification and knowledge belong to what we take to be the *right kind* of reasons for beliefs, and all other reasons are in some way extrinsic. I shall not here go into the discussion about the right kind of reasons.[11] The pluralism issue about epistemic evaluation is orthogonal to this problem. Epistemic monism assumes that epistemic evaluation is exclusive, in other words that whether a subject is justified, or has knowledge, depends only upon epistemic considerations. This assumption is often called "purism" or "intellectualism". It is contested by "subject sensitive invariantists" who claim that a subject's practical situation can affect whether he has knowledge or not and that his epistemic position can be "pragmatically encroached" (Stanley 2005; Fantl and McGrath 2012). Unlike contextualists, subject sensitive invariantists claim that the meaning of "knows" does not change with the context of attribution, but the claim that knowledge can come and go depending on the stakes. If justification and knowledge are subject to such pragmatic encroachment, epistemic evaluation will not be exclusive, and epistemic monism will be under threat. But is the threat real?

Subject sensitive invariantists give examples like the following:

Train Case 1. You are in Boston and are about to board one of the trains that ultimately go to Providence, which is your destination. You are in no hurry but ask the passenger next to you out of curiosity whether this train is the express (and goes straight to Providence) or the regular service that first stops in Foxboro. He says the ticket agent told him that it is the regular service. You have no reason to doubt him.

Train Case 2. You are in Boston and are about to board one of the trains that ultimately go to Providence. Your destination, however, is Foxboro, and it is vitally important that you get there as soon as possible. You ask the passenger next to you whether this train is the express (which goes directly to Providence and does not stop in Foxboro) or the regular service (which does make that stop). He says the ticket agent told him that it is the regular service. You have no reason to doubt him.

Although the evidence that the subject has in both cases is the same, because the stakes are high, one is tempted to say that he does not really know in the second case. The principle which underlies such diagnoses is that if one knows that *p*, then one should be able to act as if *p*, hence that if one shouldn't act as if *p*, one does not know that *p*. This principle is supposed to follow from fallibilism, together with the fact that when the stakes are high, we are less certain, or prepared to withdraw our self-attributions of knowledge. This seems to be intuitively plausible because we tend to attribute to ourselves knowledge when we take our actions to be reasonable, and to withdraw knowledge when our actions are not reasonable. But why should we conclude this from such examples as the train cases? In both cases, the subject has the same evidence. Why couldn't we continue to possess knowledge while accepting that our action is not rational? Reed (2010) gives the following counterexample:

You are participating in a psychological study intended to measure the effect of stress on memory. The researcher asks you questions about Roman history—a subject with which you are well acquainted. For every correct answer you give, the researcher will reward you with a jelly bean; for every incorrect answer, you are punished by an extremely painful electric shock. There is neither reward nor punishment for failing to give an answer. The first question is: when was Julius Caesar born? You are confident, though not absolutely certain, that the answer is 100 BC. You also

know that, given that Caesar was born in 100 BC, the best thing to do is to provide this answer (i.e. this course of action will have the best consequences—you will be one jelly bean richer!).

In such as case, given the stakes, it would not be rational to answer the question, for the prospect of an electric shock is not outweighed by the prospect of a jelly bean, although your epistemic position is very strong, close or equal to knowledge.[12] The phenomenon of pragmatic encroachment is real—our self-ascriptions of knowledge vary with what is practically at stake—but is there any reason to conclude from this that the actual epistemic position of the subject does change? The fact that the practical situation of the subject, and his psychological reactions to changing stakes, may vary and lead him to cease to self-ascribe knowledge or justification does not entail that such knowledge and justification can come and go with the stakes and their perception.[13] In many ways, as Reed (*ibid*.) remarks, the problem of pragmatic encroachment is similar to the sceptical problem: when presented with the possibility of being trumped by an Evil Demon, the epistemic—and possibly practical—stakes become very high. But does it follow that our epistemic position is threatened? No, this does not follow. Our epistemic position remains constant, whichever theory of justification is true. If reliabilism is correct, the subject's knowledge stays the same, even when he thinks himself to be in a sceptical scenario. And if internalism is correct, the subject has access to evidence which can discard the sceptical threat. So pragmatic encroachment cannot be a reason to reject epistemic monism, which goes hand in hand with "purism".

A second source of doubt about epistemic monism comes from the phenomenon which Miranda Fricker (2007) has described under the name of *epistemic injustice*. Her account starts from Craig's analysis of knowledge, which is based on the idea that knowledge has a certain function which is to sort out good informants, through what is mostly a shared practice of testimony. If knowledge has a "point" or a role, which is mostly "practical", it should not be surprising that it should be evaluated differently from various perspectives and that it may come and go depending on how it is shared. When it is shared unevenly or badly, we are, Fricker argues, in situations of epistemic injustice. According to

her, there are two kinds of such situations: *testimonial* injustice, where a subject is denied his capacity to know through a "credibility deficit", and *hermeneutical* injustice, where a subject is prevented to know about his situation as a knower. As the cases are described, they clearly involve various situations of *social* injustice, where minorities, racial, sexual, or otherwise are denied access to knowledge, or impaired in their capacities to know. That such situations are clear cases of injustice is beyond doubt. They occur in many cases within social life. But why should they be in any specific sense *epistemic*? The uneven sharing of information, of the power to testify, of truth-telling or of spreading lies and fake news, of giving testimonies which are justified or not, bullshitting or not, the denial of credibility or power to know because one belongs to a certain group or because of one's social status, all these are widespread phenomena which, without any doubt, are causes of social injustice. This kind of injustice may be called "epistemic" when it concerns the formation, transmission and maintenance of information, of possibilities of assertion and of knowledge-status. But why should it be epistemic in the sense in which the epistemic status of the subjects could vary with their testimonial or hermeneutical status? It is injustice, but the fact that it concerns an epistemic matter does not entail that there is a special kind of epistemic status which deserves to be called "unjust" *as such*. Indeed, the subjects who suffer from epistemic injustice are impaired in their capacities as knowers and as testifiers because they have fewer possibilities of access to knowledge, just as subjects who are blind or amnesic have a cognitive deficit. But these variations in the conditions to which they have access to knowledge do not entail that the status of their beliefs with respect to justification or knowledge change. In other words, the concept of epistemic injustice rests on the following inference:

If an epistemic practice (such as testimony) rests on conditions of acquisition which are unjust then it is a case of epistemic injustice.

But this inference is not valid. For what is unjust in such cases is not the evaluation of testimonial beliefs, which remain true or false, justified or not, reliable or not, warranted or not, evidentially based or not while their conditions of transmission change. The social circumstances of the production

and of the reception of knowledge may affect the quantity, the quality and the value that is granted to testimonies, but it does not affect the nature of what it takes for a belief to be knowledge. Fricker recognizes this:

> Clearly credibility deficit can constitute such a wrong, but while credibility excess may (unusually) be disadvantageous in various ways, it does not undermine, insult, or otherwise withhold a proper respect for the speaker qua subject of knowledge; so in itself it does her no epistemic injustice, and a fortiori no testimonial injustice. (Fricker 2007, p. 20)

The status of someone as a subject of knowledge does not change: it is the consequences of his social injustice which impairs this status. Epistemic justice or injustice does not constitute epistemic warrant or non-warrant. The upshot is that there is no distinctive *epistemic* property which deserves the status of *epistemic* injustice, although it certainly deserves the status of injustice.

7   Conclusion: Monism After All

I have given a few reasons to resist the powerful trend towards epistemic pluralism within contemporary epistemology, by trying to specify what it entails and by giving some reasons to reject it. Epistemic monism is the view that whatever knowledge and justification can be, these notions cannot designate properties which are independent from truth, evidence and reliability, and which do not leave room to distinct and conflicting kind of evaluations. This is not to say that epistemic evaluation operates in the void, as if the truth and knowledge tracking properties could be abstracted from the many determinants of epistemic rationality which weigh upon our knowledge ascriptions. Social, pragmatic, contextual, psychological and historical factors permeate our judgements about knowledge. Our assessments are in most cases relative to these factors. But the validity of our epistemological judgements is not relative. It is always tied to a small and invariant core of truth and evidence related epistemic properties.

Notes

1. This is not always so clear in Sosa's writings. Sometimes he suggests that animal knowledge can by itself count as knowledge. But at other times, he suggests that knowledge is achieved only at the reflective level. See below.

2. Smithies (2015) argues that "the point and purpose of using the concept of justification in epistemic evaluation is tied to its role in the practice of critical reflection".

3. The functionalism about justification sketched here should not be confounded with the view defended by Michel Bergman (2006) under the name of "proper functionalism" about justification, which is a version of reliabilism.

4. See e.g. Skorupski (2011) on the frontiers of the "domain of reasons".

5. Berneker (2006) objects to Sosa that the duality of animal and reflective knowledge mirrors the duality of justification and knowledge which affects Gettierised subjects, whose beliefs are internally justified but fail externally to be knowledge. But Sosa's two-tiered account is on the contrary meant to explain why Gettierised subjects do know for they failed to have a reflective appreciation of their otherwise competent belief. See Sosa (2015).

6. I concur in this with Peels (2010).

7. This is not the place here to rehearse the arguments. See e.g. Engel (2013a).

8. For an analysis of Price's argument, see Engel (2013b).

9. Indeed, the terminology varies in the literature. See, for an overview of the various conceptions of such commitments, Coliva (2014).

10. In Engel (2007), I have defended a kind of « biscopic » vision, arguing against scepticism from externalist premises, supplemented by an entitlement account. But the two seem to me now incompatible. In Engel (2012), I argue that these cannot be held together.

11. See e.g. Hieronymi (2005), Skorupski (2011), Engel (2007).

12. For a similar reaction to such examples, see Engel (2009).

13. For a plausible psychological explanation in terms of « epistemic anxiety » , see Nagel (2008).

References

Alston, W. 2005. *Beyond Justification*. Cornell: Cornell University Press.

Audi, R. 1988. *Belief, Justification, and Knowledge*. Belmont: Wadsworth.

Berneker, S. 2006. Prospects for Epistemic Compatibilism. *Philosophical Studies* 130: 81–104.

Bergmann, M. 2006. *Justification Without Awareness: A Defense of Epistemic Externalism*. Oxford: Oxford University Press.

Boghossian, P. 2005. *Fear of Knowledge*. Oxford: Oxford University Press.

Bonjour, L. 1985. *The Structure of Empirical Knowledge*. Cambridge: Harvard University Press.

———. 1988. Externalist Theories of Empirical Knowledge. *Midwest Studies in Philosophy* 5: 53–73.

Coliva, A. 2014. *Extended Rationality. A Hinge Epistemology*. London: Palgrave.

———. 2016. *The Varieties of Self Knowledge*. London: Palgrave Macmillan.

Craig, E. 1991. *Knowledge and the State of Nature*. Oxford: Oxford University Press.

Elgin, C. 1988. The Epistemic Efficacy of Stupidity. *Synthese* 74: 297–311.

Engel, P. 2007. *Va savoir ! De la connaissance en général*. Paris: Hermann.

———. 2009. Pragmatic Encroachment and epistemic value. In *Epistemic Value*, ed. A. Haddock, A. Millar, and D. Pritchard, 183–203. Oxford: Oxford University Press.

———. 2012. Knowledge and Reason. In *Reason and Rationality*, ed. C. Amoretti and N. Vassallo, 53–70. Frankfurt: Ontos.

———. 2013a. In Defense of Normativism about the Aim of Belief. In *The Aim of Belief*, ed. T. Chan, 32–63. Oxford: Oxford University Press.

———. 2013b. Alethic Functionalism and the Norm of Truth. In *Truth and Pluralism: Current Debates*, ed. N. Pedersen and C.D. Wright. Oxford: Oxford University Press, 69–86.

———. 2013c. Belief and the Right Kind of Reason. *Teorema* 32 (3): 19–34.

———. 2016. Epistemic Norms and the Limits of Epistemology. *International Journal for the Study of Skepticism* 6 (2–3): 228–247.

Fantl, J., and M. McGrath. 2012. *Knowledge in an Uncertain World*. Oxford: Oxford University Press.

Fogelin, R. 1994. *Pyrrhonian Reflections on Knowledge and Justification*. Oxford: Oxford University Press.

Fricker, M. 2007. *Epistemic Injustice*. Oxford: Oxford University Press.

Goldman, A. 1991. Strong and Weak Justification. In *Liaisons: Philosophy Meets the Cognitive and Social Sciences*. Cambridge: MIT Press.

Harman, G. 1973. *Thought*. Princeton: Princeton University Press.

Hieronymi, P. 2005. The Wrong Kind of Reason. *Journal of Philosophy* 102 (9): 437–457.

Kvanvig, J. 2003. *The Pursuit of Knowledge and the Goal of Understanding*. Cambridge: Cambridge University Press.

Lehrer, K. 1990. *Theory of Knowledge*. Boulder: Westview.

Lynch, M. (ed.). 2001. *The Nature of Truth*. Cambridge: MIT Press.

———. 2009. *Truth as One and Many*. Oxford: Oxford University Press.

Nagel, J. 2008. Knowledge Ascriptions and the Psychological Consequences of Changing Stakes. *Australasian Journal of Philosophy* 86: 279–294.

Peels, R. 2010. Epistemic Desiderata and Epistemic Pluralism. *Journal of Philosophical Research* 35: 193–206.

Plantinga, A. 1992. *Warrant: The Current Debate*. Oxford: Oxford University Press.

Price, H. 1998. Three Norms of Assertibility. *Philosophical Perspectives* 12, ed. J. Tomberlin, 241–254.

———. 2003. Truth as Convenient Friction. *Journal of Philosophy* 100 (4): 167–190.

Pritchard, D. 2012. *Epistemological Disjunctivism*. Oxford: Oxford University Press.

———. 2015. *Epistemic Angst*. Princeton: Princeton University Press.

Reed, B. 2010. A Defense of Stable Invariantism. *Nous* 44 (2): 224–244.

Rorty, R. 1995. Is Truth a Goal of Enquiry? Donald Davidson Versus Crispin Wright. *Philosophical Quarterly* 45: 281–200; reprinted in *Truth and Progress*: *Philosophical Papers* vol. 3. CUP (1998): 19–42.

Shah, N. 2003. How Truth Governs Belief. *Philosophical Review* 12 (4): 447–482.

Skorupski, J. 2011. *The Domain of Reasons*. Oxford: Oxford University Press.

Smithies, D. 2015. Why justification Matters, In *Epistemic Evaluation, Point and Purpose in Epistemology*, eds. David K. Henderson and John Greco, 224–244. Oxford University Press.

Sosa, E. 2009. *Reflective knowledge*. Oxford: Oxford University Press.

———. 2011. *Knowing Full Well*. Princeton: Princeton University Press.

Stanley, J. 2005. *Knowledge and Practical Interests*. Oxford: Oxford University Press.

Sutton J. 2007. *Without Justification*. Cambridge Massachusetts: MIT Press.

Weiner, M. 2005. Must We Know What We Say? *Philosophical Review* 114 (2): 227–251.

Williamson, T. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.

Wright, C.J.G. 1992. *Truth and Objectivity*. Cambridge Massachusetts: Harvard University Press.

———. 2001. Minimalism, Deflationism, Pragmatism, Pluralism. In *The Nature of Truth: Classic and Contemporary Perspectives*, ed. M. Lynch. 751–787. Cambridge: MIT Press.

Part II

Epistemic Pluralism: Justification, Knowledge, and Epistemic Values

The Pluralism of Justification

Anne Meylan

1   What Is a Justified Belief? The Higher-Order and the Basis Question

Beliefs—like the belief that your neighbour hates you, that your cat is sick, etc.—are sometimes *justified*. But what is a justified belief? Famously enough, this question has intensely preoccupied epistemologists. Despite its title, the first part of this contribution should not be considered as a simple extension of the already existing debate regarding beliefs' justification. Here is why.

When a belief is justified, it is justified in virtue of other specific properties of the belief. This is the case for normative[1] properties in general. Consider, for instance, the following normative properties: being disgusting, being elegant and being permitted. These are properties that objects, persons, states of affairs, etc., possess in virtue of

A. Meylan (✉)
Department of Philosophy, University of Basel, 4051 Basel, Switzerland
e-mail: anne.meylan@unibas.ch

their exemplifying other specific properties.[2] A dish may be disgusting in virtue of its aspect, taste, consistency, etc. A person may be elegant in virtue of the clothes she wears, the way she walks, etc. An action may be permitted in virtue of certain features, which make it such that it does not violate any obligations.

> More generally:
>
> For all entities x that have the property N (where "N" denotes a normative property), x is N in virtue of its exemplifying one[3] specific N-making property.[4]

What is true about normative properties in general is true in the specific case of justification given that—as will be emphasized below—justification is a normative property.

> Every time an entity is justified, it is justified in virtue of its exemplifying a specific justified-making property.[5]

What this principle says is that, every time a belief is justified, it necessarily exemplifies two properties:

1. the property of being justified and
2. a justified-making property.

Accordingly, when one asks (as philosophers often do): "what is a justified belief?", one may actually be raising two different questions depending on whether one is interested in the first or in the second of these two properties.

One may be asking either:

1. What is justification?
   or
2. Which property is the justified-making property?

I shall call the first question "the higher-order question" and the second question "the basis question".

The distinction between these two questions is only rarely made explicit in the literature.[6] But, clearly, epistemologists have mostly focussed their attention on the second question.[7] Their project has mainly been to determine what the justified-making properties—the properties in virtue of which a belief is justified—are. A plausible explanation of this preference is not hard to find. The prevailing interest of some epistemologists in the basis question might well result from the fact that they *already* have an answer to the higher-order question. This answer is the following: justification is the property whose exemplification turns un-Gettierized true beliefs into instances of knowledge. I have previously warned that the first part of this contribution should not simply be considered as an extension of the vast epistemological debate pertaining to beliefs' justification. The first part of this contribution differs from this more familiar way of discussing justification in the following respect: it focuses on the higher-order question (what is justification?) and not on the basis question (which property is the justified-making property?).

The Key Features of Justification

Just as we can give an informative answer to the question: "who is Petra?" by enumerating several of Petra's significant features (e.g. she grew up in India, she studied logic then law, she practices yoga), we can give an informative answer to the higher-order question by listing significant features of justification. Here I will briefly discuss two of these key features that are useful for my purposes in this paper.[8] They are:

1. the normative feature of justification;
2. the trans-categorial feature of justification.

Epistemologists frequently attribute (most often in passing) the following feature to justification:

Justification is a normative property.

This is, what I have just called, *the normative feature of justification*. Even epistemologists who deeply diverge regarding other matters agree on this part of the answer to the higher-order question.

> The term 'justify', in its application to a belief, is used as a term of epistemic appraisal (Chisholm 1977, Chap. "Varieties of epistemic pluralism", p. 6)

> The term 'justified', I presume, is an evaluative term, a term of appraisal. (Goldman 1979, p.1)

> To say that S is justified in believing that p is to imply that there is something all right, satisfactory, in accord with the way things should be, about the fact that S believes that It is to accord S's believing a positive evaluative status. (Alston 1985, p. 58)

> Epistemic justification is a normative or evaluative property. (Sosa, 1991a, p. 86)

> Justification is one among many dimensions of epistemic evaluations. (Smithies 2015, p. 225)

As we saw already, if justification is a normative property then every time a belief is justified, it is justified in virtue of the exemplification of a specific justified-making property. But the exemplification of justification is not restricted to beliefs. Just like beliefs, actions—for example, the action of breaking off diplomatic relations with a particular country, of quitting your job, of lying to your kid, etc.—are justified in certain circumstances. Justification is exemplified by various kinds of entities, including beliefs, actions and—according to some—emotions.[9] In this respect, justification is like the property of being inspiring but unlike the property of being true.[10] Let me call this feature of justification the "trans-categorial feature".

The trans-categorial feature of justification
Justification is a property that is exemplified by beliefs and actions alike when they are justified.

The trans-categorial feature of justification is not a claim about what it is appropriate to say. The idea is not simply that it is appropriate to predicate "justified" of actions as well as beliefs. The claim is rather that a justified action and a justified belief exemplify the same property: justification. Or, to formulate the idea on another level, "justified" refers to the very same property whether it is predicated of an action or predicated of a belief.[11] There is no more reason to think that "justified" refers to two different properties depending on whether it is predicated of an action or of a belief than there are reasons to think that "blue" refers to two different properties depending on whether it is predicated of a shirt or of the sky. The correct way of understanding "the sky is blue" and "this shirt is blue" is to take these two sentences to mean that the shirt and the sky have at least something (a property) in common. They both exemplify the property of being blue (even though they might differ with regard to their more precise blueness). The same is true as regards "justified". The correct way of understanding "this action is justified" and "this belief is justified" is to take these two sentences to mean that the belief and the action have a basic property in common. They are both justified.

The goal of the following Sects. (2 and 3) is to exploit the aforementioned trans-categorial feature of justification in order to show that the properties to which many contemporary accounts of justification refer with the predicate "justified" are distinct technical properties. I shall more precisely proceed as follows. I first try to show that the classical reliabilist account of justification entails that justification is a technical property. The claim to the effect that this is also true of accessibilism and of other accounts of justification is defended in a second time.

2   Classical Reliabilism and the Trans-Categorial Feature

In classical reliabilism, the justified-making property is the property of being produced by a reliable process, viz. by a process that generally produces true beliefs.[12] Let us consider this claim in the light of what has been previously stated. Recall that:

> Every time an entity is justified, it is justified in virtue of the exemplification of a specific justified-making property.

When combined with the reliabilist claim, this principle has the following (welcome) implication for reliabilism:

> Every time a belief is justified, it is justified in virtue of its exemplifying the property of being produced by a reliable process.

Now recall the trans-categorial feature of justification, viz. the idea that justification is a property that is exemplified by beliefs and actions alike when they are justified. Given the trans-categorial feature, what is true about justification when exemplified by beliefs should also be true about justification when exemplified by actions. As we have just seen, one implication of reliabilism is that every time justification is exemplified by a belief, it is exemplified in virtue of its exemplifying the property of being produced by a reliable process. Given the trans-categorial feature, this should also be true about justification when exemplified by actions. But is it the case? Is it true that every time the property of being justified is exemplified by an action, it is exemplified in virtue of the fact that this action exemplifies the property of being produced by a reliable process? Clearly, the answer is "no".[13] When one claims: "it is justified to quit your job in these circumstances", there are various imaginable justified-making properties in virtue of which the action of quitting your job could count as justified. But, obviously, when quitting your job is justified, it is not justified in virtue of its being produced by a process which generally produces true beliefs. An actions' justification does not have anything to do with the truth-conduciveness of the processes that produce it.

You might object that this is because speaking in terms of "truth-conduciveness" is appropriate for beliefs but not for actions. The reliabilist conception of the justified-making property—the objection goes—has to be adapted in order to suit actions, for instance, by stating that the justified-making property (in the case of actions) is the property of being produced by a process that generally produces pleasant (useful, beneficial, etc.[14]) outcomes.

But this does not work either. Suppose that it is justified to quit your job in some circumstances. What makes this action justified, viz. the justified-making property, cannot be identified with the property of being produced by a process that is likely to produce pleasant (useful, beneficial, etc.) results. Indeed, suppose you have been caused to quit your job by a tendency to under evaluate your capacities (a mental process that one might say has generally unbeneficial outcomes). This does not seem to necessarily make your action of quitting your job something that is not justified. The action of quitting your job can be justified even if it has been caused by a process which generally has unpleasant (useless, unbeneficial, etc.) outcomes, in this case, by the mental process of systematically underestimating your capacities. Put differently, whether an action is justified or not does not depend on the instrumental value of the process that causes it.

Let me clarify my purpose here. My intention is to show that the property to which classical reliabilism refers with the term "justification" does *not* have the trans-categorial feature. But this result is not meant—and cannot be meant—to raise an objection to classical reliabilism. The reason why it cannot be properly intended to constitute an objection to reliabilism is simply that nothing forces reliabilists to accept that the property to which they refer with the predicate "justified"—when they claim, for instance, that what makes a belief *justified* is the property of being produced by a reliable process—has the trans-categorial feature. Indeed, when formulating their views, philosophers are allowed to refer to the things that they have in mind with the terms they think appropriate. In view of that, it is perfectly acceptable to use the predicate "justified" to refer to the property that a belief exemplifies when produced by a reliable process even if this property has nothing to do with the justification that actions exemplify when *they* are justified.

Plantinga formulates the idea quite explicitly[15]:

> There is wide…. agreement that true belief is necessary but not sufficient for knowledge. But then what more is required? What is this quantity enough of which (Gettier problems, perhaps, aside) epistemizes true belief? … Perhaps the natural procedure would be to *baptize* the element, whatever it is, 'epistemic justification'. (Plantinga 1988, pp. 2–3)

Thus, to repeat, reliabilists are perfectly within their rights to use the predicate "justified" to refer to the property that a belief possesses when produced by a reliable process without accepting that the property to which they refer with this predicate also has the trans-categorial feature. However—and this is a very important remark for my purpose in this paper—to use the predicate "justified" in this way amounts to using it in a *technical* way, that is to say, to using it in order to denote some technical property called "justification". This is (perhaps unfortunately) only very rarely stated explicitly by epistemologists. One exception is BonJour:

> The specifically epistemological notion of justification is to a significant extent a technical philosophical notion, one that is not clearly and unquestionably present in common sense. (BonJour 2007, p. 36)

I have just claimed that using "justified" in order to refer to a property that does not have the trans-categorial feature (as reliabilism uses it) boils down to a *technical* use of the predicate. I turn now to the vindication of this last claim.

3   The Ordinary Vs. the Technical Use of the Predicate "Justified"

As already alluded to, uses of the predicate "justified" are not restricted to epistemological discourse. The predicate "justified" appears in everyday language as well. More precisely, when used in ordinary or in everyday language, "justified" (or "unjustified") is predicated of actions, like, for instance, in the following sentences: "quitting your job was justified in these circumstances", "slapping him was perfectly unjustified" and "the use of military force is not always justified". In these ordinary sentences, "justified" obviously means something. This meaning is the *ordinary* meaning of the predicate "justified". The ordinary meaning of the predicate "justified" is the one "justified" has when it is predicated of actions in ordinary sentences like the aforementioned ones. To use

the predicate "justified" in the *ordinary* way is simply to use it in such a way that it means the same as when predicated of actions in these ordinary sentences. And, on the other hand, to use the predicate "justified" in the *technical* (non-ordinary) way is to use it in such a way that it does not mean the same as it means in these ordinary sentences.

The reason why using the predicate "justified" in order to refer to a property that does not have the trans-categorial feature (as, we saw, classical reliabilism uses it) amounts to a *technical* use of the predicate follows from these considerations. You cannot use the predicate "justified" in order to refer to a property that does not have the trans-categorial feature while using the predicate in the ordinary way, viz. while using it in such a way that "justified" means the same as when it is predicated of actions in ordinary sentences. Indeed, to use the predicate "justified" in order to refer to a property that does *not* have the trans-categorial feature (in the way reliabilism uses it) is to use it in order to refer to a property that is necessarily *not* the one that actions exemplify when they are justified. And to use the predicate "justified" in order to refer to a property that is *not* the one that actions exemplify when they are justified is *ipso facto* not to use the predicate in the ordinary way. When "justified" is used in order to refer to something that cannot be identified with the property that actions exemplify when they are justified, the use of this predicate is *ipso facto* not ordinary. "Justified" is *ipso facto* not used in such a way that it means the same as it means when it is predicated of actions in ordinary sentences.

To recap, the explanation why the reliabilist use of the predicate "justified" is in fact a *technical* one is the following:

1. The property to which reliabilists refer when using the predicate "justified" does not have the trans-categorial feature (defended in Sect. 2);
2. To use the predicate "justified" in the ordinary way is to use it in such a way that "justified" means the same as it means when predicated of actions in sentences of everyday language like: "it was justified to quit your job in these circumstances" (definition of ordinary/technical use presented in Sect. 2.1);

3. To use the predicate "justified" in order to refer to a property that does not have the trans-categorial feature is just to use it in order to refer to a property that differs from the property exemplified by actions when they are justified (definition of what it is for justification to have the trans-categorial feature presented in Sect. 1.1);
4. The property to which reliabilists refer with the predicate "justified" is not the same property as that which is exemplified by actions when they are justified (1, 3);
5. To use the predicate "justified" in order to refer to a property that differs from the property exemplified by actions when they are justified is to use the predicate in a non-ordinary way (given what it is to use the predicate in the ordinary way, see premise 2);
6. To use a predicate in a non-ordinary way is to use it in a technical way;
7. Reliabilists use the predicate "justified" in a technical way (4, 5, 6).

The same reasoning applies to every account of beliefs' justification for which premise 1 holds. And, probably, many accounts of beliefs' justification are such that the property that they denote with the predicate "justified" does not have the trans-categorial feature. It would obviously be too tedious to work through such a demonstration. I therefore leave it to the reader to consider whether her favourite account of beliefs' justification is such that the "justification-making" property for beliefs is also the "justification-making" property for actions.

If this is not the case—as I suspect—the property that is denoted by the predicate "justified" in whatever account one favours does not have the trans-categorial feature. This means that premise 1 holds for this account and that the rest of the reasoning applies.

Now, I would like to show explicitly that the aforementioned reasoning does not only concern externalist accounts of justification. Even internalist accounts of justification are such that the property to which they refer with the predicate "justified" does not have the trans-categorial feature. Consider, for instance, a paradigmatic form of internalism: accessibilism. According to the accessibilist conception of justification, a subject S's belief is that is justified in virtue of *S having some cognitive access (mere non-conceptual awareness might be sufficient)*

of a fact (or a mental state) that speaks in favour of the truth of the belief that p.[16] The whole italicized expression is the rather complex justified-making property according to such an accessibilist account of justification. Now consider, once again, the action of quitting your job and suppose that it is justified. What makes this action a justified one is not the fact that you have access to certain facts (or mental states) that speak in favour of its rightness.[17] Quitting your job might very well be the action to perform in these circumstances even though you do not have any access to the facts (or the mental states) that justify it. In other words, the accessibilist's complex property is not a property the exemplification of which makes actions justified. The property that is denoted by the predicate "justified" in such an accessibilist account does not have the trans-categorial feature. The aforementioned reasoning applies to accessibilism as much as it applies to classical reliabilism, and the conclusion is, therefore, the same. Accessibilism and classical reliabilism use the predicate "justified" in a technical way, i.e. in order to refer to distinct technical properties that both theories call "justification". And this is true of every account of justification for which premise 1 holds. According to the aforementioned reasoning, each account of justification for which premise 1 holds—and there are probably many—refers to a distinct technical property that they all call "justification". This is a form of *technical pluralism* because it results from the fact that the ordinary meaning of the predicate "justified"—i.e. the meaning that the predicate "justified" has in sentences like "quitting your job is justified in such circumstances"—is not captured by these various accounts of beliefs' justification.

To give credit where credit is due, it must be said that William Alston somewhat anticipated this conclusion.

> There does not seem to be enough commonality in their (the distinct epistemologists, NDA) pre-theoretical understanding of the nature of epistemic justification to warrant us in supposing that there is some uniquely identifiable item about which they hold different views. It seems, rather, that they are highlighting, emphasizing, 'pushing' different concepts, all called 'justification'. (Alston 1993, p. 534)

Alston did not, however, provide a precise rationale supporting his claim. Providing such a rationale is what I have endeavoured to achieve in this paper.

4   Conclusion

Epistemologists have long been concerned with the question: "what is a justified belief?" The first part of the present paper (Sect. 1) spells out a distinction between two ways of addressing this question. We can understand it as asking: "what is justification?" (the higher-order question) or as asking: "which property is the justified-making property?" (the question of the basis). This first part also presents a partial answer to the higher-order question: justification is a property that is characterized by a normative and, more crucially, a trans-categorial feature. The second part of this article (Sects. 2 and 3) exploits the results of the first part in order to show that classical reliabilism, accessibilism and presumably other accounts of justification use the predicate "justified" in various technical ways.

Perhaps some readers' reaction to this conclusion will be the following: "I have always thought this". I do not think that this reaction reduces the importance of this paper. First, the claim that reliabilism and accessibilism use the predicate "justified" in a technical way is sufficiently central to be worthy of explicit mention. Second, this article does not just claim that reliabilism and accessibilism use the predicate "justified" in distinct technical ways. It also vindicates this claim, and the careful vindication of this claim turns out to be a complex matter.

Notes

1. In this article, "normative properties" refer to the *general* class, which has deontic properties, axiological properties, etc., as its members. See Mulligan (2009) for a similar use.
2. The relation between a normative property and the properties in virtue of which a thing possesses that normative property is very often

considered to be a relation of *supervenience*. See Lemos (1994); Sosa (1991a); Mulligan (2009). This also seems to be Moore's view (1993, 1912) even though Moore never makes use of the term "supervenience". The remaining question—which still divides philosophers—is whether the relation of supervenience is a reductive one or not, that is, whether normative properties can be reduced to the non-normative ones or not. I leave this problem aside here.

3. There might be *several* specific N-making properties. For sake of brevity, I omit to mention this possibility here and further on.

4. A common instance of this principle is the following: when a performance is a right performance, the performance is right in virtue of its exemplifying right-making properties, see Brandt (1985); Dancy (2005); Goldman (1979); Ross (1930, Chap. 1).

5. See Sosa (1991b, p. 110) for the same idea.

6. Goldman is an exception since he explicitly distinguishes between these two questions in his study (Goldman 1979).

7. It is rare for epistemologists to explicitly tackle the higher-order question. The exception is , most strikingly, Alston. See Alston (1985, 1993, 2006) but also Chisholm (1977, Chap. 1), Bedford Naylor (1988) and Brandt (1985).

8. Another significant feature of justification is its gradability. Justification is a property that things can exemplify to various degrees. Note that this is something that brings justification closer to value properties and makes it drastically different from deontological properties. Just as a painting can be more or less beautiful, a belief can be more or less justified. By contrast, an action cannot be more or less obligatory or permitted.

9. On emotions' justification, see Deonna & Teroni (2012).

10. There are correct actions, for instance, morally correct ones. But there is, *stricto* sensu, no true action. To put it roughly, actions do not "represent" and this prevents them from being *stricto* sensu assessed as true or false.

11. For a similar view regarding goodness, see Moore (1993, p. 54).

12. See Goldman (1979). Since Olsson's form of reliabilism is not expressed in terms of "justification" (see Olsson (2007), 2011), his view is not covered by the following remark.

13. The problem is *not* that actions are incapable of exemplifying the justified-making property. Actions can be produced by reliable processes,

processes that generally produce true beliefs (for instance, perceptual processes).

14. Any practically valuable property will do the job here.

15. Plantinga's final view is that the term "justification" has too much of a deontological connotation in order to properly refer to "this quantity enough of which (Gettier problems, perhaps, aside) epistemizes true belief". Plantinga prefers the term "warrant" instead. Alston partially agrees. According to Alston (1985, n. 21 and 1993, p. 533), it would be better to abandon the term "justification" because of its deontological connotation. Alston's final decision is, however, to keep using this term for the reason that it is "too firmly rooted" in philosophical habits.

16. See Chisholm (1977) or Ginet (1975) for traditional accessibilist conceptions of justification.

17. I moved from truth to rightness here because it is simply obvious that the justification of an action is not a matter of its truth.

References

Alston, W. 1985. Concepts of Epistemic Justification. *The Monist* 68: 57–89.

———. 1993. Epistemic Desiderata. *Philosophy and Phenomenological Research* 53 (3): 527–551.

———. 2006. *Beyond "Justification" Dimensions of Epistemic Evaluation.* Ithaca/London: Cornell University Press.

Bedford Naylor, M. 1988. Epistemic Justification. *American Philosophical Quarterly* 25 (1): 49–58.

Brandt, R.B. 1985. The Concept of Rational Belief. *The Monist* 68 (1): 3–23.

BonJour, L., and E. Sosa. 2007. *Epistemic Justificstion: Internalism vs. Externalism, Foundations vs. Virtues.* Oxford: Blackwell.

Chisholm, R. 1977. *Theory of Knowledge* 2nd edition. Englewood Cliffs, New Jersey: Prentice-Hall.

Dancy, J. 2005. Should We Pass the Buck. *Recent Work on Intrinsic Value*, eds. T. Rønnow-Rasmussen and M. Zimmerman. Library of Ethics and Applied Philosophy 17. Dordrecht: Springer.

Deonna, J., and F. Teroni. 2012. From Justified Emotions to Justified Evaluative Judgements. *Dialogue* 51 (01): 55–77.

Ginet, C. 1975. *Knowledge Perception, and Memory*. Dordrecht: D. Reidel Publishing Company.

Goldman, A. 1979. What is Justified Belief? In *Justification and Knowledge*, ed. G. Pappas, 1–23. Boston: D. Reidel Publishing Company.

Lemos, N. 1994. *Intrinsic Value: Concept and Warrant*. Cambridge: Cambridge University Press.

Moore, G.E. 1912. *Ethics*. London: Williams & Norgate.

———. 1993. *Principia Ethica*. revised edition, ed. T. Baldwin. Cambridge: Cambridge University Press.

Mulligan, K. 2009. Values. *The Routledge Companion to Metaphysics*, eds. R. Poidevin, P. Simons, A. McGonigal and R. Cameron, 401–411. London: Routledge.

Olsson, E. 2007. Reliabilism, Stability, and the Value of Knowledge. *American Philosophical Quaterly* 44: 343–355.

———. 2011. The Value of knowldege . *Philosophy Compass* 6 (12): 874–883.

Plantinga, A. 1988. Positive Epistemic Status and Proper Function. *Philosophical Perspectives* 2: 1–50.

Ross, W.D. 1930. *The Right and the Good*. Oxford: Oxford University Press.

Smithies, D. 2015. Why Matters. *Epistemic Evaluation: Point and Purpose in Epistemology*, eds. D. Henderson and J. Greco. Oxford: Oxford University Press.

Sosa E. 1991a. Nature Unmirrored, Epistemology Naturalized. In *Knowledge in Perspective: Selected Essays in Epistemology*, 86–107. Cambridge: Cambridge University Press.

———. 1991b. Theories of Justification: Old Doctrines Newly Defended. In *Knowledge in Perspective: Selected Essays in Epistemology*, 108–130. Cambridge: Cambridge University Press.

Knowledge Second (for Metaphilosophy)

Jennifer Nado

The concept of knowledge occupies a central role in epistemology—so much is obvious. But knowledge and its compatriots—evidence, justification, warrant, and the myriad of other concepts intimately connected with knowledge—also dominate the comparatively recent literature surrounding the use of intuitions in philosophy, and surrounding the philosophical consequences of experimental philosophy. Indeed, both experimentalist critics of intuition and their opponents have a tendency to state their conclusions using the standard epistemic concepts of mainstream analytic epistemology. This has, quite unfortunately, led to a misleading picture of experimentally based criticisms of intuition as advocating 'intuition skepticism'—in other words, as denying that intuitions are a source of knowledge. This, in turn, has led to a metaphilosophical literature dominated by debates over whether such a complete rejection of the evidential value of intuition is reasonable, or even possible.

J. Nado (✉)
Department of Philosophy, University of Hong Kong,
Run Run Shaw Tower, Centennial Campus, Pokfulam Road,
Room 10.13, 10/F, Hong Kong, China
e-mail: jennifernado@gmail.com

But concepts like knowledge, I'll argue, are in fact ill-suited for purposes of an evaluation of philosophical methodology. This is because the epistemic standards that govern inquiry in philosophy are separate from, and plausibly much more stringent than, the standards one must meet in order to know. This is not intended as a contextualist claim; the claim is not a semantic one, but one concerning the epistemic norms governing philosophical inquiry. The existence of these separate norms, in turn, invites a pluralistic approach to epistemic states. We might, for instance, employ counterpart terms—evidence$_p$, justification$_p$, knowledge$_p$, and the like—to refer to the epistemic states that result from meeting the epistemic standards generated by philosophy-specific norms.[1] We can then note that the fact that a given method or process generates evidence, justification, or knowledge does not guarantee that it will generate evidence$_p$, justification$_p$, or knowledge$_p$. A proposition might well be known without being known$_p$; it might be evidence without being evidence$_p$.

The appropriate conclusion to draw from experimental philosophy studies, I'll claim, is best stated in terms of philosophy-specific epistemic categories like these, rather than the categories familiar from standard analytic epistemology; experimental philosophy, then, would benefit from a pluralistic conception of epistemic states. Given such a conception, we can then gloss the experimentalist conclusion as follows: Though intuitions may produce knowledge, experimental findings indicate biases and epistemic flaws that suggest that our current use of intuition may not produce knowledge$_p$.

This, I would suggest, is a much less problematic stance than 'intuition skepticism.' The restated experimentalist challenge I offer can be viewed as analogous to the uncontroversial claim that the biases and epistemic flaws that afflict ordinary, uncontrolled observation render it insufficiently reliable for use in an experimental context: thus making epistemically obligatory, within the sciences, the use of, e.g., measuring instruments and procedures such as double-blinding. It's obvious that such claims can be made without threatening skepticism—*mutatis mutandis* for philosophy-specific criticisms of the use of intuition.

The reorientation of our metaphilosophical framework that I'll be suggesting implies that experimentalists and defenders of intuition

are, in fact, frequently talking past one another; while defenders busy themselves with arguments about intuition's ability to justify belief or to generate knowledge, such defenses fail to address the real challenge that experimental findings potentially present to our current methodological practices. Ultimately, all this points to a more general moral: The monolithic focus on 'knowledge' as the primary epistemic state of interest in philosophy obscures many of the subtleties involved in evaluating our multifaceted practices of inquiry, both in philosophy and in other fields. There exist many different epistemic standings of interest; it is, I think, a losing game to attempt to capture all these with any reasonably unified account of knowledge. Worse, it is a distraction from the more central question of how to best investigate reality.

1   The Skepticism Problem

Elizabeth O'Neill and Edouard Machery, in their introduction to a recent volume on experimental philosophy, present the experimentalist argument against intuition as follows:

1. The judgments elicited by philosophical thought experiments are significantly influenced by factors that do not track the fact of the matter.

2. If a judgment is significantly influenced by factors that do not track the fact of the matter, that judgment is not reliable.

3. If a judgment is not reliable, it cannot provide warrant for assuming its content.

4. Hence, the judgments elicited by thought experiments do not provide warrant for assuming their content. (Machery and O'Neill 2014, p. xvi)

The first premise of this argument is, of course, meant to be supported by the various findings generated by experimental studies on intuition—findings that suggest that intuitions may be subject to influence by cultural background (Weinberg et al. 2001; Machery et al.

2004), emotion (Wheatley and Haidt 2005; Schnall et al. 2008), order (Swain et al. 2008), personality (Feltz and Cokely 2008, 2009), gender (Buckwalter and Stich 2014), and so on. One could object to this premise on grounds of methodological flaws in the experiments (see, e.g., Kauppinen 2007; Cullen 2010; Woolfolk 2013), or by noting that several of the aforementioned studies have failed to replicate (Nagel et al. 2013; Adleberg et al. 2015; Seyedsayamdost 2015). But set such issues aside for present purposes—though the current body of experimental studies no doubt has its fair share of flaws, it's reasonably plausible that a more mature experimental research program may one day fully vindicate premise 1.

More important, for our purposes, is the dramatic claim expressed by the conclusion—that the epistemic flaws exhibited by thought experiment judgments entail that the judgments provide *no warrant*. This characterization of the experimentalist argument is in no way atypical. Consider the following statements by members of the experimentalist camp:

> Sensitivity to irrelevant factors undermines intuitions' status as evidence (Swain et al. 2007, p. 141).
>
> Experimental philosophy challenges the usefulness of [appealing to intuition] in achieving justified beliefs. (Alexander et al. 2010, p. 298)
>
> The ultimate hope is that we can use [experimental philosophy findings] to help determine whether the psychological sources of the beliefs undercut the warrant for the beliefs. (Knobe and Nichols 2008, p. 7)
>
> [Susceptibility to bias gives us] good reason to think that intuitions are unreliable. (Tobia et al. 2013, p. 631)
>
> [Experimental findings are] bad news for the standard philosophical assumption that the contents of people's intuitions are very likely to be true. (Stich and Tobia 2016, p. 13)

Such statements are quite strong. The first three explicitly deny that intuition produces the standard 'epistemic goods' that epistemologists urge us to pursue—goods such as evidence, justification, and by extension knowledge. The latter two imply a similar denial given the extremely plausible assumption that substantial unreliability is incompatible with the production of said goods.

Opponents of experimental philosophy, too, characterize its claims in similar terms. Max Deutsch claims that '[experimental findings] should lead us, or so one might think, to question whether we are properly justified in intuiting, and then believing, what we do about the cases' (2015, p. 7). Elijah Chudnoff claims that experimental philosophers believe their results to 'give us reason to think that intuitions are unreliable' (2014, p. 108). Boyd and Nagel (2014) portray experimental philosophers as arguing that 'we have empirical evidence for the unreliability of epistemic intuition' (p. 114). Finally, one has the bevy of defenses of intuition that take, as their primary target, a 'skeptic' about intuition—implying that the intended opponent denies that intuitions produce knowledge. Instances of this are found in, for example, the work of George Bealer (1992, 1996), Joel Pust (2000), David Sosa (2006), Alvin Goldman (2007), and Timothy Williamson (2004, 2007, 2016).

Framing the epistemological consequences of experimental philosophy in terms of these 'standard' epistemic categories—evidence, justification, knowledge, and so forth—is puzzling. It suggests that what is primarily at issue is intuition's 'overall' epistemic status. If intuition is a source of evidence/justification/knowledge, then the experimentalists lose; if it is not, then they win. But methodological critiques in other fields don't at all take that form. Early proponents of double-blinding in science, for instance, were not 'skeptics' about ordinary observation; they did not take their primary question to be whether or not observation generates knowledge. The focus on knowledge and its epistemic siblings is, I'll argue, just as inappropriate for debates over philosophical method. The threat of skepticism generated by said focus is wholly illusory.

Timothy Williamson's complex and powerful critique of 'intuition skepticism' provides, I think, a perfect illustration of why this is so—and though the argument succeeds on its own terms, ultimately, I claim, it misses the point. In what follows, I'll use Williamson's critique as an example of how metaphilosophical arguments framed in terms of the concepts of standard analytic epistemology are, quite generally, a red herring. An adequate evaluation of the consequences of empirical

findings on intuition, I'll argue, simply cannot be conducted in such terms; finer carvings of epistemic space are needed.

Williamson's critique of intuition skepticism is multi-dimensional, stemming in large part from various doubts he has about the use of 'intuition' as a defining feature of philosophical methodology. One such doubt involves a general anti-psychologism with regard to evidence; on Williamson's view, we should conceive of our evidence in philosophy not as consisting of propositions like 'I intuit that the Gettier case is not a case of knowledge,' but as consisting of propositions like 'the Gettier case is not a case of knowledge.'[2] Psychologizing our evidence, Williamson notes, invites skepticism—and indeed, external world skepticism derives much of its force from the difficulty of arguing from psychological premises alone to any non-psychological conclusion whatsoever. We should, then, resist the temptation to characterize philosophical evidence in terms of intuition.

Williamson is plausibly correct here, but note that this by itself isn't sufficient to render experimental philosophy problematic—so long as its proponents are relatively careful with how they frame their critique. A scientist's evidence plausibly consists of propositions about, e.g., acids, proteins, or what have you; it does not plausibly consist solely of propositions about the perceptions experienced by scientists in the lab. Yet it's obvious that epistemological deficiencies in perception suffice to problematize certain methods in the sciences—hence the use of, e.g., measuring devices rather than bare unaided observation. A methodological critique based on experimental evidence for similar deficiencies in intuition in no way requires psychologizing evidence.

A second aspect of Williamson's critique, however, appears *prima facie* to be more problematic for the experimentalist. On Williamson's view, the judgments we make in response to thought experiments are underwritten by multiple different cognitive processes, including ordinary counterfactual reasoning and the quite general ability to apply concepts to particulars. Quite obviously, such abilities also underwrite much of our everyday cognitive activity. As Williamson puts it, 'any psychological kind that includes armchair philosophical judgments includes a mass of non-philosophical judgments too'

(Williamson 2009, p. 475). He notes that 'the obvious danger is that the category of philosophical intuitions will be stretched so wide, encompassing virtually anything one says about actual cases, that experimental philosophers' critique of reliance on philosophical intuitions will become a global skepticism' (Williamson 2016, p. 24).

Such observations bring out two separate problems with the current focus, within metaphilosophical debates, on intuition's 'overall' epistemic status—that is, its status as a source of knowledge, justification, and/or evidence. The first is that, though experimentalists aim their critiques at philosophical inquiry, intuition is clearly not limited to that practice alone. 'Intuition' covers states that are used in ordinary cognition, in the sciences, and so forth, and so a wholesale ban on intuition renders those activities problematic as well. If intuition does not generate knowledge (or evidence, or what have you), then a great deal of our everyday judgments are suspect. The second problem is that 'intuition' covers states that are heterogeneous—the single term 'intuition' obscures the fact that philosophical judgment invokes a variety of quite different forms of cognition which quite plausibly differ widely in their epistemological features.[3] To urge the rejection of 'intuition,' then, is to urge the rejection of a number of quite different psychological processes in one fell swoop. I would add that even the aforementioned 'ability to apply concepts to particulars' isn't plausibly a single ability, but many. Consider the experimental evidence suggesting that disgust affects moral judgments (see, e.g., Wheatley and Haidt 2005; Schnall et al. 2008); consider how incredible it would be to argue, on that basis, that we should assume logical judgments will be similarly affected. Yet this is exactly what experimentalists in effect suggest when they make claims like those expressed in the quotes above—they implausibly treat intuition as an epistemologically uniform kind, and they implausibly lump together as 'unreliable' or 'unjustified' a variety of presumably epistemologically dissimilar states.

Ultimately, the heterogeneity of intuition is—like an anti-psychologistic account of evidence—relatively easy for experimentalists to accommodate. Rather than argue that 'intuition' is unreliable or fails to produce evidence, the experimentalist merely needs to draw her conclusions more narrowly—in terms of 'moral judgment' or 'logical judgment' rather

than 'intuition,' for instance.[4] This does, I think, give the experimentalist the resources to avoid a blanket global skepticism: Rather than condemning 'our practices of applying concepts in judgment' (Williamson 2007, p. 220), she can content herself with condemning only those judgments which are plausibly sufficiently psychologically similar to judgments which experiments suggest to be problematic.[5] How similar is 'sufficiently' similar? That's a difficult question to answer, but it isn't one exclusive to the experimental critique—we face such issues when we, for instance, attempt to determine the consequences of psychological studies on memory for our treatment of eyewitness testimony.

The trickier issue is not the risk of skepticism per se, but rather the 'overlap' between philosophical and non-philosophical cognition—and it is here that the experimentalists' (and Williamson's) reliance on the standard categories of epistemology becomes problematic. Abandoning 'intuition' for a variety of more narrowly drawn mental categories might suffice to avoid global skepticism, but it does not fully avoid the overlap issue. Consider the fact that, although the cognitive processes underlying moral and logical cognition are plausibly different, those that underlie everyday knowledge attribution ('my mother knows that today is my birthday', for example) and attributions of knowledge made in the context of a philosophical thought experiment are plausibly largely the same.[6] Even a 'mere' rejection of epistemic intuitions, then, would still invite the rather implausible view that ordinary, everyday knowledge attributions are unjustified. And *mutatis mutandis* for other cognitive processes upon which experimentalists might cast doubt. It is this issue—the *overlap* between the forms of cognition employed in everyday activity and in philosophical inquiry—that I take to be the primary challenge that Williamson's critiques present to experimentalist criticisms of philosophical method. And it is this aspect of the challenge that highlights, I'll argue, a need for a more fine-grained, pluralistic approach to epistemic states.

2   Dealing with the Overlap Problem

One possible strategy for dealing with 'overlap' problem is to appeal to an epistemic difference between different practices which make use of intuition, rather than differences between the various sub-types of intuition itself. Jonathan Weinberg (2007) has argued for just such a strategy. Perception, Weinberg notes, is quite obviously fallible. Yet this fallibility is rendered unproblematic by our everyday epistemic practices:

> our practices involving perception will by and large keep us from getting into too much epistemic trouble…[we] know a great deal about the circumstances and ways in which it is fallible and about what to do when we find ourselves in such circumstances. (Weinberg 2007, p. 324)

In Weinberg's terminology, our everyday perceptual practices are 'hopeful'—they allow us to detect and correct for the errors inherent in perception. By contrast, Weinberg suggests that the philosophical practices surrounding use of intuition are hope*less*. It is this hopelessness, rather than the unreliability of intuition itself (which is plausibly not much greater than that of perception), which renders our philosophical practices worthy of criticism.

This is a move in the right direction, but it isn't enough to avoid the overlap problem. For it's not clear that philosophical practices surrounding use of intuition are actually more hopeless than the corresponding everyday practices surrounding use of intuition. It's not as though there are error-mitigating procedures in ordinary cognition that we fail to avail ourselves of when philosophizing—if anything, our ordinary epistemic practices with intuition are plausibly less rigorous, and consequently *more* hopeless, than our philosophical practices with such. As far as I can tell, the only apparent feature of philosophical practice that is supposed to render it more problematic than everyday usage of intuition is the fact that philosophy sets 'no constraints on how esoteric, unusual, far-fetched, or generally outlandish any given case may be' (Weinberg 2007, p. 321). Yet it's not at all clear that 'far-fetched' cases are (or should be) the sole targets of the experimentalist critique. Nor is it clear exactly how to draw the line between the 'esoteric' and

the ordinary—many Gettier cases, for instance, are fairly mundane (see also Williamson (2016) on the difficulty of isolating the 'far-fetched'). If all this is more or less correct, then a condemnation of philosophical practices which employ intuition still threatens to generalize to much of ordinary cognition.

Recently, a somewhat different approach has been suggested both by Alexander and Weinberg (2014) and by myself (2011, 2015). This approach focuses on what I have called the 'epistemic demandingness' of philosophy as compared to ordinary cognition.[7] The core idea is that different tasks come with different epistemic requirements; an epistemic source that is 80% reliable, for instance, might be sufficiently reliable for some tasks but insufficiently reliable for others. In particular, if a single task requires multiple successful uses of the same source, the risk of failure quickly multiplies—an 80% reliable ability to distinguish crows from ravens results in only about an 11% chance of success if your task requires you to properly distinguish 10 such birds in a row.

The Gettier literature provides a perfect example of why philosophy may well be highly demanding, in the sense just described. Gettier judgments form a tiny minority of knowledge attributions—an intuiter who erred *only* on Gettier cases would thus be a highly reliable knowledge attributor. Yet suppose this intuiter was to use the set of knowledge attributions which she found intuitive as the sole basis for her theory of knowledge. The sliver of unreliability in her knowledge attributions would, if undetected, quite likely lead our hypothetical intuiter to endorse the widely rejected JTB theory of knowledge. Conversely, it's quite possible that JTB is in fact the *correct* theory of knowledge and that our generally reliable ability to classify cases of knowledge has, due to a small stain of error, set us off on the wild-goose chase that is Gettierology. In fact, given the heavy weight we philosophers give to single counterexample cases, it's plausible that just about any nonzero degree of unreliability brings with it the potential to massively disrupt successful theorizing.[8] Quite generally, in science as well as in philosophy, small amounts of error in a data source can lead to radically false theories—and if we want to minimize this risk, then we may need to implement especially stringent procedures to mitigate said errors.

Suppose that this line of argument is on the right track. It suggests that we need to distinguish not only different practices which employ intuition, but also different *epistemic standards* which correspond to those practices. And this, I think, is the real heart of the issue—participants in the intuition debates have falsely assumed that a single epistemic standard suffices for both ordinary epistemological questions and for metaphilosophical questions about the quality of our methods. Yet it may well be that the epistemic standards operative in philosophical contexts are unusually high; perhaps issues like the epistemic demandingness of philosophical theorizing lead to a greater-than-usual need for vigilance against error. If this is indeed so, then a given source might be of sufficiently high quality for ordinary activity while being too error-prone for philosophical purposes; the overlap problem, at least potentially, dissolves.

There are, of course, accounts of knowledge which accept that epistemological standards may shift between contexts—contextualist and subject-sensitive invariantist accounts being the most familiar. But metaphilosophers aren't primarily concerned with the nature of knowledge; they are concerned with whether or not the practices employed by philosophers meet the epistemic standards appropriate to philosophical activity. And the familiar shifting-standards accounts simply don't directly address that issue. Suppose it to be true that philosophical activity involves higher epistemic standards. Even if a contextualist or related account *were* able to successfully accommodate such varying standards, questions like 'does intuition produce knowledge' would *still* be more or less irrelevant to questions about the viability of current philosophical methods, and plausibly, our ordinary epistemic vocabulary would *still* be inappropriate in metaphilosophical contexts.

To see why, suppose for the sake of argument that the only epistemic requirement that varies by context is the minimum reliability of the process by which a belief is formed. Suppose further that subject-sensitive invariantism correctly captures this variation; in other words, suppose that the stakes are higher in philosophical contexts and that the standards for knowing are elevated in such high-stakes contexts. In particular, suppose that in ordinary contexts a reliability level of 80% suffices for knowledge, and suppose that, in philosophical contexts, a

reliability level of 99% is required for the same. Now finally, suppose intuition to be 85% reliable. Is intuition sufficiently reliable to generate knowledge (in our imagined scenario)? Yes—it produces knowledge in at least some contexts. Is it sufficiently reliable to *generally* produce knowledge? Almost certainly, given the relative rarity of philosophical contexts compared to ordinary ones. Such questions are, however, unhelpful if our goal is a critical examination of our methods. The question we *should* be asking is as follows: Is intuition sufficiently reliable to generate knowledge-in-a-philosophical-context? At which point, we may as well simply adopt a technical term—'knowledge$_p$,' say—to refer specifically to epistemic states that meet the standard we are interested in.

This is not to say that a subject-sensitive invariantist account *is* the right account of knowledge. In fact, I'll aim to argue that the epistemic norms underlying the various types of human cognitive activity are even more complex than the foregoing discussion has suggested; and that their complexity renders implausible any attempt to capture all the relevant phenomena under any kind of concept we would recognize as 'knowledge.' In that sense, the use of 'knowledge$_p$' and related terms is slightly misleading—the needed categories may not be straightforward analogs of standard epistemic categories, but rather wholly different types of epistemic state.

3   Epistemic Standards in Professional Inquiry

We don't, I think, need to commit ourselves to any claim about epistemic demandingness in order to motivate the idea that epistemic standards vary dramatically between different forms of inquiry and that epistemology's monolithic focus on knowledge is therefore problematic. We can simply consider our actual practices—for instance, the extraordinarily stringent methodological procedures that are considered obligatory in the sciences. Double-blinding is, of course, an obvious case here. The role of double-blinding in an experimental context is to eliminate potential biases resulting from the expectations of the experimenter or the subject—including such well-known phenomena as the placebo

effect. The practice of double-blinding is not employed in every scientific activity, of course. But if the threat of bias is present, and if it is possible to blind the study without negative effects (such as ethical violations), then it is arguably obligatory that the study be conducted blind. It would be wholly appropriate to criticize a study that failed to adhere to this standard.

But note that double-blinding is more or less *never* employed by ordinary cognizers during their everyday belief-forming activities. This is not because the relevant biases don't appear outside of scientific contexts. Ordinary cognition is rife with biases, including biases that result from the 'coloring' of observation by preexisting expectations. We simply don't expect anyone to do much about it outside of the laboratory; at least, we certainly don't expect anyone to go to the lengths that scientists do to mitigate said biases. It is not as though ordinary cognizers, then, are typically failing to live up to their epistemic obligations. Instead, it is that our epistemic expectations—or at least, our methodological expectations—differ between these two types of epistemic activity.

Double-blinding is only one example of many. Scientists are obligated to only rely on data that achieve statistical significance. They must control for confounding variables, where possible. They are obligated, in certain circumstances, to employ various measuring devices, videotaping protocols, multiple observers, and so forth to protect from the (minimal) errors that arise from ordinary visual observation. None of these procedures are ever required of ordinary cognizers—even when those cognizers are forming beliefs about the exact same propositions that are studied by science. Members of prescientific communities knew, for instance, all kinds of propositions about medicine, psychology, and so forth; they achieved this knowledge without the use of the procedures that are now obligatory in those fields. To deny this would be to claim that, for instance, no one knew that any of our various medicines were effective before the placebo effect became recognized and standardly controlled for in the *mid twentieth century*.

The sciences, then, plausibly are subject to heightened epistemic standards which obligate practitioners to employ extra-stringent methods to control for bias and error—standards that exceed what is

required for the attainment of knowledge in ordinary circumstances. I claim that the same is also true in philosophy—the epistemic standards to which professional philosophers ought to hold themselves are plausibly higher than the epistemic standards which must be met in order to know.

If epistemic standards really are elevated in fields like science and philosophy, then this prompts the following question: Why should there be these different standards? If it is possible, for instance, to control for biases and reduce error, why shouldn't the methods for doing so be epistemically obligatory for everyone, in all contexts? Plausibly, it is simply because the practices that must be employed to live up to the stringent standards of professional inquiry are far too overdemanding for ordinary cognitive activity. Let's focus, for the moment, on the sciences. Consider, for instance, how absurd it would be to criticize an ordinary cognizer for failing to run a double-blindcontrolled experiment before forming an ordinary causal generalization—say, the belief that the new office break policy has led to an increase in her co-workers' morale. The vast majority of non-scientists simply do not have the time, skills, or resources to subject their belief-forming to the rigorous methodological controls used in science. Employing scientific rigor in all instances of everyday belief-formation would be inefficient, needlessly overmeticulous—obsessive, even. Nothing would ever get *done*.

It's common to argue that knowledge is implicated in various norms surrounding assertion, belief, and action. For instance, I should act on a given proposition only if I know that proposition to be true; I should believe a proposition only if I am in a position to know it. Suppose such arguments to be on the right track; knowledge, then, is a concept that is intimately linked with other aspects of our ordinary lives. If knowledge required ultra-stringent procedures for eliminating error, then, this would plausibly have extreme consequences for our practices of communication, practical reasoning, and so forth. If I cannot know p before subjecting p to rigorous experimental test, then this might imply that I cannot justifiably act on p until such a test has been conducted. Again, that seems preposterously overdemanding; it seems absurd to think that ordinary individuals are epistemically required to invest so much time and effort before acting. Again, then, we have reason to think that

ordinary folk are not epistemically required to meet the methodological standards which are obligatory in scientific activity.

Obviously, though, we do not consider it overdemanding to require that *scientists* subject their *scientific* activities to said methodological requirements. Now suppose we were to employ a straightforward contextualist or subject-sensitive invariantist account to try to capture these differing expectations. If one assumed a shifting-standards explanation of the differing requirements—e.g., if one claimed that the standards for knowing are simply raised in scientific contexts—then this would suggest that the standards for believing, acting, and asserting would similarly be raised in said contexts. It's not clear to me that they are. Indeed, the methodological standards we've been discussing don't seem to place any limitations whatsoever on individual belief or action. A researcher who forms the belief that her theory is true does not thereby break the norms of science—so long as she dutifully continues subjecting her theory to the appropriate experimental tests. A researcher who takes her own experimental medication before fully verifying its effectiveness through experiment does not thereby break the norms of science—so long as she does not use her own case as data in her experiments. We might question whether said belief or action is warranted, but we would not accuse these hypothetical researchers of bad science: Their actions do not make them poor researchers.[9] Contrast this with a case where a scientist fudges her experimental data, or fails to properly calibrate her instruments.

Scientific methodology is not concerned with regulating the beliefs and actions of individuals, except where those actions specifically violate required methodological procedures. The standards of science simply do not interface with the familiar knowledge norms we have been discussing. This makes perfect sense, for science is a community effort rather than an activity conducted by an isolated researcher. The community aspect of science, in fact, plausibly explains why we don't view the stringent methods of science as overdemanding in an experimental context. The limitations on time and resources that individuals face are not faced by the scientific community as a whole; if it takes several generations to confirm a hypothesis to the satisfaction of the scientific community, then so be it. The exemplars of knowledge upon which traditional

epistemology has been built are mental states of individuals; but a successful characterization of the epistemic features of group inquiry may well differ from traditional epistemology quite substantially.

I've argued that ordinary cognizers are not required to employ the strict bias-reducing and error-eliminating practices that are obligatory in scientific inquiry. But this should not be taken to imply that biases cannot undermine knowledge in ordinary circumstances—of course they can. If my belief that Smith is the most qualified candidate is based on a cherry-picked set of sources all of which back the view I favor, this might well prevent my belief from qualifying as knowledge (even if Smith is *in fact* the most qualified candidate). Indeed, it *may* well be that a belief must be free of bias in order to count as knowledge (though I personally suspect that this is too strong a requirement). But notice that the requirements of science are higher still. It is not enough that a scientific study simply *be* free of bias; the experimenter must have taken explicit steps to *ensure* that the study is free of bias. A non-blinded study is still a flawed study, even if experimenter expectation did not *in fact* affect the experiment's results.

This might suggest to us that at least in the scientific case, what matters is not *knowing*, but knowing-that-one-knows. This might, in turn, enable an explanation of the apparent variation in methodological standards without requiring any sort of contextualist or subject-sensitive invariantist stance on knowledge. In fact, Williamson has advocated just this sort of strategy for explaining our intuitions in standard cases of raised stakes (such as the well-known bank and airport cases):

> [S]uppose that the agent in C knows p and the agent in C* knows p*, but the agent in C is in no position to know that she knows p and the agent in C* is in no position to know that she knows p*. Since stakes are higher in C* than in C, we as theorists may view the failure of second-order knowledge in C* more sternly than its failure in C, and therefore regard p as appropriate in C but p* as inappropriate in C*. (Williamson 2005, p. 231)

In other words, when the stakes are higher, an agent has more reason to 'double-check' on her knowledge—and we judge that such an agent has

erred when they fail to do so. This need not entail that knowledge does not suffice for epistemically warranted action; locking one's door is sufficient for securing one's belongings from theft, but in some contexts, we might nonetheless think it prudent to double-check whether one *has* in fact met this standard. It would be open to Williamson to treat the variable standards among fields of inquiry in the same way—claiming that, in such fields, we think it is important to 'double check' that our data consists of genuinely known propositions.

Suppose, then, that we consider this as a possible framework for discussing metaphilosophical issues. Suppose, in other words, that we frame the experimentalist argument as follows: The flaws of intuition don't prevent it from being a source of knowledge, but they do prevent us from achieving second-order knowledge—that is, knowledge that we know. And, in philosophy, what we are really concerned with is knowing that we know. Note that, if this approach succeeds, there is no need for 'new' epistemic vocabulary, or for any departure from a classic knowledge-centric approach to epistemology. We could simply conduct metaphilosophical debate in terms of ordinary iterations of knowledge.

But unfortunately, iterations of knowledge aren't likely to provide a sufficiently fine-grained set of epistemic categories to capture the standards at issue. An examination of non-scientific fields of inquiry and their proprietary standards will, I think, demonstrate just how much variation we may need to countenance in order to capture the methodological requirements I've been discussing. This will be the focus of the next section.

4   The Multiplicity of Epistemic Standards

Science is not the only type of inquiry that plausibly operates with elevated epistemic and/or methodological standards. Consider some of the methodological obligations present in courtroom proceedings: 'leading' questions are forbidden, restrictions are put on admissible evidence, and verdicts must be held to be beyond reasonable doubt. Needless to say, these methodological constraints are not obligatory in most ordinary instances of belief-formation involving testimony or physical evidence.

Or consider fact-checking expectations and standards of objectivity that hold in journalism. Or standards surrounding the use of written sources in history. And so on. Plausibly, nearly every academic and professional field has at least some methodological strictures aimed at improving epistemic position—strictures which they do not expect non-participants to either know about or hold to.

Indeed, some fields seem to have even higher epistemic standards— or, perhaps, higher epistemic aims—than the sciences. Consider the case of Goldbach's conjecture in mathematics, which states that every even integer greater than 2 is the sum of two primes. We have direct verification that this conjecture holds of all primes up to 4×10^{18}. Thus, we have tested the claim on *millions* of cases—many more confirming cases than most ordinary empirical experiments ever receive. Yet the conjecture lacks a proof, and mathematicians thus *continue* to work on the problem, and take it to be *wholly worthwhile to do so*. They will be satisfied with nothing less than either complete deductive certainty, or a demonstration that no proof is possible.

I find it plausible that we *know* Goldbach's conjecture to be true (or, more carefully, assuming that it is *in fact* true, I find it plausible that our epistemic situation with respect to the conjecture suffices for knowledge). There are, of course, complicated issues surrounding the use of inductive evidence in mathematics. But similar cases suggest that we really should at least sometimes take ourselves to know in cases where proof is possible but absent. Consider the fact that several ancient cultures seem to have been aware of the truth of the Pythagorean theorem long before the notion of a proof was even first conceived—the well-known Plimpton 322 papyrus from c. 1800 BCE Babylon, for instance, lists a large number of triples of integers satisfying $a^2 + b^2 = c^2$, suggesting that Babylonians possessed both a method for generating such triples and an understanding of their import. Presumably, this understanding of the relationship between the sides of a right triangle was put to use by such cultures in practical contexts—say, in land surveying tasks. The fact that such actions seem epistemically warranted suggests that pre-Greek mathematics *is* genuine knowledge. Yet knowing is not enough for current mathematicians—Pythagoras's proof was not superfluous. Continued pursuit of proof in such cases (and others—consider

the quest for a proof of Euclid's parallel postulate) indicates that mathematicians don't aim at mere knowledge, but at something much higher—perhaps complete certainty. Needless to say, nothing approaching this level of certainty is plausibly achievable in, say, science.

Even within a field, different standards presumably apply to different subfields, to different tasks, or even to different individual propositions.[10] The conditions under which a scientist may assert 'our experiment confirms our hypothesis' are obviously different from the conditions under which she may assert 'our experimental group contained 235 subjects.' Given this cornucopia of wildly different epistemic expectations, I think we should be a bit pessimistic about the idea that iterations of ordinary knowledge provide enough categories to capture this epistemic complexity. Do all fields aim for knowledge-of-knowledge? Are the standards in place in science exactly as rigorous as the standards in place in journalism, law, or history? What about mathematics—is there *any* iteration of knowledge that maps perfectly onto the deductive certainty that seems to play a central role in mathematics?

I can't fully consider such questions here. But the following seems, to me, to be a better overall characterization of the situation. There exist various methods and/or epistemic policies that have the potential to increase, to various degrees and plausibly along various dimensions, the quality of our epistemic position. When we are considering implementing some such rule, we are faced with a trade-off: Is the increase in epistemic position worth the non-epistemic costs (time, effort etc.) that come along with the adoption of the method? The answer to that question is likely to be different for an individual and for a group; it might be different for different fields, and for different tasks within a field. Stretching 'knowledge' and its iterations to capture all such acceptable trade-offs may be possible, but it's not clear why it should be necessary. As metaphilosophers we can, I think, simply adopt discipline-specific epistemic categories and not concern ourselves with the relation that such states may or may not have to knowledge.[11]

5 The Epistemic Standards of Philosophy

Let's return, at long last, to the hypothesized elevated standards of philosophy. Does philosophy, like the other disciplines we've discussed, really aim higher than 'mere' knowledge? It is true that philosophy lacks the carefully articulated methodological requirements that operate within the sciences—there is no obvious analog for practices such as double-blinding, or even the use of measuring apparati. But it seems plausible that philosophy is nonetheless subject to epistemic standards that far exceed what we expect in ordinary cognition. In order for their work to be considered publishable, philosophers are expected to carefully define any crucial terms they employ in their arguments; they are expected to demonstrate familiarity with relevant literature; they are expected to anticipate and respond to obvious potential objections. They are subject to expectations of rigor, clarity, objectivity, and the like that are more or less beyond the abilities of the majority of untrained, ordinary cognizers. Yet surely ordinary cognizers have knowledge of at least some propositions that have been, or could be, subject to philosophically rigorous debate. It is wholly implausible that no non-philosopher has ever known, for instance, any of the various ethical propositions, religious propositions, or so on over which philosophers have argued—not to mention potentially debatable propositions regarding the mind, beauty, causation, existence, and all the rest.

And even if we leave aside the question of whether philosophy *currently* employs elevated methodological standards, we can clearly make the case that it *should*. For note that the exact same considerations surrounding the elevated methodological standards in the sciences apply in philosophy, as well. Philosophy is a group effort, conducted by persons who are paid to spend large portions of their waking hours on philosophical inquiry. Constraints on available time, resources, skills, and so forth don't clearly apply to the philosophical effort as a whole. Supposing that at least some of the many mental processes that fall under the banner of 'intuitive' cognition do display biases of various sorts, it seems obvious that we ought to seriously concern ourselves with devising methods to mitigate such epistemic flaws. This is, essentially,

the same argument that one might give in support of the use of double-blinding in the sciences.

It's not a very large step from there to the following claim: It is worthwhile for the philosophical community to divide its cognitive labor such that at least some members devote their time to empirical investigation of potential biases in philosophical cognition. It is also worthwhile for at least some members to devote their time toward consideration of possible mitigating methodological changes, or to impressing upon their fellows the potential pitfalls of overconfidence in a source whose epistemic credentials are not yet wholly established. It may well be that all current work in experimental philosophy is flawed or that the findings thus far give us no good reason to accept the claim that 'intuitive' judgment displays bias. But our general awareness of the foibles inherent in nearly all areas of human cognition is, alone, enough reason to undertake the task of experimentally investigating the cognitive processes that we rely on as philosophers. One might even argue that doing so is obligatory—not for any individual philosopher, but for the philosophical community as a whole.

Here, then, is how I would frame an experimentalist critique of current philosophical methods. Intuition (or, better, most types of cognitive process that we lump together as 'intuitive') presumably generates knowledge—where this is understood as, say, the epistemic state that an individual must meet in order to assert or act in ordinary everyday circumstances. But empirical findings suggest that intuition has substantial flaws, and these flaws are incompatible with knowledge$_p$—where this is understood as something like the epistemic standing we ought to aim for as professional philosophers (again, an even more careful argument might distinguish even more epistemic standings, corresponding to the aims/obligations of different philosophical subfields, or required by different philosophical tasks, and so forth). This is, quite clearly, a more moderate position than 'intuition skepticism.' But it is also, I think, both a more subtle and a more plausible one, too.

Finally, I think it is an *obvious* position. It is, to my eyes, rather baffling that so much philosophical ink should have been spilled debating whether or experimental findings license a complete rejection of intuition. Obviously, they do not. But just as obviously, they license a

critique of intuition—one that cannot be easily formulated, however, if one insists on focusing solely on the presence or absence of knowledge, justification, or evidence. It is for that reason, then, that I have urged the introduction of novel epistemic concepts. Categories like the one picked out by 'knowledge$_p$,' do not correspond to our pre-theoretic 'folk' epistemological notions in any obvious way; but this, I think, should be no objection to their use. 'Knowledge$_p$' and its ilk are novel technical terms—but ones that delineate categories of genuine use to epistemologists and methodologists.[12] By means of such categories, we step out of the monistic, knowledge-centered epistemic framework which has obscured the mundanity of the methodological 'crisis' surrounding intuition—a crisis which exactly parallels the methodological challenges faced by scientists aiming to evade the biases of observation, or of courts aiming to evade the biases of memory. As with those challenges, skepticism is wholly beside the point.[13]

Notes

1. In the case of knowledge$_p$, it would presumably also be a requirement that the content of said state is true—though even this might be questioned if one were, e.g., inclined to some form of anti-realist interpretation of philosophical inquiry (which I am not, I might add).
2. Throughout the paper, I will assume (with Williamson) that evidence consists of propositions.
3. See also Nado (2014).
4. In fact, I think the needed categories will be much narrower than this—even moral judgment plausibly involves multiple psychologically and epistemologically distinct mental processes.
5. This does tone down the ambition of experimental philosophy's negative program quite significantly. But the failure of the armchair-in-flames approach to methodological criticism doesn't at all show that philosophers should rest on their laurels—even a piecemeal criticism of problematic mental states has the potential to be quite devastating.
6. I leave aside here questions of expertise; such issues would favor the experimentalist on this particular point, regardless.

7. Alexander and Weinberg's term for the same phenomenon is 'error fragility.'

8. None of this, of course, demonstrates that any of our theories *are* so affected. But the likelihood of such a scenario may be significant enough to warrant concern about our methods—just as the likelihood of experimenter bias is significant enough to warrant obligatory double-blinding. Again, nothing about this requires any form of skepticism; if such an argument mirrors that of the skeptic (as might be suggested by discussion in Williamson (2007)), then so too does the argument for stringent methodological standards in science.

9. The case of assertion is somewhat more complex. Certain assertions do seem to be prohibited by the methodological norms of science—at least in a publication context.

10. I take it this leaves open the possibility that the very same intuited proposition might be acceptable for use in one field but not another—especially if it is put to different tasks in those different fields. Williamson (2009) responds to Weinberg by noting that science makes use of intuitive epistemic judgments, when, e.g., concluding that data renders a hypothesis probable. Yet it is possible that the use of intuition is acceptable for that task while being insufficiently rigorous for the purposes of, say, formulating an overall theory of knowledge.

11. That's not to say that the question holds no interest—but it is an epistemological question, not a methodological one.

12. This need not imply that categories like knowledge$_p$ denote 'natural kinds' or anything of the like; I claim merely that they reflect useful carvings of the multi-dimensional spectrum of possible epistemic standings.

13. Acknowledgement: The work described in this paper was fully supported by a grant from the Research Grants Council of the Hong Kong Special Administrative Region, China (Project No. LU 359613).

References

Adleberg, T., M. Thompson, and E. Nahmias. 2015. Do Men and Women Have Different Philosophical Intuitions? Further Data. *Philosophical Psychology* 28 (5): 615–641.

Alexander, J., and J.M. Weinberg. 2014. The "Unreliability" of Epistemic Intuitions. In *Current Controversies in Experimental Philosophy*, eds. E. Machery and E. O'Neill. New York: Routledge.

Alexander, J., R. Mallon, and J. Weinberg. 2010. Accentuate the Negative. *Review of Philosophy and Psychology* 1 (2): 297–314.

Bealer, G. 1992. The Incoherence of Empiricism. *Proceedings of the Aristotelian Society, Supplementary volume* 66: 99–138.

Bealer, G. 1996. On the Possibility of Philosophical Knowledge. *Noûs*: 1–34.

Boyd, K., and J. Nagel. 2014. The Reliability of Epistemic Intuitions. In *Current Controversies in Experimental Philosophy*, eds. E. Machery and E. O'Neill. New York: Routledge.

Buckwalter, W., and S. Stich. 2014. Gender and Philosophical Intuition. In *Experimental Philosophy* vol. 2, eds. J. Knobe and S. Nichols, 307–346. Oxford: Oxford University Press.

Chudnoff, E. 2014.*Intuition.* Oxford: Oxford University Press.

Cullen, S. 2010. Survey-driven Romanticism. *Review of Philosophy and Psychology* 1 (2): 275–296.

Deutsch, M. 2015. *The Myth of the Intuitive*. Cambridge: MIT Press.

Feltz, A., and E. Cokely. 2008. The Fragmented Folk: More Evidence of Stable Individual Differences in Moral Judgments and Folk Intuitions. In *Proceedings of the 30th Annual Conference of the Cognitive Science Society*, eds. B.C. Love, K. McRae, and V.M. Sloutsky, 1771–1776. Austin: Cognitive Science Society.

Feltz, A., and E. Cokely. 2009. Do Judgments About Freedom and Responsibility Depend on Who You Are? Personality Differences in Intuitions About Compatibilism and Incompatibilism. *Consciousness and Cognition* 18 (1): 342–350.

Goldman, A. 2007. Philosophical Intuitions: Their Target, Their Source, and Their Epistemic Status. *Grazer Philosophische Studien* 74: 1–26.

Kauppinen, A. 2007. The Rise and Fall of Experimental Philosophy. *Philosophical Explorations* 10 (2): 95–118.

Knobe, J., and S. Nichols. 2008. An Experimental Philosophy Manifesto. In *Experimental Philosophy*, ed. J. Knobe and S. Nichols, 3–14. New York: Oxford University Press.

Machery, E., R. Mallon, S. Nichols, and S. Stich. 2004. Semantics, Cross-Cultural Style. *Cognition* 92: B1–B12.

Machery, E., and E. O'Neill. 2014. Introduction: Experimental Philosophy: What is it good for? In *Current Controversies in Experimental Philosophy*, eds. E. Machery and E. O'Neill. New York: Routledge.

Nado, J. 2011. *Intuition and Inquiry*. PhD thesis, Department of Philosophy, Rutgers University.

———. 2014. Why Intuition ? *Philosophy and Phenomenological Research* 89: 15–41.

———. 2015. Intuition, Philosophical Theorising, and the Threat of Scepticism. In *Experimental Philosophy, Rationalism, and Naturalism: Rethinking Philosophical Method*, eds. E. Fischer and J. Collins, 204–221. New York: Routledge.

Nagel, J., V. San Juan, and R. Mar. 2013. Lay Denial of Knowledge for Justified True Beliefs. *Cognition* 129 (3): 652–661.

Pust, J. 2000. *Intuitions as Evidence*. New York: Routledge.

Schnall, S., J. Haidt, G.L. Clore, and A.H. Jordan. 2008. Disgust as Embodied Moral Judgment. *Personality and Social Psychology Bulletin* 34: 1096–1109.

Seyedsayamdost, H. 2015. On Normativity and Epistemic Intuitions: Failure of Replication. *Episteme* 12 (01): 95–116.

Sosa, D. 2006. Scepticism About Intuition. *Philosophy* 81 (04): 633–648.

Stich, S., and K. Tobia. 2016. Experimental Philosophy and Philosophical Tradition. In *Blackwell Companion to Experimental Philosophy*, eds. W. Buckwalter and J. Sytsma. Malden: Blackwell.

Swain, S., J. Alexander, and J. Weinberg. 2008. The Instability of Philosophical Intuitions: Running Hot and Cold on Truetemp. *Philosophy and Phenomenological Research* 76 (1): 138–155.

Tobia, K., W. Buckwalter, and S. Stich. 2013. Moral Intuitions: Are Philosophers Experts? *Philosophical Psychology* 26 (5): 629–638.

Weinberg, J. 2007. How to Challenge Intuitions Empirically without Risking Skepticism. *Midwest Studies in Philosophy* 31: 318–343.

Weinberg, J., S. Nichols, and S. Stich. 2001. Normativity and Epistemic Intuitions. *Philosophical Topics* 29: 429–460.

Wheatley, T., and J. Haidt. 2005. Hypnotic Disgust Makes Moral Judgments More Severe. *Psychological Science* 16 (10): 780–784.

Williamson, T. 2004. Philosophical "Intuitions" and Skepticism about Judgment. *Dialectica* 58 (1): 109–153.

———. 2005. Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge. *The Philosophical Quarterly* 55 (219): 213–235.

———. 2007. *The Philosophy of Philosophy*. Oxford: Blackwell.

———. 2009. Replies to Ichikawa, Martin and Weinberg. *Philosophical Studies*, 145 (3): 465–476.

———. 2016. Philosophical Criticisms of Experimental Philosophy. In *A Companion to Experimental Philosophy*, eds. J. Sytsma and W. Buckwalter. Oxford: Wiley Blackwell.

Woolfolk, R.L. 2013. Experimental Philosophy: A Methodological Critique. *Metaphilosophy* 44: 79–87.

Pluralism About Knowledge

Robin McKenna

1   Introductory Remarks

Monism about some property P or relation R is the view that there is a single P-property or R-relation. Pluralism about P or R is the view that there is a plurality of P-properties or R-relations. There is a burgeoning literature on whether we should be monists or pluralists about truth (see, for instance, Lynch 2009; Pedersen and Wright 2013; Wright 1992). In this chapter, I explore whether we should be monists or pluralists about knowledge. The standard view in epistemology is that the knowledge relation is a two-place relation between a subject S and a proposition *p* of the form K(S, *p*).[1] The *knowledge monist* says that there is a single knowledge relation of this form. The *knowledge pluralist* says that there is a plurality of knowledge relations of this form. I will focus on a kind of knowledge pluralism I call *standards pluralism*. Put roughly,

R. McKenna (✉)
University of Vienna, 1010 Vienna, Austria
e-mail: rbnmckenna@gmail.com

standards pluralism is the view that one never knows anything *simpliciter*. Rather, one knows relative to this-or-that epistemic standard. Because there is a plurality of epistemic standards, there is a plurality of knowledge relations (knowing relative to this standard, knowing relative to that standard and so on).

Many readers will already be wondering about the connection between standards pluralism and two familiar views in contemporary epistemology. The first view is *epistemic contextualism*, according to which utterances of sentences of the form "S knows that *p*" express different propositions in different contexts of utterance (see Blome-Tillmann 2014; Cohen 1999; DeRose 2009; Lewis 1996). For instance, in a context where low epistemic standards are appropriate, the proposition expressed would be something like *S knows that p by low standards*, whereas in a context where high standards are appropriate, it would be something like *S knows that p by high standards*. The basic difference between these views is that standards pluralism is a metaphysical thesis about the knowledge relation, whereas epistemic contextualism is a linguistic thesis about utterances of sentences containing the "knows".

The second view is *epistemic relativism*, according to which one is never justified in believing something (or never knows anything) *simpliciter*, but only relative to this-or-that epistemic system, and there is a plurality of epistemic systems, all of equal validity (for a defence, see Kusch 2016; for a critique, see Boghossian 2006). The basic difference between these views is that, while both "relativise" knowledge to something (epistemic systems and epistemic standards), the epistemic relativist insists that all systems are equally valid, whereas the standards pluralist does not insist that all standards are equally valid.

Here is the plan. In §2, I start by reviewing some views that either involve or seem to entail a kind of pluralism about knowledge. I then introduce standards pluralism and construct a case for it that largely parallels the case for contextualism. In §3, I address the question of the relationship between standards pluralism, contextualism and epistemic relativism. In §4, I argue that standards pluralism faces a serious objection. The gist of the objection is that standards pluralism is incompatible with plausible claims about the normative role of knowledge. In §5,

I finish by sketching the form that a standards pluralist response to this objection might take.

Before continuing, I want to say a little more about the motivations for my project. First, my aim is to outline what can be said for and against pluralism about knowledge. Accordingly, my conclusion will be somewhat inconclusive: there is plenty to be said for pluralism about knowledge, but there is also plenty to be said against it. Nevertheless, I intend this chapter to be a spur to further work on this topic.

Second, the literature on the semantics and pragmatics of knowledge ascription is at something of an impasse. The costs and benefits of the various views are well-understood, but different authors disagree as to how they should be weighed up. It is my view that approaching these issues from a rather different angle—at the level of metaphysics rather than semantics—will help to advance the debate. While contextualism does not entail standards pluralism, or vice versa, it is worth considering the prospects for one in the light of the prospects for the other. To the extent that this chapter gives us reason to be optimistic/pessimistic about the prospects for standards pluralism, it gives us reason to be equally optimistic/pessimistic about the prospects for contextualism.

2   The Case for Knowledge Pluralism

I will start with some clarifications and a review of some relevant literature (§2.1). I then introduce standards pluralism (§2.2) and try to make a case for it (§2.3).

2.1   Types of Epistemic Pluralism

First, I am interested in pluralism about *propositional* knowledge, that is, knowledge-*that*. It is plausible (though not uncontroversial; see, e.g. Stanley 2011) that some types of knowledge are non-propositional, e.g. knowledge-*how*, knowledge-*whether,* etc. While this is a sort of pluralism about knowledge, it is not the kind of pluralism I am interested in.

Second, pluralism about knowledge should not be confused with the truism that there are many ways of knowing a proposition. I might know that my glasses are on the kitchen table because I can see them or because I remember leaving them there. But it would be odd to say that this shows there are two different knowledge relations: knowing-through-perception and knowing-through-memory. Rather, there are two ways in which one can come to stand in the same relation. What they have in common is, roughly, that they are cognitive successes of the same kind. To say that one can know *p* through perception or memory is to say that there are different routes to the same success. In contrast, to say that there is a plurality of knowledge relations is to say that there are different relations that are successful in different ways.

There are two kinds of pluralism about epistemic properties and relations that are familiar in epistemology. The first is pluralism about justification. For instance, consider William Alston's (1993, 2005) view that there are several "desiderata" for a theory of justification, but no one property that can satisfy all these desiderata. For instance, a belief may have the property of being produced by a reliable process, thereby satisfying a reliabilist desideratum, yet lack the property of resulting from the satisfaction of some doxastic obligation, thereby failing to satisfy a deontological desideratum. If we assume that justification (of whatever variety) entails knowledge, this seems to entail a sort of pluralism about knowledge.

The second is Ernest Sosa's (2007, 2009) distinction between animal and reflective knowledge. Sosa says that a belief qualifies as animal knowledge just in case it is *apt*, which means that it is correct because of an exercise of cognitive skill by the believer, whereas a belief qualifies as reflective knowledge just in case it is *meta-apt*, that is, aptly held to be apt, which means that the believer has acquired a perspective on her belief from which they can reflectively endorse the reliability of the faculties that produced it. Animal knowledge is relatively easy to obtain. Joanie has animal knowledge that she has hands so long as she correctly believes that she has hands as a result of the reliable functioning of her perceptual system. Reflective knowledge is a lot harder. Joanie may lack reflective knowledge that she has hands because she can't reflectively endorse the reliability of her perceptual system (she is plagued by

sceptical doubts). So, animal knowledge and reflective knowledge are two distinct relations that one can stand into a proposition *p*. Sosa's distinction leads to a kind of knowledge pluralism.[2]

2.2   Standards Pluralism

My focus in this chapter is on a kind of knowledge pluralism I call "standards pluralism". Put roughly, standards pluralism is the view that one can never know anything *simpliciter*. Rather, one only knows relative to this-or-that epistemic standard. Say that the set of epistemic standards is $\{E_1, E_2 \ldots E_N\}$. Standards pluralists hold that there is a plurality of knowledge relations of the form $K_{Ei}(S, p)$. In the rest of this section, I will say a little bit more about standards pluralism, explain how it differs from both Alston's and Sosa's brand of pluralism and outline the motivations behind it.

We can start with the idea of some subject S's "strength of epistemic position" with respect to some proposition *p*. How strong S's epistemic position with respect to *p* depends on things like how much evidence S has for *p*, how reliable S's belief-forming processes are and so on. How strong does S's epistemic position with respect to *p* need to be for S to know that *p*? The knowledge monist holds that there is a unique strength of epistemic position required (though monists will differ on what that strength is). Because there is a unique strength of epistemic position required for knowledge, one can talk about S as either knowing or not knowing *simpliciter*. If S's epistemic position is strong enough, they know; if it isn't, they don't. The standards pluralist holds that there cannot be a unique strength of epistemic position required for knowledge (for their reasons why, see below). Because there is not a unique strength required, one cannot talk about S as either knowing or not knowing *simpliciter*. However, one can talk about whether S's epistemic position is strong enough to satisfy this-or-that epistemic standard. The standards pluralist says that, if S's epistemic position satisfies standard E, then S knows relative to E. We can think of epistemic standards as lying on a continuum, with the highest standards requiring the strongest possible epistemic position for knowledge and lower standards requiring weaker epistemic positions.

Standards pluralism differs from Alston's brand of pluralism in that it is based on the thought that one can know relative to one epistemic standard but not relative to another, not the thought that there are different desiderata one might want one's account of knowledge (or justification) to meet. It would be wrong to construe "reliabilist knowledge" (knowledge as something produced by a reliable belief-forming process) as "knowledge by low standards" and "deontological knowledge" (knowledge as something requiring the meeting of doxastic obligations) as "knowledge by high standards" (or vice versa). These are different conceptions of what it is to know, not cognitive states that can be ranked in terms of the strength of epistemic position required to instantiate them. While one could construe Sosa's distinction between animal and reflective knowledge in terms of epistemic standards—animal knowledge is "low grade" knowledge, whereas reflective knowledge is "high grade" knowledge—the operative notion of epistemic standard is rather different from the standards pluralist's notion of epistemic standard. Reflective knowledge is "higher grade" in the sense that it involves a kind of perspective on one's first-order beliefs. One's epistemic position is stronger when one has reflective knowledge than when one has mere animal knowledge, but the difference is one of kind, not degree.

The main motivation for standards pluralism is that it can resolve three puzzles. First, consider this pair of cases:

LOW: Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit a check. It's not important that they do so, as they have no impending bills. But, as they drive past the bank, they notice that the lines inside are very long. Realising that it isn't very important that the check is deposited right away, Hannah says, 'I know that the bank will be open tomorrow. I was there just two weeks ago on Saturday morning. So we can deposit our check tomorrow morning'.

HIGH: Same setup, but here Hannah and Sarah have an impending bill and very little in their account, so it's very important that they deposit their check by Saturday. Hannah notes that she was at the bank two weeks before on a Saturday morning, and it was open. But, as Sarah

points out, banks do change their hours. Hannah says, 'I guess you're right. I don't know that the bank will be open tomorrow'. (Stanley 2005, pp. 3–4; for the original cases see DeRose 1992)

Intuitively, Hannah's knowledge ascription in LOW and her knowledge denial in HIGH are both true. But her epistemic position is no stronger in HIGH than in LOW. This raises a problem. If we assume that any given epistemic position is either strong enough or not strong enough for knowledge, we must choose: Hannah either knows in LOW and HIGH, or does not know in either. The standards pluralist solves the problem by rejecting this assumption. One can only know relative to this-or-that standard, and any given epistemic position can be strong enough to satisfy one standard but not strong enough to satisfy another. Thus, the standards pluralist doesn't need to choose. In LOW, Hannah and Sarah adopt an epistemic standard based on their needs, interests and purposes (they want to cash the check, but it isn't urgent). Call this standard E_{LOW}. Hannah's epistemic position is strong enough to satisfy E_{LOW}, so she knows that the bank is open on Saturdays relative to E_{LOW}. Now take HIGH. Hannah and Sarah adopt a different epistemic standard based on their different needs, interests and purposes (it is now urgent that they cash the check). Call this E_{HIGH}. Hannah's epistemic position is not strong enough to satisfy E_{HIGH}, so she does not know relative to E_{HIGH}. If we specify the different knowledge relations involved, the puzzle evaporates.

Second, consider this inconsistent triad:

(1)   Joanie knows that she has hands.
(2)   If Joanie knows that she has hands, she knows that she is not a handless brain in a vat.
(3)   Joanie does not know that she is not a handless brain in a vat.

Imagine Joanie has the usual kind of evidence that she has hands. The anti-sceptic argues, convincingly, that (1) is true. If Joanie doesn't know this, what does she know? The sceptic argues, also convincingly, that (3) is true. How could Joanie know that she is not a handless brain in a

vat, given that she seems unable to rule sceptical hypotheses out?[3] (2) is common-ground between the sceptic and the anti-sceptic.[4] This raises a problem. If we assume that any given epistemic position is either strong enough or not strong enough for knowledge, we must choose between anti-scepticism and scepticism. Joanie either does not know that she has hands, or she does know that she is not a handless brain in a vat. Again, the standards pluralist solves the problem by rejecting this assumption, and so "splits the difference" between the anti-sceptic and the sceptic. Who is "right" depends on which epistemic standard is appropriate given our needs, interests and purposes. Imagine that Joanie has recently had an accident and is checking that everything is in working order. She looks down and sees hand-shaped objects. Her epistemic position is strong enough to satisfy the epistemic standard appropriate to her purpose. So, she knows that she has hands relative to this standard. Now imagine that Joanie is discussing Descartes' *Meditations* with Chris. Chris is pushing the line that Joanie can't rule out Descartes' sceptical scenarios. Her epistemic position is not strong enough to satisfy the epistemic standard appropriate to this purpose. So, she does not know that she has hands relative to this standard. Where both the sceptic and anti-sceptic go wrong is in assuming that Joanie either knows or does not know *simpliciter*.

Third, consider what Laurence BonJour (2010) calls the "threshold problem". For the sake of argument, let's assume that knowledge is non-accidental justified true belief. A problem arises for those who want to combine two popular claims:

- Fallibilism: S may know that *p* on the basis of justification that does not guarantee the truth of *p*.
- Knowledge is a very valuable and desirable cognitive state.[5]

If knowledge does not require justification that guarantees the truth of the relevant belief, what level of justification does it require? Bonjour says that "it is very difficult or, I believe, impossible to see what could give any level of justification that is short of being conclusive the kind

of special significance that [fallibilism] requires it to have" (2010, p. 61). The point is not that we couldn't decide that there is some threshold t such that, if S truly believes that *p* on the basis of justification that is at least as strong as t, then S knows that *p*. Rather, the point is that any choice of t would be largely arbitrary, and so it would be implausible to hold that the difference between true belief based on justification that is at least as strong as t and true belief based on justification weaker than t marks a difference in value.[6] Thus, it is hard to combine fallibilism with the claim that knowledge is valuable. If, like BonJour, we assume that knowledge is valuable, we must conclude that fallibilism is false.

BonJour's discussion implicitly assumes that the fallibilist holds that there is a unique level of justification that is required for knowledge. The standards pluralist rejects this assumption and so can solve the problem. Even if there is no level of justification that is appropriate in all situations, there may be reasons to privilege particular levels in particular situations. These reasons would give particular levels of significance in some situations, but not necessarily any significance in other situations. Take Hannah and Sarah. In LOW, their needs, interests and purposes give them a reason to adapt a relatively lax epistemic standard such that Hannah's epistemic position is strong enough to satisfy it. However, in HIGH, their different needs, interests and purposes give them a reason to adopt a stricter epistemic standard such that Hannah's epistemic position is not strong enough to satisfy it. On the plausible assumption that people generally have reasons to privilege one standard over another in particular situations, the problem vanishes.[7]

All of these arguments need further development, but I lack the space to pursue this here. My claim is just that standards pluralism offers *prima facie* plausible responses to all three puzzles. This is not to say that they are *ultima facie* plausible, or that the standards pluralist responses are preferable to alternatives.[8] One thing that is striking is how closely the standards pluralist responses parallel the contextualist responses to these puzzles. In the next section, I clarify the relationship between standards pluralism and two other views in epistemology: contextualism and epistemic relativism.

3   Standards Pluralism, Contextualism and Epistemic Relativism

I will argue for three claims. First, while the standards pluralist isn't forced to adopt a contextualist semantics for the word "knows", they are forced to adopt some kind of non-standard semantics, like contextualism. Second, while contextualism does not entail standards pluralism, it is good evidence that standards pluralism is true. Third, while both standards pluralism and epistemic relativism relativise knowledge to something (standards and systems), they differ in that the standards pluralist does not hold that standards cannot be ranked. I will take each claim in turn.

3.1   Knowledge Pluralism and the Semantics of Knowledge Ascriptions

There is an important structural difference between truth pluralism and standards pluralism. According to the truth pluralist, sentences are either true or false, but different properties "make for" truth in different domains of discourse. For instance, Lynch (2009) proposes a functionalist version of truth pluralism according to which different properties play the "truth role" in different domains of discourse, in the same way that different physical states might play the "pain role" for different biological species. Perhaps, in discourse about "medium-sized dry goods", correspondence plays the relevant functional role, whereas in discourse about what is funny, the role is played by what Wright (1992) calls "super-assertability". In contrast, the standards pluralist does not think that different relations "make for" knowledge in different domains of discourse. Rather, the view is that, no matter what the domain is, one can only know relative to this-or-that standard.

This structural difference raises a problem. We talk about subjects *knowing that p*, not about subjects *knowing that p relative to this-or-that standard*. What does the standards pluralist say about knowledge ascriptions of the form "S knows that *p*"? Consider:

(4)   Joanie knows that she has hands.
(5)   Joanie does not know that she has hands.

Does the standards pluralist think that (4) is true? If so, on pain of contradiction they must think that (5) is false. But the idea behind the response to the sceptical puzzle was that (4) is sometimes true and other times false. Clearly, the standards pluralist needs a story about the semantics of knowledge ascriptions such that (4) can sometimes be true and other times be false. In principle, any number of views about the semantics of knowledge ascriptions would serve the standards pluralist's purposes here.[9] For the sake of simplicity, I will opt for a contextualist semantics according to which the word "knows" refers to different relations in different contexts. Take a context where "everyday" epistemic standards are operative. Utterances of (4) and (5) in such contexts refer to the relation of knowing relative to everyday standards, and because Joanie knows relative to everyday standards, utterances of (4) are true and utterances of (5) are false in such contexts. Now take a context where "sceptical" standards are operative. Utterances of (4) and (5) in such contexts refer to the relation of knowing relative to sceptical standards, and, because Joanie does not know relative to sceptical standards, utterances of (4) are false and utterances of (5) are true in such contexts.

3.2   Contextualism as Evidence for Standards Pluralism

Standards pluralism is a view about the knowledge relation, whereas contextualism is a view about the word "knows". To get a handle on this difference, compare what the two views say about LOW and HIGH. The standards pluralist says that there are two relations, knowing relative to E_{LOW} and knowing relative to E_{HIGH}. Hannah knows relative to E_{LOW} that the bank is open on Saturdays, but does not know relative to E_{HIGH}. The contextualist says that Hannah's knowledge ascription in LOW and knowledge denial in HIGH don't express incompatible propositions. In LOW, she expresses something like the proposition *Hannah knows that the bank is open on Saturdays by E_{LOW}*; in HIGH, she expresses the

proposition *Hannah does not know that the bank is open on Saturday by E_{HIGH}*. Her epistemic position is such that both propositions are true.[10]

The contextualist's view about the propositions expressed by Hannah's knowledge ascription and denial doesn't entail that there are two relations: knowing relative to E_{LOW} and knowing relative to E_{HIGH}. But, if contextualism is true, it puts pressure on the view that there is a single two-place knowledge relation. If there were a single two-place knowledge relation, then Hannah would either know that the bank is open on Saturdays, or she would not know. But, if Hannah knows, it should be true to say that she "knows" in any context, and if she doesn't know, it should be true to say that she doesn't "know" in any context. The only way to avoid this would be to hold that, even though there is a single two-place knowledge relation, uses of the word "knows" don't always refer to it.[11] But this would mean divorcing the semantics of knowledge ascriptions from the metaphysics of knowledge.

One can buttress this argument by considering the contextualist response to the sceptical problem. Like the standards pluralist, the contextualist tries to "split the difference" between the sceptic and the anti-sceptic. In some contexts, it can be true to say that Joanie "knows" that she has hands, whereas in other contexts, it can be true to say that Joanie does not "know" that she has hands.[12] Now imagine there was just a single two-place knowledge relation. Presumably, Joanie either knows that she has hands, or she doesn't. If she knows, scepticism is false. If she doesn't, it is true. Either way, it is unclear how the contextualist can "split the difference" between the sceptic and the anti-sceptic, irrespective of the correct semantics for the word "knows".[13]

I conclude that, while contextualism doesn't entail that standards pluralism is true, it is good evidence that standards pluralism is true.

3.3   Standards Pluralism and Epistemic Relativism

By epistemic relativism, I mean a metaphysical view about epistemic properties and relations (as in e.g. Boghossian 2006) rather than a view about the semantics of sentences involving epistemic vocabulary (as in e.g. MacFarlane 2014). Discussions of this kind of epistemic relativism usually focus on a view about epistemic justification, according to

which one's beliefs can never be justified *simpliciter*, but only relative to this-or-that epistemic system, and there is a plurality of epistemic systems, all of equal validity (see, for instance, Boghossian 2006; Kusch 2016; Seidel 2014; Williams 2007). But one can consider a structurally analogous view about knowledge, according to which one never knows *simpliciter*, but only relative to this-or-that epistemic system, and there is a plurality of epistemic systems, all of equal validity. How does standards pluralism differ from this view?

There is a relatively superficial difference. Where the epistemic relativist talks of epistemic systems, the standards pluralist talks of epistemic standards. We can think of epistemic systems as consisting in sets of principles that encode acceptable (and unacceptable) epistemic practice. For instance, our epistemic system tells us that we are entitled to rely on the evidence of our senses absent reason to doubt their reliability. Alternative epistemic systems might tell their adherents that they are entitled to rely on the deliverances of oracles. The crucial contrast between epistemic systems and epistemic standards is that, where epistemic standards mark different degrees or levels of justification that beliefs can enjoy, epistemic systems constitute frameworks against which we can talk of beliefs as being justified in the first place.

There is also a deeper difference. The epistemic relativist holds that all justification is system-relative, and, because it is impossible to provide a non-circular justification for any epistemic system, all these systems are equally valid. The standards pluralist holds that all knowledge is standards-relative. But the kind of standards pluralism considered here does not hold that all these standards are equally valid, or at least it doesn't hold that they are all equally valid in any particular situation. Rather, in any given particular situation, some standards are appropriate, whereas others are inappropriate. For instance, given Hannah and Sarah's situation in LOW, a low epistemic standard is appropriate, not a high standard; given their different situation in HIGH, a high standard is appropriate, not a low standard. But this requires that all standards are not equally valid. One could say that all standards are appropriate in *some* situation, but not all standards are appropriate in all situations. While one could consider a version of standards pluralism on which all standards are equally valid, it is unclear whether one could argue for it in the way I have argued for standards pluralism here.

4   The Case Against Standards Pluralism

In this section, I argue that standards pluralism faces a serious objection. The standards pluralist posits a plurality of knowledge relations. But what makes these knowledge relations rather than some other kind of relation? One would expect the answer is that they all have something in common: they satisfy certain conceptual constraints on what can count as a knowledge relation. The objection is that it is, at best, unclear whether the relations posited by the standards pluralist meet these constraints. So, it is unclear whether standards pluralism really offers us a theory of knowledge. This challenge parallels a challenge that Michael Lynch ([2013](#)) raises for the truth pluralist. So, I start with Lynch's articulation of his challenge.

4.1   What Makes a Theory of Truth a Theory of Truth?

Truth pluralism is the view that, while there are properties the instantiation of which makes sentences true, there is a plurality of these properties. But in virtue, of what are all of these properties *truth* properties, as opposed to some other kind of property? Lynch's thought is that there are conceptual constraints on the kinds of properties that the theorist can identify with our pre-theoretic notion of truth. These constraints are captured in three "core platitudes" about truth:

- Objectivity (OBJ): True propositions are propositions such that, when we believe them, things are as we believe them to be.
- Aim of inquiry (AIM): True propositions are propositions such that we should aim to believe them when engaging in inquiry.
- Norm of Belief (NORM): True propositions are propositions such that they are correct to believe.

A theory that identifies truth with some property T or set of properties $\{T_1, T_2 \ldots T_N\}$ is a theory of truth iff propositions that instantiate T or one of $\{T_1, T_2 \ldots T_N\}$ satisfy OBJ, AIM and NORM. Thus, Lynch's

challenge for the truth pluralist is to show that propositions that instantiate one of $\{T_1, T_2 \ldots T_N\}$ satisfy OBJ, AIM and NORM. If they can't do that, then it is unclear in what sense they have given us a theory of truth, as opposed to a theory of something else.

Lynch motivates his challenge by appealing to Frank Jackson's (1998) view that our tacit beliefs about the object of a metaphysical inquiry can be codified in a set of platitudes, which put constraints on the subsequent theories about that object that can be developed in the course of inquiry. One might worry that, for any philosophically interesting notion X, there are very few genuine platitudes about X. Take truth. Is it a platitude that truth is the aim of inquiry or the norm of belief? Some have explicitly denied that truth is the aim of inquiry or norm of belief (e.g. Rorty 1995). While we can accuse them of conceptual confusion, or of just talking about something different to the rest of us, these accusations feel strained. Aren't we just fighting about whether to call something a theory of truth or not?

While I admit that this worry has some force, I don't think it necessarily causes a problem for Lynch's challenge. One can tweak Jackson's picture to avoid a commitment to platitudes. Rather than think of OBJ, AIM and NORM as platitudes, we can think of them as principles that codify why we think truth is theoretically important. Truth is theoretically important because of its connection to objectivity (OBJ) and because of its normative role (AIM and NORM). We may *call* a theory that does not satisfy these platitudes a theory of truth, but it is a theory of truth on which truth is of little theoretical interest. If *that* is what truth is, we might as well stop talking about it. So, one can think of Lynch's challenge as a challenge to the theorist to demonstrate that they are giving us an account of some philosophically interesting notion rather than reasons to abandon that notion. While I will stick with Lynch's own understanding of his challenge, I am happy to retreat to this alternative formulation if the reader regards it as necessary.

I will now construct a parallel challenge for the standards pluralist.

4.2   What Makes a Theory of Knowledge a Theory of Knowledge?

Just as there are conceptual constraints on the kinds of properties that the theorist can identify with our pre-theoretic notion of truth, there are conceptual constraints on the kinds of relations that the theorist can identify with our pre-theoretic notion of knowledge. If the plurality of relations posited by the standards pluralist doesn't satisfy these constraints, then it is unclear in what sense they have given us a theory of knowledge, as opposed to a theory of something else.

While I think one can motivate Lynch's challenge on these kinds of general grounds, I want to emphasise that the challenge is particularly pressing for the standards pluralist. Recall that the standards pluralist thinks that there is no such thing as knowing something *simpliciter*. One might think that if there is no such thing as knowing *simpliciter*, we should eliminate the notion of knowledge. But the standards pluralist does not propose eliminating the notion of knowledge. Rather, they propose holding that there is a plurality of knowledge relations. Thus, the burden is on the standards pluralist to show that this is a *replacement* rather than an *elimination* of our standard picture of knowledge. If the standards pluralist can answer Lynch's challenge, they have shown this is a replacement. If they can't, the suspicion is that the standards pluralist has really eliminated knowledge and replaced it with something else entirely.

The first step is to identify the platitudes governing knowledge. Some of these platitudes are of little interest for my purposes. Take, for instance, the platitude that knowledge is factive. This poses no problem for the standards pluralist. The standards pluralist can hold that, while there is a plurality of knowledge relations, all of these relations are factive: if S knows that p relative to some standard E, then p is true. I am more interested in these three platitudes about the "normative role" of knowledge:

>   *INQ*:   If S knows that p at t, S may stop inquiring into whether p at t (see Kappel 2010; Kelp 2011; Rysiew 2012).

ACT: If S knows that *p* at t, then, usually, S may treat *p* as a reason for acting at t.[14]

ASS: If S knows that *p* at t, then, usually, S may assert that *p* at t.[15]

The basic thought behind INQ is that knowledge is the natural stopping-point of inquiry. Once an inquiry has yielded knowledge, you may stop. The basic thought behind ACT and ASS is that, if you know some proposition *p*, there is usually no epistemic barrier to you treating *p* as a reason for acting or asserting that *p*. You may (in the epistemic sense of "may") treat *p* as a reason for acting or assert that *p*. I want to make two comments here. First, I have chosen these "weaker" formulations of ACT and ASS because some have argued that there are cases where knowing some *p* is insufficient for properly treating *p* as a reason for acting or asserting that *p*. But these cases are unusual.[16] Second, for the sake of simplicity, I will assume that there are no other epistemic norms governing inquiry, practical reasoning and assertion. So, knowing that *p* at t is necessary as well as sufficient for being permitted to terminate inquiry into whether *p* at t, and knowledge is typically required for proper practical reasoning and assertion.

The second step is to raise Lynch's challenge. There is an initial problem here. The standards pluralist can't accept these platitudes as they stand because they talk about knowledge *simpliciter*, rather than knowing relative to this-or-that standard. I am going to argue that this initial problem leads to a bigger problem. While there are a number of ways in which the standards pluralist can reformulate the platitudes, all of these ways are either (a) independently implausible or (b) yield platitudes that the vast majority of knowledge relations posited by the standards pluralist fail to satisfy. Thus, the standards pluralist either holds an implausible view about the normative roles of knowledge, or fails to satisfy Lynch's challenge.

Here are three options for the standards pluralist. To make things simple, I focus on INQ. But similar remarks can be made about ASS and ACT. They are:

INQ*:   If S knows that *p* relative to *some* standard at t, then S may stop inquiring into whether *p* at t.

*INQ*****:   If S knows that *p* relative to *every* standard at t, then S may stop inquiring into whether *p* at t.

*INQ******:   If S knows that p relative to the *relevant* standard at t, then S may stop inquiring into whether *p* at t.

To evaluate these proposals, it will be helpful to consider some cases:

> TRAIN STATION: Catriona is getting the train from Edinburgh to Glasgow. She wants to get the fast train, but it isn't particularly important that she does so. Because the fast train has always left from platform 7 and she hasn't heard or seen any announcements to the contrary, she heads to platform 7 and boards the train.

> RISKY TRAIN STATION: Same as TRAIN STATION, except that it is vitally important that Catriona catch the fast train because she has a very important appointment that she will otherwise miss.

It seems clear that Catriona may terminate her inquiry into whether the train leaves from platform 7 in TRAIN STATION (she can board the train without asking) but may not terminate her inquiry in RISKY TRAIN STATION (she needs to find out more before boarding). Catriona's epistemic position with respect to the proposition that the fast train leaves from platform 7 is the same in both cases. But, while it is strong enough for the purposes at hand in TRAIN STATION, it is not strong enough for the purposes at hand in RISKY TRAIN STATION. Call the proposition that the fast train leaves from platform 7 "TRAIN", the epistemic standard appropriate in TRAIN STATION "E_{TS}" and the standard appropriate in RISKY TRAIN STATION "E_{RTS}". Catriona knows TRAIN relative to E_{TS}, but she does not know TRAIN relative to E_{RTS}. So, in RISKY TRAIN STATION, she knows relative to *some* standard—E_{TS}—but, clearly, can't stop inquiring into whether TRAIN is true. Thus, INQ* is false. The problem with INQ* is that it is far too weak. A subject can stand in *a* knowledge relation to some proposition *p* yet be in no position to stop inquiring into whether *p*.

The problem with INQ** is that it is far too strong. Remember we are assuming that INQ** is the only epistemic norm governing inquiry.

Thus, it entails that one may terminate inquiry into whether *p* *only* if one knows that *p* relative to any standard whatsoever. The upshot is that we may rarely, if ever, terminate inquiry. Consider this case:

> BETTER TRAIN STATION: Same as TRAIN STATION, except that Catriona has asked the ticket inspector who has told her that this is indeed the fast train.

Plausibly, Catriona both knows relative to the standards appropriate in BETTER TRAIN STATION (whatever else could she do?) and may terminate her inquiry into whether TRAIN is true. But she does not know relative to any standard whatsoever. Consider "sceptical standards", that is, standards that require Catriona to rule out scenarios in which the ticket inspector is part of a vast conspiracy to trick Catriona into taking the wrong train. Despite having excellent evidence for TRAIN, Catriona can't rule this possibility out. So, INQ** is implausibly strong.

The obvious fix to both problems is to require that the subject stand in the *relevant* knowledge relation (INQ***). At a rough first pass, we can say that the knowledge relation that is relevant for a subject S at a time t is the relation of knowing relative to whatever epistemic standards are appropriate for S at t. The thought is that something about S's practical situation "picks out" a single knowledge relation as most appropriate. So, the relevant knowledge relation for Catriona in RISKY TRAIN STATION is E_{RTS}, not E_{TS}. Because Catriona does not know relative to the relevant epistemic standard, it is not the case that Catriona may terminate her inquiry. INQ*** gets the right result.

Even though INQ*** gets the right result, it is unclear whether the standards pluralist who endorses it can answer Lynch's challenge. Remember that the platitudes put conceptual constraints on what can count as a knowledge relation. Platitudes like INQ tell us that knowledge relations have a particular normative role or profile. For instance, INQ tells us that the knowledge relation is a relation such that, when one stands in it to some proposition *p* at a time t, one may terminate inquiry into whether *p* at t. So, one would expect that the pluralist about knowledge posits a set of knowledge relations all of which are such that, when one stands in them to some *p* at t, one may terminate

inquiry into whether p at t. But INQ*** tells us that only the relation of knowing by the relevant epistemic standard is such that, when one stands in it to some p at t, one may terminate inquiry into whether p at t. So, at any given time, many of the knowledge relations posited by the standards pluralist don't play the normative roles of knowledge relations. But this means that the standards pluralist can't answer Lynch's challenge, and so it is unclear in what sense they offer a theory of knowledge, rather than a theory of something else.[17]

In conclusion, no matter how the standards pluralist tries to formulate INQ, they face problems. Either the reformulations are implausible, or some of the relations they posit fail to satisfy them. Similar considerations apply to ACT and ASS. I am going to finish by looking at how the standards pluralist could try to respond to this problem.[18]

5   Sketch of a Response

The basic idea behind ACT, ASS and INQ is that knowing that p gives one certain rights—to treat p as a reason for acting, to assert that p, to terminate one's inquiry into whether p—in a very wide range of situations. The problem for the standards pluralist is that, if you know that p relative to some standard E, this only gives you these rights in a narrower range of situations, viz. situations where E is the appropriate standard. The only way that the standards pluralist can salvage the basic idea behind ACT, ASS and INQ is if they can find a way to expand the range of situations in which knowing relative to some standard E gives one of these rights. I am going to finish by suggesting how the standards pluralist might be able to do this. As before, I will focus on INQ, but similar considerations apply to ACT and ASS.

We can start by thinking about other properties and relations that give one certain rights in a wide, but still circumscribed, set of situations. Take the relationship of being someone's legal guardian. If X is Y's legal guardian, then X has various rights. For instance, X has the right to make certain decisions on Y's behalf; what school Y goes to, what medical treatments they get, etc. But legal guardianship is a status that is bestowed on one by a legal system or framework. Because

different countries have different legal systems, it may be that X is recognised as Y's legal guardian by one system, yet isn't recognised as Y's legal guardian by another system. Thus, the legal guardianship relation gives one certain rights in a wide, but still circumscribed, set of situations. Roughly, one enjoys these rights when and only when one is under the jurisdiction of the legal system that recognises one as someone's legal guardian. For instance, imagine Catriona is recognised as Morven's legal guardian in country 1, where she has adopted Morven together with her same-sex partner. While she lives in this country, she has various rights towards Morven. However, she takes a new job in country 2, where same-sex adoption isn't legal, and so she is not recognised as Morven's legal guardian in country 2. Thus, she lacks these rights in this new country.

The standards pluralist might suggest that we should think of knowledge as being a little like legal guardianship. So far, I have talked as if individuals determine epistemic standards, and as if the standards pluralist holds there are as many knowledge relations as there are situations that individuals find themselves in, e.g. knowing relative to the standards of LOW, of HIGH, etc. But we could instead say that it is *communities* that determine epistemic standards. If this is right, then there are as many knowledge relations as there are distinct communities, and subjects may know relative to one community, but not know relative to another.

The underlying idea here is that each of us is a part of a complicated epistemic world. We move from epistemic community to epistemic community. In some of these communities, we enjoy rights that we lack in others. But these rights, and our implicit recognition of them, are key to our ability to navigate the world around us. Take Neil, a layperson who is reasonably well informed about climate science, and the proposition that sea levels have risen 6.7 inches over the past century as a result of human-induced global warming ("SCARY" for short).[19] Neil knows SCARY relative to the epistemic standards governing general public discourse about climate change, but he does not know SCARY relative to the rather more stringent standards governing scientific discourse. This means that, as a member of the general public, he need not inquire further into whether SCARY is true; he has done all that he needs. But, if

he wanted to engage with the scientific community and enter into scientific discussion, he would need to do a lot more inquiring. Thus, that he knows relative to the standards governing public discourse means he enjoys a right—to terminate his inquiry—which obtains in a wide but still circumscribed range of situations. He enjoys this right when "speaking as a member of the public", but would not enjoy it if he tried to present himself as a scientist.

If this works, the standards pluralist at least has the beginnings of a solution to the problem. Maybe she can't honour the letter of the platitudes ACT, ASS and INQ. But she can do justice to the idea that knowing that p gives one certain rights in a wide range of situations. Remember that the Jacksonian methodology that Lynch's challenge appeals to permits modifying the platitudes. Of course, any modification must be independently motivated, and, in general, minor modifications are to be preferred to major modifications. But the case for standards pluralism provides some motivation for modifying the platitudes to fit with the standards pluralist's picture of knowledge, and I have just argued that the necessary modifications do some justice to the original idea behind the platitudes.

One might object that, while this might help with Lynch's challenge, it somewhat undermines the original motivations for standards pluralism. In particular, on any plausible account of what communities are, Hannah and Sarah are in the same community in both LOW and HIGH. So what about the argument that the standards pluralist can resolve this puzzle? I would like to make two points in response. First, even if the point is granted, this does not touch the standards pluralist solution to the second and third puzzles. The second puzzle concerns our conflicting "sceptical" and "anti-sceptical" intuitions. The standards pluralist response is that different communities may choose to favour more or less sceptical standards, depending on their particular needs, interests and purposes. The third puzzle concerns how the "threshold" for knowledge is set. The standards pluralist response is that different communities may set the threshold in different places, also depending on their particular needs, interests and purposes.

Second, it is not clear that the point should be granted. The first puzzle is similar to the second and third in that it raises the question of where, precisely, the standards for knowledge should be set. Should they be set such that Hannah knows that the bank is open, or should they be set such that she does not know? Different communities may legitimately give different answers to this question. Which answer they give will depend on their needs, interests and purposes. For instance, are they often in "high stakes" environments? Put very roughly, the more often a community's members find themselves in high stakes environments, the higher their standards for knowledge are likely to be.

6   Concluding Remarks

It is time to take stock. I have argued that one can construct a viable case for standards pluralism. This case, in large part, parallels the case for contextualism. I have also argued that while neither view entails the other, standards pluralism and contextualism are mutually supporting: standards pluralism requires a non-standard semantics such as contextualism, and contextualism is good evidence for standards pluralism. I then turned to the case against standards pluralism. The basic problem is that standards pluralism posits knowledge relations that don't have the right kind of normative profile. This means it is unclear in what sense the standards pluralist offers us a theory of knowledge, rather than an elimination of knowledge. Finally, I gave a rough sketch of a standards pluralist response to this problem.

My aim in this chapter has been to explore a neglected question: Is there a single knowledge relation, or is there a plurality of knowledge relations? I have outlined both the case for thinking there is a plurality of knowledge relations, and the case against. I hope that this chapter will lead to further work on how the considerations appealed to in favour of contextualism might be a reason to modify our metaphysical picture of knowledge, as well as a reason to modify our standard assumptions about the semantics of knowledge ascriptions.[20]

Notes

1. There is some dissent on this (see, e.g. Schaffer 2004a). I comment on this further in fn. 11 and 17.
2. While Sosa denies that the word "knows" is ambiguous (see 2009, p. 135), this is no reason to deny that Sosa is a knowledge pluralist. Many truth pluralists deny that "true" is ambiguous (see Alston 2002; Lynch 2009; Wright 1996).
3. Of course, this is controversial. I am following the standard set-up of the sceptical problem in the contextualist literature (see, e.g. DeRose 1995).
4. That is, they both accept the principle that knowledge is closed under known entailment: if S knows that p, and knows that p entails q, then S knows that q. For defences of closure principles, see DeRose (1995), Hawthorne (2004), Williamson (2000). For criticisms, see Dretske (2005).
5. Of course, some epistemologists deny fallibilism (for instance, Williamson 2000), and some deny that knowledge is particularly valuable or desirable (for instance, Kaplan 1985). But Williamsonian infallibilism is highly controversial (see Brown 2013) and, while BonJour talks as if the problem requires holding that knowledge is the supremely valuable and desirable cognitive state, it is enough to hold that knowledge is merely very valuable.
6. This is why the threshold problem isn't just about vagueness. Even if there were a sharp boundary between knowledge and its absence, that boundary would be arbitrary, and so wouldn't support a difference in value between knowledge and cognitive states that fall just short of it.
7. This is just a sketch of the standards pluralist's response to BonJour's threshold problem. For a more developed account that I largely agree with, see Grimm (2015), Hannon (2017).
8. For some alternatives, see Brown (2006), Fantl and McGrath (2009), Gerken (2011), Hawthorne (2004), MacFarlane (2014, Chap. 8), Nagel (2008), Rysiew (2001), Stanley (2005).
9. The standards pluralist could endorse the kind of Austinian semantics for knowledge ascriptions proposed in Lawlor (2013), the non-indexical contextualism proposed in Brogaard (2008), Kompa (2002) or MacFarlane (2009), and maybe even the assessment relativism proposed in MacFarlane (2014). I lack the space to discuss these options here.

10. For versions of this account, see Cohen (1999) and DeRose (2009).

11. That is, unless the contextualist denies that the knowledge relation is a two-place relation, and proposes instead that is a three-place relation of the form K(S, p, E_i), where E_i is a parameter fixed by the context (e.g. an epistemic standard, a contrast class). Hannah's ascription and denial could then refer to the same relation, but in LOW, the third argument place is filled by E_{LOW}, whereas in HIGH, it is filled by E_{HIGH}. Though most contextualists don't take this route, some do (e.g. Schaffer 2004a; Schaffer and Szabo 2013). For further discussion, see fn. 17.

12. For versions of this account, see Blome-Tillmann (2014), DeRose (1995), Lewis (1996).

13. For a similar argument, see Schaffer (2004b).

14. For defences of somewhat stronger versions of ACT, see Hawthorne and Stanley (2008), Williamson (2005).

15. For a defence of a somewhat stronger version of ASS, see Williamson (2000).

16. For instance, these are cases where the "stakes" are unusually high, or one is required to follow certain institutional procedures. For such cases, see Brown (2008a, 2008b), Gerken (2011), Lackey (2007).

17. Does the problem arise if there is a single three-place relation of the form K(S, p, E_i)? According to this view, the relation of knowing relative to the standard E_{TS} and the relation of knowing relative to the standard E_{RTS} are not distinct relations. Rather, they are instances of the same relation with different relata. I think that it does. The key question is: What makes all of these instances *knowledge* relations? Given Lynch's challenge, the answer must be that they are knowledge relations because they satisfy certain conceptual constraints. But the argument just given tells us that only one instance of this single relation—knowing relative to the relevant epistemic standard—is such that, if one stands in it to some p at t, one may terminate inquiry into whether p at t. So, at any given time, many of the instances of this single relation don't satisfy the relevant conceptual constraints.

18. It is worth noting that, while contextualists face a parallel problem, the standard solutions to that problem won't help the standards pluralist. Hawthorne (2004) argues that contextualists licence the truth of "abominable conjunctions" such as "Catriona knows that it is the fast train, but she may not assert that it is the fast train" because they disconnect the truth-conditions of knowledge ascriptions (which

are tied to the context of utterance) from the propriety of assertion and practical reasoning (which is tied to the situation of the subject). Contextualists have dealt with this problem by "explaining away" the propriety of these conjunctions (see Blome-Tillmann 2013). Because this solution is geared towards explaining linguistic data, it is of no use in dealing with Lynch's challenge.

19. http://blogs.scientificamerican.com/cross-check/climate-change-facts-versus-opinions/.

20. Thanks to Natalie Ashton, Delia Belleri, Tom Fery, Anna-Maria Asunta Eder, Michael Hannon, Dirk Kindermann, Katherina Kinzel, Anne-Kathrin Koch, Martin Kusch, Thomas Raleigh, Niels Wildschut and audiences in KU Leuven and Warwick. Research on this chapter was assisted by funding from the ERC Advanced Grant Project "The Emergence of Relativism" (Grant No. 339382).

References

Alston, William. 1993. Epistemic Desiderata. *Philosophy and Phenomenological Research* 53 (3): 527–551.

———. 2002. Truth: Concept and Property. In *What Is Truth?* ed. R. Shantz, 11–26. Berlin: Walter de Gruyter.

———. 2005. *Beyond Justification: Dimensions of Epistemic Evaluation*. Ithaca: Cornell University Press.

Blome-Tillmann, Michael. 2013. Contextualism and the Knowledge Norms. *Pacific Philosophical Quarterly* 94 (1): 89–100.

———. 2014. *Knowledge and Presuppositions*. Oxford: Oxford University Press.

Boghossian, Paul. 2006. *Fear of Knowledge: Against Relativism and Constructivism*. Oxford: Oxford University Press.

BonJour, Laurence. 2010. The Myth of Knowledge. *Philosophical Perspectives* 24 (1): 57–83.

Brogaard, Berit. 2008. In Defence of a Perspectival Semantics for "Know". *Australasian Journal of Philosophy* 86 (3): 439–459.

Brown, Jessica. 2006. Contextualism and Warranted Assertibility Manoeuvres. *Philosophical Studies* 130 (3): 407–435.

———. 2008a. Knowledge and Practical Reason. *Philosophy Compass* 3 (6): 1135–1152.

———. 2008b. The Knowledge Norm for Assertion. *Philosophical Issues* 18 (1): 89–103.

———. 2013. Infallibilism, Evidence and Pragmatics. *Analysis* 73 (4): 626–635.

Cohen, Stewart. 1999. Contextualism, Skepticism, and the Structure of Reasons. *Philosophical Perspectives* 13: 57–89.

DeRose, Keith. 1992. Contextualism and Knowledge Attributions. *Philosophy and Phenomenological Research* 52 (4): 913–929.

———. 1995. Solving the Skeptical Problem. *Philosophical Review* 104 (1): 1–52.

———. 2009. *The Case for Contextualism: Knowledge, Skepticism, and Context.* Vol. 1. Oxford: Oxford University Press.

Dretske, Fred. 2005. Is Knowledge Closed Under Known Entailment? The Case Against Closure. In *Contemporary Debates in Epistemology*, ed. Matthias Steup, and Ernest Sosa, 13–26. Malden: Blackwell.

Fantl, Jeremy, and Matthew McGrath. 2009. *Knowledge in an Uncertain World*. Oxford: Oxford University Press.

Gerken, Mikkel. 2011. Warrant and Action. *Synthese* 178 (3): 529–547.

Grimm, Stephen R. 2015. Knowledge, Practical Interests, and Rising Tides. In *Epistemic Evaluation: Purposeful Epistemology*, ed. John Greco, and David Henderson. Oxford: Oxford University Press.

Hannon, Michael. 2017. A Solution to Knowledge's Threshold Problem. *Philosophical Studies* 174 (3): 607–629.

Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.

Hawthorne, John, and Jason Stanley. 2008. Knowledge and Action. *Journal of Philosophy* 105 (10): 571–590.

Jackson, Frank. 1998. *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Oxford University Press.

Kaplan, Mark. 1985. It's Not What You Know That Counts. *Journal of Philosophy* 82 (7): 350–363.

Kappel, Klemens. 2010. On Saying That Someone Knows: Themes from Craig. In *Social Epistemology*, ed. Adrian Haddock, Alan Millar, and Duncan Pritchard. Oxford: Oxford University Press.

Kelp, Christoph. 2011. What's the Point of "Knowledge" Anyway? *Episteme* 8 (1): 53–66.

Kompa, Nikola. 2002. The Context Sensitivity of Knowledge Ascriptions. *Grazer Philosophische Studien* 64 (1): 1–18.

Kusch, Martin. 2016. Epistemic Relativism, Scepticism, Pluralism. *Synthese*, 1–17.

Lackey, Jennifer. 2007. Norms of Assertion. *Noûs* 41 (4): 594–626.

Lawlor, Krista. 2013. *Assurance: An Austinian View of Knowledge and Knowledge Claims*. Oxford: Oxford University Press.

Lewis, David. 1996. Elusive Knowledge. *Australasian Journal of Philosophy* 74 (4): 549–567.

Lynch, Michael. 2009. *Truth as One and Many*. Oxford: Oxford University Press.

———. 2013. Three Questions for Truth Pluralism. In *Truth and Pluralism: Current Debates*, ed. Nikolaj Jang Lee Linding Pedersen and Cory D. Wright, 21–41. Oxford: Oxford University Press.

MacFarlane, John. 2009. Nonindexical Contextualism. *Synthese* 166 (2): 231–250.

———. 2014. *Assessment Sensitivity: Relative Truth and Its Applications*. Oxford: Oxford University Press.

Nagel, Jennifer. 2008. Knowledge Ascriptions and the Psychological Consequences of Changing Stakes. *Australasian Journal of Philosophy* 86 (2): 279–294.

Pedersen, Nikolaj Jang, Lee Linding, and Cory D. Wright (eds.). 2013. *Truth and Pluralism: Current Debates*. Oxford: Oxford University Press.

Rorty, Richard. 1995. Is Truth a Goal of Enquiry? Davidson vs. Wright. *Philosophical Quarterly* 45 (180): 281–300.

Rysiew, Patrick. 2001. The Context-Sensitivity of Knowledge Attributions. *Noûs* 35 (4): 477–514.

———. 2012. Epistemic Scorekeeping. In *Knowledge Ascriptions*, ed. Jessica Brown, and Mikkel Gerken. Oxford: Oxford University Press.

Schaffer, Jonathan. 2004a. From Contextualism to Contrastivism. *Philosophical Studies* 119 (1–2): 73–104.

———. 2004b. Skepticism, Contextualism, and Discrimination. *Philosophy and Phenomenological Research* 69 (1): 138–155.

Schaffer, Jonathan, and Zoltan Gendler Szabo. 2013. Epistemic Comparativism: A Contextualist Semantics for Knowledge Ascriptions. *Philosophical Studies* 2: 1–53.

Seidel, Markus. 2014. *Epistemic Relativism: A Constructive Critique*. Palgrave-Macmillan: London.

Sosa, Ernest. 2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, vol. I. Oxford: Oxford University Press.

———. 2009. *Reflective Knowledge: Apt Belief and Reflective Knowledge*, vol. II. Oxford: Oxford University Press.

Stanley, Jason. 2005. *Knowledge and Practical Interests*. Oxford: Oxford University Press.

———. 2011. *Know How*. Oxford: Oxford University Press.

Williams, Michael. 2007. Why (Wittgensteinian) Contextualism Is Not Relativism. *Episteme* 4 (1): 93–114.

Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford: Oxford University Press.

———. 2005. Knowledge, Context, and the Agent's Point of View. In *Contextualism in Philosophy: Knowledge, Meaning, and Truth*, ed. Gerhard Preyer and Georg Peter, 91–114. Oxford: Oxford University Press.

Wright, Crispin. 1992. *Truth and Objectivity*. Cambridge: Harvard University Press.

———. 1996. Precis of "Truth and Objectivity" and Response to Commentators. *Philosophy and Phenomenological Research* 56 (4): 911–941.

Part III

Epistemic Pluralism and Epistemic Relativism

Epistemic Relativism and Pluralism

Martin Kusch

1   Introduction

In his influential anti-relativist treatise, *Fear of Knowledge*, Paul Boghossian suggests that epistemic relativism is committed to a principle he calls 'Epistemic Pluralism' (henceforth 'EP') and formulates as follows:

> There are many fundamentally different, genuinely alternative epistemic systems, but no facts by virtue of which one of these systems is more correct than any of the others.(2006, p. 73)

Boghossian claims to find EP, for example, in Richard Rorty's *Philosophy and the Mirror of Nature* (2006, p. 61). In the relevant pages (pp. 322–332), Rorty discusses the controversy between Cardinal Bellarmine and Galileo Galilei (henceforth 'B/G controversy') over the

M. Kusch (✉)
Department of Philosophy, University of Vienna, 1010 Vienna, Austria
e-mail: martin.kusch@univie.ac.at

'two chief world systems' as an instance of a clash between incommensurable Kuhnian paradigms. Rorty appeals to this case study in order to convince his readers of two claims. First, Bellarmine and Galileo used different 'grids' for determining 'what sorts of evidence there *could be* for statements about the movements of planets' (Rorty 1981, p. 330). And second, the conflict between the two grids was too deep for it to be decided by epistemic considerations acceptable at the time. For Rorty, the dispute therefore did not 'differ in kind' from the *political* controversy around 1917 between the liberal-socialist Alexander Kerensky and the Bolshevik Vladimir Lenin over such issues as Russia's involvement in World War I, or the need for a radical land reform (1981, p. 331). Boghossian takes it that Rorty is here defending epistemic relativism by drawing an epistemic-pluralist lesson from the B/G-disagreement.

Boghossian's talk of EP and epistemic systems consisting of epistemic principles is meant as a rational reconstruction of Rorty's position. And yet, ultimately Boghossian is not interested in defending Rorty but intent on criticizing both epistemic relativism in general and Rorty's case study argument in favour of EP in particular. Boghossian seeks to show that the rationally reconstructed version of Rorty's position is not supported by the B/G-dispute. Bellarmine and Galileo did not use 'fundamentally different, genuine alternative epistemic systems'. In developing this criticism, Boghossian is not alone. Markus Seidel (2014) seeks to improve on Boghossian's argument.

In this paper, I shall try to undermine Boghossian's and Seidel's criticism of Rorty's (alleged) use of the B/G-controversy. I write 'alleged use' since I do not think that Boghossian's rational reconstruction of Rorty's position is adequate. First, I shall suggest that *even if* we leave unchallenged Boghossian's reconstruction, it is far from clear that his and Seidel's criticism succeeds. Second, I shall argue that Boghossian's reconstruction makes the problematic—'isolationist'—assumption that epistemic systems can be clearly separated from non-epistemic systems of beliefs, principles or values. Boghossian's assumption is questionable both in the light of what we know about the B/G-controversy and in the light of how Rorty renders it. I shall propose replacing Boghossian's isolationism with a form of 'holism'. And third, and in a similar vein, I shall object to Boghossian's and Seidel's 'foundationalist' view of

epistemic systems, according to which such systems have a permanent and fixed structure of principles, ending with fundamental principles that are epistemically independent of all other principles. I shall suggest that a 'coherentist' view fits much better both with Rorty's wording and with the historical evidence concerning the Galileo trials.

I shall conclude—in a 'Postcript'—by offering some fragmentary observations on the relationship between the epistemic pluralism at issue in the main bulk of this paper and the epistemic pluralism associated primarily with the work of William Alston (1993).

2   Boghossian on Epistemic Systems and the Galileo-Bellarmine Controversy

Boghossian spends a whole chapter of *Fear of Knowledge* on developing a reconstruction and prima facie defence of epistemological relativism (2006, Chap. 5). The starting point is Rorty's idea of a multitude of Kuhnian 'paradigms' or Foucauldian 'grids' that are, in some sense, 'equally valid'. Rorty explains his view by reminding us of the B/G-dispute. Cardinal Roberto Bellarmine is often alleged to have believed that the Bible is a better source of evidence about the heavens than are telescopes. Rorty defends the cardinal against the charge of being 'illogical and unscientific'. According to *Philosophy and the Mirror of Nature*, Bellarmine had adopted a paradigm or grid of beliefs and principles that was fundamentally different from both Galileo's and our own. For instance, Bellarmine's grid did not allow for our principled distinction between science and religion. Rorty goes on to suggest that there is no absolute vantage point from which our grid can be judged to be superior. That we believe our grid to be more 'objective' or more 'rational' is nothing but an accident of history. From within Galileo's or our system, it is epistemically justified to believe in the Copernican theory; from within Bellarmine's epistemic system is it justified to stick to the Ptolemaic view (Rorty 1981, pp. 328–329; Boghossian 2006, p. 63).

Boghossian seeks to make Rorty's thought more precise. In a first step, Boghossian reconstructs in more detail what he takes to be the

constituents of epistemic systems, that is, epistemic principles. He distinguishes between 'generation' and 'transmission' principles on the one hand, and 'fundamental' and 'derived' principles on the other hand. Generation principles produce justified beliefs on the basis of something that is not itself a belief; transmission principles prescribe how to move from one justified belief to another. A fundamental principle is one 'whose correctness cannot be derived from the correctness of other epistemic principles' (2006, p. 67). This contrasts with derived principles. Here are examples of the four categories:

(*Observation*) [a fundamental generation principle] For any observational proposition *p*, if it visually seems to *S* that *p* and circumstantial conditions *D* obtain, then *S* is *prima facie* justified in believing *p*. (2006, p. 64)

(*Deduction*) [a fundamental transmission principle] If *S* is justified in believing *p* and *p* fairly obviously entails *q*, then *S* is justified in believing *q*. (2006, p. 66)

(*Observation-dog*) [a derived generation principle] If it visually appears to *S* that there is a dog in front of him, and circumstantial conditions *D* obtain, then *S* is *prima facie* justified in believing that there is a dog in front of him. (2006, p. 64)

(*Modus Ponens-rain*) [a derived transmission principle] If *S* justifiably believes that it will rain tomorrow, and justifiably believes that if it rains tomorrow the streets will be wet tomorrow, *S* is justified in believing that the streets will be wet tomorrow. (2006, p. 66)

Finally, Boghossian also proposes a formulation of Bellarmine's central principle:

(*Revelation*) For certain propositions *p*, including propositions about the heavens, believing *p* is *prima facie* justified if *p* is the revealed word of God as claimed by the Bible. (2006, p. 69)

Attributing this principle to Bellarmine might be supported the fact that he defended Ptolemy's system with passages from the *Bible* like 'The words "The sun rises and the sun sets, and hurries back to where it

rises, etc." were those of Solomon, who not only spoke by divine inspiration but was a man wise above all others …' (Bellarmine 1615).

Having introduced and clarified epistemic principles and systems, Boghossian proceeds to formulate more explicitly what he takes to be Rorty's argument in favour of epistemic relativism. This reconstruction is not of interest to us here. In this paper, I am only concerned with Boghossian's claim that it is wrong to think that the B/G-conflict is a case to which EP applies.

For EP to apply to the B/G-controversy, the two men's epistemic systems would have to be 'fundamentally different' (Or, since Boghossian assumes throughout that Galileo and 'we' have the same system, Bellarmine's system would have to be fundamentally different from our own.). In order for Bellarmine's system to qualify as fundamentally different from ours or Galileo's, Boghossian proposes that his system must contain *at least one fundamental* epistemic principle which we do not recognise. And here the obvious candidate is *Revelation*.

At this point in the dialectic, Boghossian switches from exposition and reconstruction to criticism. He argues against a fundamental difference between Galileo's (or our) and Bellarmine's epistemic systems. He does so by denying that *Revelation* is a fundamental principle even for Bellarmine. Boghossian reasons as follows. If *Revelation* were fundamental then it would trump *Observation* with respect to *some* statements about the heavens (e.g. Jupiter has moons) but not with respect to others (e.g. there are clouds in the sky). Here Boghossian rightly takes for granted that the cardinal uses his eyes for determining the degree of cloudiness. This is bad news for the relativist, Boghossian announces: the proposed division of labour between *Revelation* and *Observation*—*Revelation* for the stars, *Observation* for the clouds—is epistemologically unmotivated and therefore arbitrary. And thus the epistemic relativist faces a dilemma. First horn: If she treats *Revelation* as fundamental for Bellarmine, and as occasionally trumping *Observation* in an arbitrary fashion, then she renders his epistemic system incoherent and in consequence open to criticism. In other words, she then has every reason to give up her relativistic take on his system. Second horn: If the relativist instead seeks to remove the incoherence by 'downgrading' *Revelation* to a mere derived principle—a principle governed in its range of

application by fundamental principles—then she has lost the right to count Bellarmine's system as fundamentally different from Galileo's (or our own).

Boghossian himself opts for the second alternative by suggesting a way in which *Revelation* might have been derived from more fundamental principles:

> … we had better regard his [i.e. Bellarmine's] system as differing from ours only in some derived sense, attributing to him the view that there is evidence, of a perfectly ordinary sort, that the Holy Scripture is the revealed world of the Creator of the Universe. And it is only natural for someone with that belief to place a great deal of stock in what it has to say about the heavens …. (2006, pp. 104–105)

This is a little sketchy but presumably Boghossian is submitting that Bellarmine's religious beliefs in general, and his belief in the Bible in particular, are due (primarily) to testimony, observation and inference to the best explanation. All these are governed by fundamental epistemic principles, principles that are parts of both Bellarmine's and our secular epistemic systems. Once Bellarmine's religious belief, and his belief in the Bible as the Word of God, are in place, he indeed has reason to also accept *Revelation* as a further principle. And yet, without the more fundamental principles, *Revelation* could not have been motivated.

Finally, since Bellarmine's system differs from Galileo's and our own only slightly, and only with respect to one derived principle, there can be a *rational* debate over the justifiability of *Revelation* between him, Galileo and us. The question is simply whether there is 'evidence of a perfectly ordinary sort for believing that what was written down in some book by a large number of people over a vast period of time, internal inconsistencies and all, is really the revealed word of the Creator' (2006, p. 105).

3   Seidel's Further Development of Boghossian's Argument

Seidel (2014) follows Boghossian's lead but tries to make his ally's main arguments more precise. He does so by formulating what he takes to be two 'intuitive' criteria for concluding that given epistemic principles are not fundamental.

The first such criterion Seidel calls 'Instance'. Assume we have one epistemic system, ES_1, with norm N' and another epistemic system, ES_2, with norm N". Allow further that N' and N" are instances of a further norm N that is part of both ES_1 and ES_2. In such a situation, Seidel thinks we would all find it intuitive that ES_1 and ES_2 are not fundamentally different epistemic systems, at least not in virtue of their differing with respect to N' and N". Seidel explains and justifies the principle with the following example: ES_1 is the epistemic system of Platonism; ES_2 is the epistemic system of Aristotelianism; N' is 'If Plato says p, then I am prima facie justified in believing that p'; and N" is 'If Aristotle says p, then I am prima facie justified in believing that p'. N' and N" are instances of N: 'If an ancient philosopher says p, then I am prima facie justified in believing that p'. Belief B', occurring in ES_1 but not in ES_2, is: 'Plato is an ancient philosopher' and B", occurring in ES_2 but not in ES_1, is: 'Aristotle is an ancient philosopher'. Seidel's verdict: however different the Platonists' and the Aristotelians' beliefs or derived norms may be, 'we would not say' that they have 'fundamentally different epistemic systems' (Seidel 2014, p. 169).

Seidel's second principle is called 'Derive'. Assume we have again ES_1 with norm N' and ES_2 with norm N". This time N' and N" can both be derived from a further norm N that is part of both ES_1 and ES_2. Here too Seidel is confident that we would deem it intuitive that ES_1 and ES_2 are not fundamentally different epistemic systems, at least not in virtue of their differing with respect to N' and N". Seidel uses the same example as before except that N is now: 'If an epistemologist says that p, then I am prima facie justified in believing that p'. B' is: 'An epistemologist told me that: "If Plato says p, then I am prima facie justified in believing that p"'. And B" mutatis mutandis for the Aristotle.

Seidel maintains that we have here no 'fundamentally different epistemic systems'. We have different beliefs, different derived norms, but the same fundamental norms.

Following in Boghossian's footsteps, Seidel applies his principles to Bellarmine's epistemic system. To make his case, Seidel contrasts Bellarmine's *Revelation* with a principle he calls *Science:*

> (Science) For certain propositions p, including propositions about the heavens, believing p is prima facie justified if p is included in the best physics books available. (Seidel 2014, p. 175)

How does *Revelation* relate to *Science*? Seidel argues that both are instances of, or derived from, a more fundamental principle we might call '*Reliability of Books*':

> (Reliability of Books) For certain propositions p, including propositions about the heavens, believing p is prima facie justified if p is included in books who have been assessed as highly reliable by appropriate experts.

The upshot is of course that since *Science* and *Revelation* are both instances of, or derived from *Relialibity of Books*, they are not to be taken as fundamental. And hence the relativist case for EP collapses (2014, p. 177).

4   Mystical Perception

I have more than one worry concerning these arguments. Some relate to Boghossian's and Seidel's interpretation of what was at issue between Bellarmine and Galileo, others focus on their static and crystalline conception of epistemic systems. But I begin with an objection that grants the two authors this (problematic) conception.

Consider an epistemic principle I propose calling 'Mystical Perception':

(Mystical Perception): If it seems to S that God is telling him that p; and if S is not already fully committed to atheism; and if circumstantial conditions D obtain, then S is prima facie justified in believing that God is telling him that p.

Mystical Perception is not part of my epistemic system but it is a fundamental principle in the epistemic systems of others. And at least some of the others cannot be easily dismissed as fools or religious fanatics. After all, the most detailed defence of the principle of mystical perception comes from the pen of the distinguished epistemologist William Alston who wrote almost four hundred pages on this topic Alston (1991). Amongst other things, Alston argues in great detail that mystical perception has parallels with sensory perception in that neither have noncircular demonstrations of their reliability: both are self-supporting; both have over-rider systems; both are sometimes inconsistent; and both cohere with other epistemic practices. To my mind, the argument that there is no noncircular demonstration of the reliability of mystical perception makes a good case for treating it as a fundamental principle—in Alston's epistemic system.

Seidel disagrees (in response to Kusch (draft)). As Seidel has it, mystical perception and sensory perception are both *instances of perception*; hence the principle of Instance applies, and rules out the option of treating Mystical Perception as fundamental (2014, p. 167). I am not convinced. It is true of course that in some sense mystical perception has always been modelled on sensory perception. That is after all why we call mystical perception 'mystical *perception*'. But it is not obvious to me that we should take our epistemic guidance from such vague analogies. It also is not clear to me how we should think of perception once we have abstracted from both the 'mystical' and the 'sensory'. In any case, I cannot see why these considerations should be weightier than Alston's argument to the effect that mystical perception has no noncircular demonstration of its reliability.

Can Seidel's argument be improved? Rather than saying that *Mystical Perception* and *Sensory Perception* are instances of *Perception* why not say that *Mystical Perception* and *Sensory Perception* are instances of a principle called '*Seeming*':

(Seeming) If it seems to S that p, and circumstantial conditions D obtain, then S is prima facie justified in believing that p.

It obviously is right to say that *Mystical Perception* and *Sensory Perception* (as well as some other principles) can be construed as instances of *Seeming*. But I am not convinced that this fact tells against the possibility of fundamentally different epistemic systems. The problem is that if the principles common to different epistemic systems become too abstract, too thin, then it is no longer plausible to assume that the common principles prevent the respective epistemic systems from being, intuitively, fundamentally different. At least if we mean by 'intuitively, fundamentally different' something like the idea that to switch from one system to the other would feel like epistemic-cum-metaphysical trauma, dislocation or revolution in thought.

Does *Seeming* rule out epistemic relativism? One might think so on the following grounds. Epistemic principles differ in what they regard as appropriate conditions for a seeming to confer justification. These differences trace back to factual beliefs about when seemings track the truth. These beliefs can be tested. Moreover, if two incompatible principles (belonging to two different epistemic systems) involve contradictory beliefs about which seemings are truth-tracking, what sense can be made of the relativist's claim that the two principles could be equally valid?

To answer this worry the epistemic relativist needs to insist again that the testing of factual beliefs does never happen in isolation but only against the background of specific epistemic systems. Does *Revelation* enable us to track the truth? That depends on what we mean by 'truth' and what we mean by 'tracking'. Moreover, remember that Alston argues that neither mystical nor sensory perception has noncircular demonstrations of their reliability. If Alston is right, then the fact that both are instances of seeming does not show that they can be tested and compared in a neutral way.

5   Rethinking the Role of *Revelation* the Galileo-Bellarmine Controversy

Up to this point, I have challenged Boghossian's and Seidel's arguments without scrutinizing their readings of Rorty or their rendering of the dispute between Bellarmine and Galileo. I now turn to that latter task.

Open any standard historical account of the episode, and invariably you will find something like the following observation: 'This was a controversy involving issues of methodology, epistemology, and theology as well as astronomy, physics, and cosmology' (Finocchiaro 2005: 1). Authors influenced by the sociology of knowledge add period-specific relationships between patrons and courtiers; traditions of instrument-making; the tensions between different religious orders; the politics of the papal court; the political problems between Spain and the Vatican; the Thirty-Year War; and much else besides (Biagioli 1993, 2006).

In reducing the episode to a clash over the epistemic status of *one* epistemic principle, *Revelation*, Boghossian strips away pretty much all of this complexity. In his reconstruction, all non-epistemic considerations are set aside as irrelevant. In the process, religion is turned into an epistemic system; epistemic principles are treated as isolated or isolatable entities with fixed implications.

Boghossian's sketchy and quick reconstruction has its costs. One casualty is the interesting fact that Galileo and Bellarmine did not disagree over *Revelation* as such—only over its scope. As we saw above, as Boghossian has it, Bellarmine's belief in geocentrism resulted from an application of *Revelation*. And since Galileo rejected *Revelation* he was free to believe in Copernicanism on the basis of empirical evidence. This interpretation of the episode contradicts the best recent scholarship (Boghossian's only reference is Giorgio de Santillana's *The Crime of Galileo* (1955)[1]).

To cut a long story short, both Bellarmine and Galileo accepted the following tripartite distinction between propositions about the natural world:

(i)   propositions about the natural world that have been *demonstrated* (by our natural lights, that is, by reason);

(ii)   propositions about the natural world that in principle are *demonstrable*, but that have not yet been demonstrated; and

(iii) propositions about the natural world that are *beyond demonstration* (Blackwell 1991: 3328).

Consider propositions in these three categories that, *given a literal reading*, seem to contradict the Bible. Bellarmine and Galileo agreed on what to do about cases in (i): in such cases, the Biblical passages were to be reinterpreted figuratively in such a way that they would come out true, and that they would not contradict the demonstrated proposition. This procedure was motivated by a belief, shared by Bellarmine and Galileo, that the Bible—a text allegedly dictated by the Holy Spirit—spoke the truth and nothing but the truth. Galileo and Bellarmine also agreed on how to react to category (iii): when such propositions contradicted the Bible, then they had to be rejected as false and heretical. In this case, revelation invariably trumped philosophical speculation.

To repeat, Bellarmine and Galileo both rejected *Revelation* for category (i), and they both accepted *Revelation* for category (iii). The point of contention was what to do about category (ii). Here Bellarmine and the Church insisted that, when such propositions contradicted a literal reading of the Bible, then they had to be considered false and heretical. This did not, however, preclude using these propositions (and the theories to which they belonged) as useful fictions. But no *realist* commitments to these propositions were acceptable. Galileo differed. He urged the Church not to treat such propositions as false and heretical. This proposal was of course inseparable from Galileo's belief that Copernicanism was an instance of category (ii). In other words, Galileo accepted that the truth of heliocentrism had not yet been demonstrated. Nevertheless, he asked that the Bible—read literally—should not be the yardstick for judging Copernicanism. Instead, the Church ought to suspend judgement until a demonstration for either Copernicanism or the Ptolemaic system had been found. For Bellarmine (and later for Pope Urban VIII), this position was unacceptable. But this was not only because they rejected Galileo's rendering of category (ii) as problematic;

the deeper reason was that they likely considered the question of the correct 'world system' to belong to category (iii).

Why is all this important? Boghossian sees Bellarmine's and Galileo's judgements about geocentrism as guided by epistemic rules like *Revelation* (in Bellarmine's case) or *Observation* plus *Inference to the Best Explanation* (in Galileo's case). Someone who had adopted *Revelation* had to end up opposing Copernicanism; someone who had instead adopted *Observation* and *Inference to the Best Explanation* could opt for Copernicanism. My brief excursion into the historical scholarship shows that this reading of the event is mistaken. It was not the adoption or rejection of *Revelation* that made the difference—both Bellarmine and Galileo accepted it. But they differed in how they interpreted this rule. Their interpretations differed with respect to the categories (i), (ii) and (iii); they differed concerning the proper position of Copernicanism as falling into either (ii) or (iii). Historians have shown us in great detail the great variety of considerations that influenced both streams of judgements: the context of the Counter-Reformation; relationships between patrons and courtiers; traditions of instrument-making; the tensions between different religious orders; the politics of the papal court; the political problems between Spain and the Vatican; the Thirty-Year War; and much else besides (Biagioli 1993, 2006). Interpretative decisions made in the light of such complex and intricate considerations cannot be reduced to a simple—or even a complicated—rule.

6   The Areas of Agreement and Disagreement Between Galileo and Bellarmine

We can get at the same complexity also via a different route, that is, by dissecting more carefully the areas of agreement and disagreement of Bellarmine's and Galileo's views. Again I am taking my lead from historical scholarship of the episode (Blackwell 1991; McMullin 2005). I do not here have the space to cite the textual evidence these authors put forwards in support of their historical claims. But it is readily

available in these two sources in particular. There was substantial agreement between Galilei and Bellarmine over many important scientific and theological issues:

  (I)   The Bible is the Word of God.
  (II)   The Bible cannot be in conflict with natural philosophy.
  (III)   No 'demonstration' (in natural philosophy) can trump the Bible in 'matters of faith or morals'.
  (IV)   In cases other than (III), when a demonstration contradicts the literal meaning of a Biblical passage, the latter must be reinterpreted figuratively.
  (V)   Concerning issues where no demonstration is possible, and where Bible passages apply, the latter are to be believed.
  (VI)   One should not prematurely commit the Church to interpretations of difficult Biblical passages, lest these interpretations later conflict with demonstrations.
  (VII)   At least sometimes it is appropriate to support cosmological views with Biblical passages.
  (VIII)   Copernican theory does an excellent job of 'saving the phenomena'.
  (IX)   To date, no 'demonstration' of Copernicanism has been offered.
  (X)   *Revelation*.

Bellarmine's and Galileo's agreement on (X) *Revelation* is clear in the light of (III), (V) and (VII). Turning from agreement to disagreement, the following propositions were all believed by Bellarmine but denied by Galileo:

  (i)   Copernicanism contradicts common sense.
  (ii)   The natural-philosophical case for Copernicanism is weak.
  (iii)   Mathematical astronomy (Copernicus, Galileo) cannot decide issues in physical astronomy (Aristotle).
  (iv)   The case for Aristotelian scholastic philosophy is strong.
  (v)   The consensus of the 'Church Fathers' on Biblical literalism is important.

(vi)   'Matters of faith or morals' include the belief that God is the truth-telling author of the Bible.

(vii) In assessing, the Bible's cosmological claims, we need not give much weight to the fact that its addressees include the uneducated.

(viii) The same is true concerning the Bible's primary concern with human salvation.

(ix)   When a demonstrable, but as yet undemonstrated, belief contradicts a literal reading of a Biblical passage, it is right to stick to the latter.

(x)   According to the Bible, Solomon was the wisest of men. Thus, his beliefs cannot have been contrary to what is demonstrable. Solomon thought that the Sun moves around the Earth. Solomon's belief is crucial cosmological testimony.

7   Isolationism Vs. Holism; Foundationalism Vs. Coherentism

Assume the historians of the episode are roughly on the right track regarding the above. What follows for the concerns of this paper?

The first thing to note is that Boghossian and Seidel are wrong to present the clash between Galileo and Bellarmine as one between a secular scientist (like 'us') and an unscientific religious believer. The conflict was one between two religious believers, both of whom had considerable knowledge of contemporary 'natural philosophy'—that is, what we today would classify as 'natural science' *and* what we today would call 'epistemology and metaphysics of natural science'.

Second, although we *could* rephrase the listed propositions as *rules*, this would be a bit artificial. At least it is not clear which rendering— the doxastic-descriptive or the normative—was more fundamental for the historical actors. This suggests that it is at best one-sided to insist that the 'systems' of natural philosophy consisted only of norms or principles rather than beliefs or that the more fundamental contents

of an epistemic system are norms. This position is not defended by Boghossian and Seidel in any way; it is simply assumed to be correct.

Third, in Bellarmine's and Galileo's respective 'systems of beliefs and principles' ('systems$_{b\&p}$' from now on), epistemological issues were tightly woven together with concerns in theology, natural philosophy, logic, common sense, metaphysics and epistemology. Take (IX) in an epistemic-normative format:

> Do not cease believing a literal reading of a Biblical passage just because it is contradicted by a demonstrable but as yet undemonstrated belief.

To understand and apply this principle, the epistemic subject needed to appreciate: what counts as a *literal reading* of a Biblical passage; this was far from straightforward since there existed competing theological schools of Biblical hermeneutics, each with its own criteria of the literal); what was meant by a 'demonstration' (this was of course a key term of Aristotelian epistemology and metaphysics, variously interpreted by different Church Fathers and philosophers); or how to determine that a demonstration was possible even though it had not yet been carried out (this question was tied to different views on the metaphysics of modalities and theological premises).

We can of course always insist on 'filtering out' the epistemological from the metaphysical, the theological or the ethical. With enough patience and ingenuity, we can often come up with a rational reconstruction or idealization of a given system$_{b\&p}$ such that only the epistemological side of a given system$_{b\&p}$ is salient. But we should be clear that to analyze a rational reconstruction of Bellarmine's system$_{b\&p}$ is not to analyze Bellarmine's system$_{b\&p}$. Conclusions drawn about the former are not automatically adequate to the latter. In other words, a rational reconstruction imposes our criteria of rationality on a system$_{b\&p}$ that—absent our reconstruction—may well encode a different form of rationality.

This is the point where it seems useful to introduce a distinction that captures where I differ from Boghossian and Seidel. I am referring to the distinction between 'isolationism' and 'holism'. The isolationist concerning systems$_{b\&p}$ deems it possible and fruitful to filter

out, with respect to such systems, the epistemic from all other dimensions. The holist finds this distorting and unhelpful, even as an idealization. I submit that the historical material adduced above makes at least a *prima-facie* case for holism. After all, it shows that both Bellarmine and Galileo thought and reasoned in ways that resist a separating out of epistemic, metaphysical, natural-philosophical and theological concepts and principles.

Recall also that Boghossian's starting point is Rorty's discussion of the disagreement between Bellarmine and Galileo. Boghossian reconstructs Rorty as an isolationist. But it seems that a holist reading would fit much better with Rorty's train of thought. Rorty's main point is to insist that it is anachronistic to evaluate the disagreement between the two sixteenth/seventeenth-centuries figures in terms of what we today call 'scientific' criteria. The sharp distinction between Bellarmine's (allegedly illegitimate) theological considerations and Galileo's (allegedly legitimate) scientific reasons, is, Rorty insists, a product of the 'grid' that emerged in the 'later seventeenth and eighteenth centuries' and this grid 'was not there to be appealed to in the early seventeenth century, at the time that Galileo was on trial' (1981, p. 330). For thinkers in the early-seventeenth century there was no such distinction; what we wish to separate for them formed a unity.

Note also that Rorty uses the Foucauldian concept of 'grid' instead of 'epistemic system'. 'Grid' is variously used as a translation for Foucault's concepts 'épistème', 'dispositif' or 'grille d'intelligibilité'. In each case, Foucault is at pains to pick out a (by our, later eyes) *heterogeneous* sets of factors, ranging from the abstractly theoretical through architectural arrangements to forms of power (Dreyfus and Rabinow 1983; Gutting 1989; Kusch 1991). Let us also remember that Rorty's whole discussion occurs in a section entitled 'Kuhn and Incommensurability'. And scientific paradigms qua 'disciplinary matrices' are constituted by epistemological, metaphysical, mathematical and scientific elements.

Fourth, my insistence above that both Bellarmine and Galileo accepted *Revelation* should not be read as support for Boghossian and Seidel's contention according to which the B/G-dispute took place on the basis of 'a common epistemic system'. Even in the absence of clear identity conditions for 'epistemic systems' or 'system$_{b\&p}$', it seems forced

to suggest that two people can disagree over (i)–(x) above and yet do so on the basis of one and the same system. At least it is unclear what purpose talk of systems can serve, if it does not roughly track differences between the kinds of bodies of belief as different as Galileo's and Bellarmine's.

The assessment that Bellarmine and Galileo had *different* system$_{b\&p}$ can be supported by the observation that it is hard to imagine either men mounting a rational argument for their respective view that would have convinced their opponent. This, in any case, is the assessment of one of the leading contemporary historians and philosophers of Galileo's trial of 1616 (McMullin 2005). Bellarmine and Galileo each had a coherent system of beliefs and principles, and neither could rationally force the other to change their assessment of Copernicanism. Moreover, each side could comfortably account for what the other side counted as evidence. And there was no neutral vantage point from which one could decide which of the two webs of beliefs was right. Rorty agrees: this is why he suggests that it needed 'three hundred years of rhetoric' to convince us that Galileo was right (1981, p. 330).

Fifth, and finally, Boghossian's and Seidel's rendering is not only 'isolationist' but also 'foundationalist'. They insist on a strict distinction between fundamental or underived, and non-fundamental or derived epistemic principles. *Observation* or *Inference to the Best Explanation* is fundamental, *Revelation* is not. And *Revelation* is not fundamental since other epistemic principles must have played a role in its rational adoption.

My central misgiving in this. As epistemologists using the history of science, we must be clear about what we are trying to do. Are we trying to identify the systems of beliefs and norms of our actors—as these actors understood these systems—or are we attempting to construct, in our terms, a maximally abstract and parsimonious system of principles from which we can derive, or in terms of which we can justify, our actors' normative judgements? Note that these two ways of proceeding come with two different ways of understanding 'fundamentality' or 'derivation': if we go by the *actors' categories*, a principle is fundamental if the actors do not regard it as derived; if we go by our *analysts' categories*, a principle is fundamental if we—within our constructed parsimonious

system—would not derive it. It seems to me obvious from what I have already quoted from Rorty, that he is opting for the actors' categories. This is after all the whole point of his warning of illicit attempts to condemn Bellarmine by standards that were 'not there to be appealed to in the early seventeenth century'.

And yet, if we go with actors' categories then it is just not clear how we should divide up the principles involved into the more or less fundamental. Bellarmine accepted (I)–(X) and (i)–(x). Galileo endorses (I)–(X) and the negations of (i)–(x). But it is difficult if not impossible to decide which of these various beliefs were derived from which others. It rather seems that the respective beliefs formed two sets, webs or systems, such that Galileo regarded the one, and Bellarmine the other, as coherent. And depending on the specific challenge they faced from specific opponents at different times, the cardinal and the courtier would take different paths through their respective webs. If this is roughly on the right track, then it seems that we should replace foundationalism with coherentism. And this is of course fully in line with Rorty's own insistence on giving up an 'epistemologically centered "foundational" philosophy' (1981, p. 329).[2]

8   Conclusions

In this paper, I have argued for the following propositions. Even if Boghossian's and Seidel's fundamentalist and isolationist rendering of epistemic systems were right, their criticism of EP in the light of the B/G-controversy would not be decisive. I offered the epistemic principle of *Mystical Perception* (as studied and justified by William Alston (1991)) as an example of a fundamental principle that might occur in one epistemic system but not in another. More principally, the central intuition pump or case study behind Boghossian's and Seidel's reasoning—Rorty's interpretation of the B/G-dispute—does not justify the foundationalism and isolationism that characterizes Boghossian's and Seidel's theorizing about epistemic systems and epistemic relativism. Moreover, this observation seems apt both as far as Rorty's interpretation of the B/G-dispute is concerned, and as far as the currently best

historical scholarship of the episode is in question. Rorty and the best historical scholarship suggest a coherentist and holist rendering.

If I am right about all this—admittedly a big 'if'—what follows for epistemic relativism in general and epistemic pluralism in particular? Needless to say, very little 'follows' in any strict sense of 'following'. And yet, perhaps my argument at least *suggests* the following hypotheses as worthy of further investigation. First, of all, the relativist position seems able to survive Boghossian's and Seidel's specific onslaughts. Furthermore, formulating the relativist position in terms of holism and coherentism is more promising than rendering it in terms of isolationism and foundationalism. It is, in any case, noteworthy that card-carrying relativists like Rorty or the advocates of the 'Strong Programme' in the sociology of scientific knowledge (Barnes, Bloor, Henry 1996) have favoured the holist-coherentist formulation.

Shifting to a holist-coherentist version of relativism demands of course a rethinking of EP:

> There are many fundamentally different, genuinely alternative epistemic systems, but no facts by virtue of which one of these systems is more correct than any of the others. (Boghossian 2006, p. 73)

As it stands, this formulation—as well as the related criterion for 'fundamental difference'—is tied too closely to Boghossian's foundationalism and isolationism. We do well to replace it with a principle that befits holism and coherentism. The B/G-dispute, e.g., might be taken to motivate a principle of 'scientific pluralism'. It is broader in scope than EP since it does not seek to filter out the epistemic domain. But it is also more specific in that it homes in on one important realm, namely science (and its predecessor, 'natural philosophy'):

> (Scientific pluralism) There are some fundamentally different, genuine alternative systems$_{b\&p}$ in the sciences. The fundamental difference between two systems S1 and S2 is not determined by fundamental principles appearing in S1 and not in S2 but by the difficulty of imaging a rational proponent of S1 convincing a rational proponent of S2 to switch her allegiance from S2 to S1. We have a fundamental difference between

two systems when a switch would feel like an epistemic-cum-metaphysical trauma, dislocation or revolution in thought. (cf. van Fraassen 2002)

9   Postscript: Two Forms of Epistemic Pluralism

Contemporary philosophers use the label 'epistemic pluralism' for two rather different views. One is at issue in the paper above. The other is addressed by most authors in this collection and was first sketched in William P. Alston's little classic, 'Epistemic Desiderata' (1993). Alston argues against the hope of finding the one and only correct account of epistemic justification. Instead, he urges epistemologists to focus their attention on a variety of different 'epistemic desiderata', such as 'coherence', 'reliability' or 'cognitive accessibility'. It is natural to go one step beyond Alston by suggesting that different epistemic desiderata apply in different areas of discourse (Similar steps have been suggested in the related but different realm of truth pluralism (cf., e.g. Wright 1992)).

How do Alston's pluralism and Boghossian's pluralism relate to one another? Perhaps surprisingly this question has not to date been systematically discussed. One initially tempting idea for distinguishing between the two forms of pluralism is to point to Alston's rejection of what he calls 'latitudinarianism' concerning epistemic desiderata. The latitudinarian is happy to let a thousand epistemic desiderata bloom, without scrutiny and evaluation. Alston's position is different. He holds that some of the conditions previously proposed as epistemically valuable 'may be eliminated as unattainable (or not sufficiently attainable), while others may be plainly more important than others' (1993, p. 543). In other words, Alston is very much concerned with the evaluation of theories concerning epistemic desiderata. And this invites the thought that latitudinarianism is really just another word for relativism, and that in rejecting the former, Alston has also rejected the latter.

On closer inspection, however, it turns out that the 'tempting idea' works only as long as we commit the epistemic relativist to abstain altogether from the epistemic evaluation of epistemic systems or desiderata.

And yet, it is not obvious that the relativist needs to apply her 'equal validity' thesis so widely. Boghossian's pluralism says that there are 'many fundamentally different, genuinely alternative epistemic systems'—but that still leaves the option of rejecting some candidate alternatives as epistemically problematic in some way.

A different suggestion for distinguishing between the two forms of epistemic pluralism (a suggestion first put to me by Annalisa Coliva) builds on the notion that different epistemic desiderata apply in different domains.[3] This might make the different epistemic desiderata *compatible* with one another. The relativist pluralism studied by Boghossian is different. It is the view that there is a plurality of epistemic systems or desiderata that are *not compatible* and yet in some sense equally valid. This suggestion seems to me to be on the right track, but needs a bit more development. Three points seem especially important to me.

First, relativism is typically invoked when we face (what appear to be) *irresolvable disagreements*. Boghossian's book is a case in point. On Boghossian's rendering, Galileo and Bellarmine are disagreeing over which epistemic principles to apply in the study of the heavens. They are disagreeing over issues in one and the same domain. And the relativist (in Boghossian's reconstruction) seeks to explain the irresolvable character of the disagreement by attributing different 'fundamental' epistemic principles (and thereby fundamentally different epistemic systems) to the two men. Clearly, if we construe Alstonian pluralism as allowing for different epistemic desiderata in different domains (only), then the disagreement between Galileo and Bellarmine is not a case on which Alston's pluralism can be fruitfully brought to bear.

Second, it is nevertheless illuminating to relate Alston's distinction between different epistemic desiderata to relativistic pluralism as reconstructed by Boghossian. This is the point where we might separate Alston's original proposal from the further idea according to which different epistemic desiderata apply only in <u>different</u> domains. For Bellarmine 'accordance with scripture (literally understood)' was an epistemic desideratum for claims about the natural world, including the heavens; for Galileo, it was not. This is of course just to offer a translation of 'epistemic-principle' talk into 'epistemic-desiderata' lingo. But the fact that such translation is easy and natural, suggests that there is

an overlap of concerns between the epistemologists studying Alstonian pluralism and the epistemologists concerned with relativistic pluralism: both have an interest in discovering just how many plausible epistemic goods (desiderata, principles) there are, and how they relate to one another.

Third, note also that Alstonian and relativistic pluralisms might have a common motivation, that is, to account for irresolvable disagreements. Thus the Alstonian might explain the gulf between internalists and externalists by pointing to their incompatible commitments to different epistemic desiderata. This is not the way Alston himself uses his pluralism. Alston is more concerned to explain away the disagreement between internalism and externalism: contrary to their initial self-conception, the two sides do not disagree over what one and the same concept of epistemic justification requires from us; they simply invoke different (and compatible) such concepts of epistemic justification.

Fourth, epistemic relativism need not commit to the idea that fundamentally different epistemic systems are committed to incompatible epistemic desiderata. At least this is so if we specify epistemic desiderata at the level of generality that Alston's original paper worked with. Epistemic relativism might also apply in cases where two epistemic communities (with their respective epistemic systems of principles) favour the same epistemic desiderata *but apply them in different ways*. Irresolvable epistemic disagreement can exist between two communities that both give great value to reliability as the central epistemic desideratum in a given domain. And yet, one community might favour reliability of predictions that allow us to minimize false positives, while another community might favour reliability of predictions that enable us to minimize false negatives.

These are just some preliminary ideas. The matter clearly deserves an extended discussion elsewhere.[4]

Notes

1. The following studies can be regarded as constituting or reflecting (in good part) the state of the art on Galileo and his conflict with the Catholic Church: Biagioli (1993, 2006), Blackwell (1991), Finocchiaro (1980, 2005), Drake (1978), Heilbron (2010), Koyre (1978), Machamer (ed.) (1998), McMullin (2005), Redondi (1987), Renn (2002), and Wallace (1984).

2. Some provisos are worth noting. First, although both Boghossian and Seidel seem to commit to a foundationalism of epistemic principles, Seidel's two criteria of 'Instance' and 'Derive' are not conclusive evidence for foundationalism *on their own*. It is Seidel's general agreement with the way Boghossian sets up epistemic systems that is conclusive. Second, it is not clear whether Boghossian's and Seidel's analysis of the B/G-controversy necessarily requires a foundationalist (rather than a coherentist) conception of epistemic systems. After all, even the coherentist can allow that some epistemic principles are *more fundamental*, closer to the centre of the web, than others. Even on a coherentist rendering of epistemic systems Boghossian and Seidel could insist that *Revelation* is *less fundamental* than *Observation*. Note though that this would force Boghossian and Seidel to change their definition of what makes two epistemic systems fundamentally different. And the question remains why we should accept the claim with that respect to each and every dispute over epistemic principles we are able to find more fundamental principles in terms of which the dispute can be rationally resolved.

3. I tentatively go along with this proposal here, though I am sceptical whether it can ultimately be made to work. Can we neatly assign different desiderata—like coherence or reliability—to different domains? I doubt it. But I grant that the value of the desiderata might differ from domain to domain.

4. Sections 3 and 4 are drawn from Kusch (2016); Sect. 5 is drawn from Kinzel and Kusch (submitted). For comments and suggestions, I am grateful to Annalisa Coliva, Natalie Ashton and Robin McKenna. Work on this paper was made possible by ERC Advanced Grant 'The Emergence of Relativism' (#339382).

References

Alston, W.P. 1991. *Perceiving God: The Epistemology of Religious Experience*. Ithaca and London: Cornell University Press.

Alston, W.P. 1993. Epistemic Desiderata. *Philosophy and Phenomenological Research* LIII: 527–551.

Barnes, B., D. Bloor, and J. Henry. 1996. *Scientific Knowledge: A Sociological Analysis*. London: Athlone.

Bellarmine. 1615. Letter to Foscarini. http://www.historyguide.org/earlymod/foscarini.html (accessed 3rd March 2017).

Biagioli, M. 2006. *Galileo's Instruments of Credit: Telescopes, Images, Secrecy*. Chicago: University of Chicago Press.

Biogioli, M. 1993. *Galileo, Courtier*. Chicago: University of Chicago Press.

Blackwell, R.J. 1991. *Galileo, Bellarmine, and the Bible. Notre Dame: University of Notre Dame Press.* (Kindle edition).

Boghossian, P. 2006. *Fear of Knowledge: Against Relativism and Constructivism*. Oxford: Oxford University Press.

de Santillana, G. 1955. *The Crime of Galileo*. Chicago: University of Chicago Press.

Drake, St. 1978. *Galileo at Work*. Chicago: University of Chicago Press.

Dreyfus, H., and P. Rabinow. 1983. *Beyond Structuralism and Hermeneutics*, 2nd ed. Brighton: Harvester Press.

Finocchiaro, M.A. 1980. *Galileo and the Art of Reasoning*. Dordrecht: Reidel.

Finocchiaro, M.A. 2005. *Retrying Galileo, 1633–1992*. Berkeley: University of California Press.

Gutting, G. 1989. *Michel Foucault's Archaeology of Scientific Reason*. Cambridge: Cambridge University Press.

Heilbron, J.L. 2010. *Galileo*. Oxford: Oxford University Press.

Koyré, A. 1978. *Metaphysics and Measurement*. London: Chapman & Hall.

Kinzel, K., and M. Kusch (submitted). De-idealizing Disagreement, Rethinking Relativism.

Kusch, M. 1991. *Foucault's Strata and Fields: An Investigation into Archaeological and Genealogical Science Studies*. Dordrecht: Kluwer.

———. 2016. Epistemic Relativism, Scepticism, Pluralism. *Synthese*. doi:10.1007/s11229-016-1041-0.

Machamer, P. (ed.). 1998. *The Cambridge Companion to Galileo*. Cambridge: Cambridge University Press.

McMullin, E. (ed.). 2005. *The Church and Galileo*. Notre Dame: Notre Dame University Press.

Redondi, R. 1987. *Galileo Heretic*. Princeton: Princeton University Press.

Renn, J. (ed.). 2002. *Galileo in Context*. Cambridge: Cambridge University Press.

Rorty, R. 1981. *Philosophy and the Mirror of Nature*. Oxford: Blackwell.

Seidel, M. 2014. *Epistemic Relativism: A Constructive Critique*. London: Palgrave Macmillan.

van Fraassen, B.C. 2002. *The Empirical Stance*. New Haven and London: Yale University Press.

Wallace, W.A. 1984. *Galileo and His Sources: The Heritage of the Collegio Romano in Galileo's Science*. Princeton: Princeton University Press.

Wright, C. 1992. *Truth and Objectivity*. Cambridge: Harvard University Press.

Epistemic Pluralism, Epistemic Relativism and 'Hinge' Epistemology

J. Adam Carter

1 Introduction

According to Paul Boghossian (2006, p. 73), a core tenet of epistemic relativism is what he calls *epistemic pluralism*, the thesis that (i) 'there are many fundamentally different, genuinely alternative epistemic systems', but (ii) 'no facts by virtue of which one of these systems is more correct than any of the others'.[1] Meta-epistemological absolutists often take for granted that the less controversial (i) does not entail the more controversial (ii), and with this assumption in hand, insist that (ii) is false even if the descriptive claim captured by (i) is true.[2]

Interestingly, though, this stock absolutist response to the pluralist component of epistemic relativism is plausibly not available to the Wittgensteinian 'hinge' epistemologist,[3] for whom the more controversial element of epistemic pluralism (ii), is (arguably) inevitable once it

J. Adam Carter (✉)
Philosophy, School of Humanities, University of Glasgow, G12 8QQ
Glasgow, UK
e-mail: adam.carter@glasgow.ac.uk

229

is affirmed that there are many, or indeed *any*, fundamentally different, genuinely alternative epistemic systems. This raises an important dilemma for the would-be Wittgenstenian epistemologist: either grant that a hinge epistemology is essentially a relativistic epistemology (one which embraces both legs of the epistemic pluralism thesis), or alternatively deny at least some version of the widely accepted descriptive claim that there are many fundamentally different, genuinely alternative epistemic systems. In the face of this dilemma, Duncan Pritchard (2009, 2010, 2015) has recently defended a non-relativistic reading of Wittgenstein's epistemology, and accordingly, has sought to reject, with some qualifications, the descriptive component of the epistemic pluralist thesis.

Here is the plan. §1 distinguishes between different versions of descriptive epistemic pluralism, clarifying which version (paired with a Wittgenstenian epistemology) arguably commits one to the stronger 'equal standing' leg of epistemic pluralism. §§2–3 outline and critique Pritchard's anti-relativistic response to relativist interpretations of Wittgenstein's epistemology, with a focus on Pritchard's Davidsonian and '*über hinge*' strategies—viz. strategies that can be used to resist with some qualifications the version of descriptive leg of the epistemic pluralist thesis that would seem to commit the hinge epistemologist to epistemic incommensurability. §4 suggests how, regardless of whether Pritchard's strategy succeeds, the Wittgenstenian hinge epistemologist might, nonetheless, have available a meta-epistemological rationale for accepting (within a Wittgenstenian framework) the very version of descriptive pluralism which Pritchard's brand of hinge epistemology rejects, whilst nonetheless avoiding any further commitment to epistemic relativism, understood along certain specific lines.

2   Descriptive Epistemic Pluralism, Epistemic Incommensurability and Relativism

Everyone ought to be an epistemic pluralist in a very minimal (and uninteresting) sense: we should all agree that not all epistemic systems are *identical*. Define an 'epistemic system' as a set of epistemic principles. Epistemic principles are to follow here Paul Boghossian (e.g. 2001, 2006) general normative propositions that specify the conditions under which certain beliefs have positive or negative epistemic merit.[4] Call the thesis that not all epistemic systems are identical *weak descriptive epistemic pluralism*.

We also needn't stray from what's entirely obvious to embrace a slightly stronger version of descriptive epistemic relativism according to which some epistemic systems are *significantly* different from other epistemic systems, where 'significantly' means more than simply 'non-identical'. Call this *moderate descriptive epistemic pluralism.* Consider briefly a sample of some 'wayward[5]' features of certain contemporary epistemic systems:

> *Zetetic Astronomy:* Maintains that the earth is a 'flat disk centered at the North Pole' and that the 'sun, moon, planets, and stars only a few hundred miles above the surface of the earth'.[6]

> *Amondawa tribe:* Lacks linguistic structures that relate time and space, and furthermore, lacks any abstract concept of 'time'.[7]

> *Yaohnanen:* This tribe is convinced that Prince Phillip, the current Duke of Edinburgh, is the embodiment of a spirit that was born on a volcano on their island, the 'pale-skinned son of the mountain spirit'.[8]

The Zetetic Astronomers, the Amondawa and the Yaohnanen part ways with us not just peripherally, but fundamentally, when it comes to the nature of the planet we occupy, the relatedness of time and space, and Prince Philip's divine attributes, respectively. We should all be moderate descriptive epistemic pluralists because there's plenty of evidence for the *existence* of these kinds of belief systems.

This said, let's revisit more carefully the component of epistemic relativism that Paul Boghossian calls *epistemic pluralism*, a doctrine that is comprised of two theses:

> **Epistemic pluralism**: (i) 'there are many fundamentally different, genuinely alternative epistemic systems', but (ii) 'no facts by virtue of which one of these systems is more correct than any of the others'.[9]

The first leg is effectively the thesis we are calling *moderate descriptive epistemic pluralism* (hereafter, MDEP), a thesis we should all accept. The second thesis is, however, deeply controversial. Whereas absolutists and relativists alike should accept MDEP, absolutists deny (ii) while relativists accept it. From the absolutists' perspective, there needn't be anything philosophically significant about diversity of opinion, even when it is radical. The absolutist's stock diagnosis of radically divergent epistemic systems (and ensuing disagreements owing to such divergent systems) will be that at least one such system is riddled with error. The relativist is inclined to a different diagnosis. For ease of exposition, let's call the second leg of the epistemic pluralist's thesis *equal standing*, to capture the claim that there are no facts by virtue of which any epistemic system is more correct than any other.

One reason many epistemic relativists are inclined to *equal standing* once MDEP is granted is that they accept, in addition to MDEP, a further thesis about the possibility conditions of rational persuasion—*epistemic incommensurability*.

> **Epistemic incommensurability (EI)**: It is possible for two agents to have opposing beliefs which are rationally justified to an equal extent where there is no rational basis by which either agent could properly persuade the other to revise their view.

Consider that if MDEP was false—that is, if everyone more or less shared the same epistemic system—then this fact would undercut the putative motivation for embracing EI, at least, in so far as the kinds of agents we're interested in are creatures like us (rather than, say, possible agents very different from us). After all, if everyone accepted (more or

less) the same epistemic system (with only peripheral differences), then it's hard to see why rational persuasion *wouldn't* at least in principle be possible.

However, the very possibility of rational persuasion does indeed look much bleaker once we grant that epistemic systems can and do radically diverge, *a la* MDEP. To appreciate why a relativist might be inclined to accept EI, given MDEP, just consider Richard Rorty's (1980) notable diagnosis of the famous seventeenth-century dispute between Galileo and Cardinal Bellarmine concerning the truth of heliocentrism, where each reached different and incompatible conclusions. Galileo claimed that heliocentrism was true, and he appealed to telescopic evidence (and more generally to Western science) in support of his claim; Bellarmine denied heliocentrism, embracing instead geocentrism, on the basis of Scriptural evidence, supplied by the Catechism of the Catholic Church.[10] Each verdict came out justified, respectively, with reference to each's *own* accepted epistemic system, but not to the other's. And further, neither seemed very well positioned to rationally persuade the other.[11]

In short, MDEP, taken to its logical conclusion, gives us cases of deep disagreement—viz. not only disagreement about some target proposition(s) but also disagreement about what kinds of evidence are even relevant to adjudicating the target proposition(s).[12] The relativist takes such deep disagreements that are the natural consequence of MDEP as evidence for EI, a thesis that itself implies the *equal standing* leg of epistemic pluralism. The absolutist by contrasts blocks the move from MDEP to EI to equal standing by simply denying that MDEP gives rise to EI. There are various ways to do this, both sceptical and non-sceptical,[13] though these strategies needn't concern us here.

What I want to suggest now is that certain substantive commitments in epistemology, when paired with MDEP, might arguably at least lead straight to EI (and thus to *equal standing*), even for epistemologists who, in embracing these substantive commitments, would prefer to steer clear of relativism.

Enter here Wittgensteinian 'hinge' epistemology—an epistemological approach inspired by Wittgenstein's posthumous writings in *On Certainty* (hereafter, OC), and which encourages a novel way of

thinking about the structure of rational support, one which has profound implications for the status of our foundational beliefs—viz. beliefs of the sort G.E. Moore (1925, cf., 1939) had, erroneously in Wittgenstein's view, appealed to in an attempt to prove the existence of the external world. The gist of Wittgenstein's positive project can be framed against a background dissatisfaction with Moore. As Wittgenstein saw it, Moore was in no position to provide evidence for the claim that he has hands, a claim for which Moore took himself to be absolutely certain.[14] Wittgenstein writes:

> My having two hands is, in normal circumstances, as certain as anything that I could produce in evidence for it. *That is why I am not in a position to take the sight of my hand as evidence for it.* (OC, §250, my italics)

If the structure of rational support relations is such that rational support flows only from what is more certain to what is less certain, but not the other way around, then what is to be said for the epistemic status of those beliefs which are most certain to one? On Wittgenstein's view, such certainties—what he calls 'hinges'—are themselves *arational* (not subject to rational support or doubt[15]) though also entirely necessary for the practice of rational evaluation. As Wittgenstein puts it:

> the *questions* that we raise and our *doubts* depend upon the fact that some propositions are exempt from doubt, are as it were like hinges on which those turn … that is to say, it belongs to the logic of our scientific investigations that certain things are *in deed* not doubted. But it isn't that the situation is like this: We just *can't* investigate everything, and for that reason, we are forced to rest content with assumption. If I want the door to turn, the hinges must stay put. (OC §§341−3)

As Duncan Pritchard (2015, p. 66) notes, what emerges from Wittgenstein's thinking about hinge propositions and their role within the practice of rational evaluation is that all rational evaluation is *essentially* local. Call this the *locality of rational evaluation thesis*, which Pritchard articulates as follows:

> **Locality of Rational Evaluation Thesis (LRET)**: '[…] all rational evaluation is essentially local, in that it takes place relative to fundamental commitments which are themselves immune to rational evaluation, but which need to be in place for a rational evaluation to occur' (2015, p. 66).

The descriptive claim MDEP, against a background commitment to LRET, receives a more sophisticated gloss. Call the version of MDEP that is available to proponents of LRET, MDEP+:**MDEP+:** There are many fundamentally different, genuinely alternative sets of hinge propositions, which are themselves immune to rational evaluation, and which need to be in place for essentially local (i.e. hinge-relative) rational evaluation to occur.

It looks very much like, even though MDEP itself doesn't entail EI, MDEP+ *does* entail EI. After all, if rational evaluation is essentially local, and if there are fundamentally different, genuinely alternative sets of hinge propositions, then—in cases of deep disagreements (viz. where a dispute concerns both some first-order proposition as well as what kinds of evidence are even relevant to settling it—it seems as though there would be no rational basis by which either agent could properly persuade the other to revise their view. After all, each would be justified relative to her own (local) set of hinge propositions, which are fundamentally different from one's interlocutor's hinges.

Putting this all together, it looks as though the Wittgensteinian hinge epistemologist, in virtue of embracing the thesis that all rational evaluation is essentially local (LRET), cannot accept moderate descriptive epistemic pluralism—the thesis that 'there are many fundamentally different, genuinely alternative epistemic systems' without also going all in for the more controversial 'equal standing' leg of epistemic pluralism, viz., the leg that is entailed by the epistemic incommensurability thesis, and which insists that there are 'no facts by virtue of which one of these systems is more correct than any of the others'.[16]

If the foregoing is right, then a consequence is that 'hinge' epistemology is *essentially* a relativistic epistemology—and indeed, this is the interpretation that many commentators have preferred.[17] As Annalisa Coliva (2010, p. 1) has suggested, 'Relativists and anti-relativists alike

are nowadays mostly united in considering Wittgenstein an epistemic relativist'.[18]

However, things on this score might not be so straightforward. In recent work, Duncan Pritchard has defended a creative anti-relativistic manoeuvre available to the Wittgensteinian, one that will require that we examine more carefully the relationship between MDEP+ and EI.

3   Pritchard on Descriptive Pluralism and Radically Divergent Hinges

MDEP+, to which the Wittgensteinian is committed, insists that there are many 'fundamentally different', genuinely alternative sets of hinge propositions. In §1, we considered some examples of epistemic systems that are not *merely* non-identical (as was implied by weak descriptive epistemic relativism), but whose differences to each other and our own are striking. It is this fact that seems incontestable. Although it seems apparent that once MDEP+ is granted, EI looks inevitable, this transition between MDEP+ and EI might be too quick.

As Duncan Pritchard (2015) puts it:

> That all rational evaluation takes place relative to hinge commitments is entirely compatible with there being a great deal of overlap in subjects' hinge commitments, *even when they are from very different cultures*. So the question we need to ask is whether there can be a *radical divergence* in one's hinge commitments[19]. (2015, p. 109, my italics)

As Pritchard is using 'radical' here and elsewhere (see, e.g. Pritchard 2010, pp. 279–80), he is taking radical to exclude at least some shared background. MDEP+ is compatible with significant divergence between epistemic systems that are not radically different, in this sense. But given that a shared background could at least potentially be appealed to in the course of rational persuasion, it will take a slightly stronger version of MDEP+ to imply EI (and a fortiori equal standing). Call this stronger version MDEP++.

> **MDEP++**: There are many *radically* and fundamentally different, genu-inely alternative sets of hinge propositions, which are themselves immune to rational evaluation, and which need to be in place for essentially local (i.e., hinge-relative) rational evaluation to occur.

But, Pritchard denies MDEP++. This is important because, if it takes MDEP++, and not merely MDEP+ to commit the hinge epistemologist to EI, then a viable rationale for rejecting MDEP++ allows the hinge epistemologist, who happily embraces LRET, to *reject* the equal standing leg of the epistemic pluralist thesis—viz., the relativistic leg of the thesis—even while maintaining a plausible version of the descriptive leg of the thesis, MDEP+. Pritchard's own rationale for rejecting MDEP++ is grounded in Wittgenstein's own thinking about the necessity of a shared background for the possibility of disagreement. For example, Wittgenstein says:

> The truth of my statements is the test of my understanding of these statements. That is to say: if I make certain false statements, it becomes uncertain whether I understand them. (OC, §§80–81)

> In order to make a mistake, a man must already judge in conformity with mankind. (OC, §156)

In his assessment of these passages, Pritchard remarks:

> If this is right, then we can at least count on a dispute between two agents who are intelligible to one another as involving a shared background of beliefs, and hence as having shared hinge commitments (2015, p. 210).

Furthermore, he adds:

> The general idea in play here is a familiar one: that wholesale error in an agent's beliefs or statements is not even intelligible. As Donald Davidson (1983, p. 432) famously put this point, "belief is in its nature veridical." Such error would thus be a reason for regarding the agent concerned as not being a believer or as not being someone who is making statements in the first place (e.g. it could be a reason to think that the agent is mad). At

the very least, wholesale error of this sort would be a reason for thinking that the agent in question is not making the statements that she seems to be making, or doesn't have the beliefs that were previously credited to her.

In summary, the picture that Pritchard offers is one that combines two very different kinds of descriptive claims, so as to yield a nuanced version of descriptive epistemic pluralism that—even for the Wittgensteinian who insists that all rational evaluation is essentially local—arguably does not lead to the epistemic incommensurability thesis. The first descriptive claim, which Pritchard (uncontroversially) takes for granted, is that—for the hinge epistemologist—it must be countenanced that there are fundamentally different sets of hinges; this is a concession of epistemic diversity. The second descriptive claim—one which implies a rejection of MDEP++— is that at least some shared background must underlie even disputes framed against very different sets of hinges, and such a shared background brings with it possibility conditions for (in principle) rational adjudication of the sort incompatible with EI.

4   Shared Backgrounds and the Über Hinge Commitment

Pritchard's position represents an innovative strategy for reconciling, on the one hand, a concession of epistemic diversity, with, on the other, the Wittgenstenian's contention that all rational assessment is essentially local, *without* devolving into EI and thus to the relativistic *equal standing* leg of the epistemic pluralism thesis.

But let's take a closer look at the claim relied upon to get this result—viz. that recognising a minimal shared background, even in cases where interlocutors are committed to very different hinges, blocks the passage from MDEP+ (which the hinge epistemologist *is* committed to) to EI. At the core, the claim can be pared down to the following *prima facie* plausible idea: that a shared background supplies conditions for possible rational adjudication. Put another way: for two interlocutors, A and B, and dispute D, D is not in principle rationally inadjudicable for A and B if there is some shared background, B, which is a necessary precondition for D.

I think this is more or less right. However, in recent work, I've criticised a version a special case of this general claim–not by rejecting it outright, but by opting for a modification of it.[20] I want to briefly review my proposed modification and then outline how it influences Pritchard's strategy. To this end, just imagine a dispute between—say—a Western scientist, and a member of the Yaohnanen tribe considered in §1 (i.e. the tribe that is convinced Prince Phillip is a divine being). Suppose we grant Pritchard the broadly Davidsonian point that any meaningful disagreement between the scientist and the Yaohnanen requires at least some shared background. Let's imagine that the scientist the leader of the Yaohanen (call them Cy and Yao, respectively) attempt to rationally dispute the claim that Prince Phillip is divine. Cy insists that there is no scientific evidence that Phillip is divine. Yao appeals to Yaohanen folk legend, according to which a 'pale-skinned' (Phillip is Caucasian) son of a mountain spirit is said to have travelled to a faraway land (e.g. England), married a powerful woman (in this case, Queen Elizabeth) and then returned to them (Phillip visited the island of Vanuatu in 1974). Prince Phillip seemed to the Yaohnanen to fit these criteria to a tee. Of course, Cy rejects that Yaohanen folk lore is relevant to the truth of the proposition that Phillip is divine, and the Yaohanen think it clearly is. As things stand, rational persuasion looks unpromising.

In a case like this one, it's important to note that *merely* identifying certain points of agreement between Cy and Yao needn't be sufficient for demonstrating the possibility that rational adjudication is possible in the light of their respective starting points. Suppose, in the dialectical situation described, it is pointed out to both parties that there is some epistemic norm both are happy to accept—viz. a tautological norm that says *Infer A from A.* Both Cy and Yao can agree that such a norm is appropriately neutral between them. However, it's not yet evident that this fact is significant *vis-à-vis* whether the two parties could rationally adjudicate their dispute. The reason is that such norm, though appropriately neutral, is not appropriately *discriminatory* in the sense that such a norm would fail to favour one position over the other; as such, it would plausibly fail to play the kind of role that it would need to in order to facilitate bringing interlocutors locked in an otherwise

irreconcilable position into a (non-question-begging) resolution. There is another side to this coin: a norm (or proposition) that that does very well in the 'discriminatory role' will plausibly fail to be appropriately neutral and thus is not something that can by the lights of both parties be introduced into a shared background.[21]

The lesson to be learned is that for two interlocutors, A and B, who antecedently accept very different hinges, and dispute D, D is possibly rationally inadjudicable for A and B *even if* there is some shared background, B, which both A and B are free to appeal to. Relevant to whether rational adjudication would be possible in such a circumstance is what, specifically, is shared in the background. It's not clear that the kind of shared background that is implied by Pritchard's Davidsonian line will be enough to vouchsafe the possibility of rational adjudication in such circumstances; this is because it's unclear whether what would be secured in such a background would be appropriately discriminatory, in the sense just articulated. And if that's right, then it becomes less clear whether MDEP++ can be rejected by the hinge epistemologist on broadly Davidsonian grounds.

The dialectic at this point can be summarised as: the challenge for the Wittgensteinian who wants to (i) concede substantial epistemic diversity (of the sort captured by MDEP+) while embracing LERT (that rational evaluation is essentially local) without also accepting the epistemic incommensurability thesis (and a fortiori, *equal standing*) that seems to follow from these claims, will be to show how MDEP++ is false. MDEP++, recall, says that there are many *radically* and fundamentally different, genuinely alternative sets of hinge propositions, which are themselves immune to rational evaluation, and which need to be in place for essentially local (i.e. hinge-relative) rational evaluation to occur. I've argued that the Davidsonian line embraced by Pritchard, which involves merely highlighting that some shared background is a necessary precondition for even seemingly deep forms of disagreement, is insufficient for demonstrating the falsity of MDEP ++ because it is in principle compatible with MDEP++.

Interestingly, Pritchard's recent anti-sceptical work (e.g. 2015, Chap. 4) reveals a further possible anti-relativistic move that the hinge epistemologist could make at this juncture, one that appeals to what

Pritchard calls an *über hinge commitment.* Pritchard's *über hinge* strategy might appear promising even if the Davidsonian line was not for the Wittgensteinian who wishes to reject the EI-entailing MDEP++. Pritchard articulates the notion of an über hinge (in connection to ordinary hinge commitments) as follows:

> […] closer inspection of this apparently heterogeneous class of hinge commitments reveals that they all in effect codify, for that particular person, the entirely general hinge commitment that one is not radically and fundamentally mistaken in one's beliefs. Call this commitment the *über hinge commitment*, and call the proposition endorsed by the über hinge commitment the *über hinge proposition*.

Perhaps, if as Pritchard elsewhere (e.g. 2015, p. 206) puts it, 'A commitment to the absence of […] systematic deception is […] a plausible manifestation of one's general über hinge commitment', we could locate a very specific item that will be in any shared background *beyond* whatever must be in the shared background between interlocutors simply (*a la* Davidson) for genuine disagreement to be meaningful. After all, we can imagine champions of very different epistemic systems failing to *radically* diverge (in the sense relevant to MDEP++) simply because their shared commitment to regarding themselves as not radically deceived might generate for them a stock of mutually recognisable propositions, some of which could (perhaps) be not only appropriately neutral but appropriately discriminatory *vis-à-vis* their epistemic disagreement.

Or course, one line of criticism to this strategy proceeds as follows: for the *über-hinge* strategy to be better positioned than the Davidsonian strategy for rejecting MEDP++, further argument is needed for why shared commitment to the *über-hinge proposition* (on behalf of both parties representing very different epistemic systems) is such that, by appealing to it, rational adjudication of the dispute would be possible. Without additional 'proof of concept', the epistemic-incommensurability-entailing MDEP++ seems to remain at least potentially on the table as a plausible commitment of the hinge epistemologist (given her acceptance of LRET).

There might, however, be a further problem for such a strategy. Here it will be helpful to consider how a recent worry raised by Crispin Wright's (2012 to an anti-sceptical strategy employed by Annalisa Coliva (2012) could potentially be 'redeployed' as a challenge for advancing the über-hinge strategy as an 'anti-relativistic' strategy on behalf of the hinge epistemologist. Here's Wright (2012):

> if it really were constitutive of our conception of rational empirical enquiry to assume that there is an external material world, then there should be a kind of unintelligibility about a sceptical challenge to the rationality of this assumption which would be at odds with the sense of paradox created by the best sceptical arguments that challenge it. There is, it seems to me, an implicit tension in the very notion that elements which are constitutive of a concept—which belong primitively to its identity and are not sustained by other features of it—should be sufficiently opaque to be controversial and apparently vulnerable to philosophical challenge. If *free* action, to take a parallel example, is, conceptually constitutively, simply action performed with a sense of freedom, for normal human reasons, without external force or duress, why does anyone feel the familiar kind of challenge posed by determinism as any kind of problem? (2012, p. 479)

Here, Wright submits as problematic any anti-sceptical strategy that renders unintelligible something that appears to be intelligible. In this case, Wright is taking it that a sceptical challenge to the assumption that there is an external world is at least intelligible, and it is *ceteris paribus* problematic if a given anti-sceptical strategy fails to preserve this. Note that I am here taking no stand as to whether Coliva's own anti-sceptical strategy, to which Wright is levelling this critique, actually succumbs to it. Rather, I want to register the desideratum being highlighted here as a valid one.

I want to now consider that the worry Wright poses to Coliva (framed in terms of this desideratum) can potentially be recast as a worry for the *über hinge* strategy, in so far as it is meant to block the move from descriptive to controversial epistemic pluralism, via a rejection of MDEP++. To this end, let's take as a starting point that, just as

the sceptical problem is an intelligible one (one that raises a genuine philosophical tension), so likewise is the problem of epistemic relativism. That is, there is, as anti-relativist Paul Boghossian (2001) sums it up, '[…] a serious difficulty seeing how there could be objectively valid reasons for belief, a difficulty that has perhaps not been adequately faced up to in the analytic tradition' (2001, p. 1). But, if it really was constitutive of the practice of rational assessment (as per the über-hinge strategy) that certain hinges be in common between even the most radically different epistemic systems, then there would be a kind of unintelligibility about the relativist's challenge, construed as a challenge to vindicate rational assessment as objective (in a manner incompatible with the 'equal standing' leg of epistemic pluralism). To the extent that this concern is on the right track, an über-hinge styled rationale for rejecting MDEP++ will need to be paired with an accompanying story for how the problem posed by epistemic relativists is as intelligible as it seems.

5   Non-relativistic 'Hinge' Epistemology? Some Further Thoughts

Thus far, I have not commented on the vexing question of what is *distinctive* of a philosophical position rightly called epistemic relativism. Rather, I've simply noted that the more provocative (non-descriptive) leg of the epistemic pluralism thesis (no facts by virtue of which any epistemic system is more correct than any of the others) is often taken to be a feature of epistemic relativism. Indeed, most self-ascribed epistemic relativists have embraced this thesis.[22]

However, I want to close by considering how—potentially at least—the hinge epistemologist could actually welcome MDEP++, and thus embrace EI along with the 'equal standing' leg of the epistemic pluralist thesis, all while maintaining that doing so is not sufficient for being an epistemic relativism of a genuinely philosophically interesting sort. This suggestion perhaps sounds bizarre initially, but to appreciate how it might go, it will require that we look more carefully at Boghossian's

epistemic pluralist thesis, as it stands embedded (along side an epistemic non-absolutist thesis and an epistemic relationist thesis) within the wider position Boghossian defines as epistemic relativism, as follows:

Epistemic Relativism (Boghossian's Formulation)

A. There are no absolute facts about what belief a particular item of information justifies. (*Epistemic non-absolutism*)

B. If a person, S's, epistemic judgments are to have any prospect of being true, we must not construe his utterances of the form

"E justifies belief B"

as expressing the claim

E justifies belief B

but rather as expressing the claim:

According to the epistemic system C, that I, S, accept, information E justifies belief B. (*Epistemic relationism*)

C. There are many fundamentally different, genuinely alternative epistemic systems, but no facts by virtue of which one of these systems is more correct than any of the others. (*Epistemic pluralism*)

The matter of whether this is in fact the right way to characterise the epistemic relativist's core insight—as a conjunction of these three claims—has been a matter of contemporary dispute. In particular, MacFarlane (2014) and Wright (2008) have called into doubt Boghossian's inclusion of the *epistemic relationist* thesis, and whether its inclusion can satisfactorily preserve the difference between relativism proper and contextualism.

Consider that epistemic contextualists (e.g. DeRose 1992, 2009) insist that the extension of 'knows' or 'justified' in attributions of the form 'S knows that p' or 'S is justified in believing P' varies with the context in which these terms are used, and accordingly sentences that attribute these terms can express different propositions and have different truth conditions in different contexts (of use). However, while contextualists agree that knowledge/justification ascribing sentences do not

get a truth value, *simpliciter*, but only relative to a 'standards parameter' whose value is supplied by the context in which they are used, contextualists are happy to agree that justification/knowledge ascribing sentences have their truth values *absolutely*. That is to say, once the value of the relevant standards parameter is supplied (i.e. in the case of attributor contextualism, the standards operant in the attributor's context), there is a once-and-for-all answer to the question of whether a given justification/knowledge attribution is true.

According to MacFarlane, this is what a proper relativist denies. On MacFarlane's brand of epistemic relativism about 'knows', whether a particular use of a knowledge-ascribing sentence, e.g. 'George knows he has a hand' is true depends on the epistemic standards at play in the *assessor*'s context—viz. the context in which the knowledge ascription is being assessed for truth or falsity. But, given that the very same knowledge ascription can be assessed for truth or falsity from indefinitely many perspectives, when I say that George knows he has a hand, what I've said does not get a truth value absolutely, but only relatively.[23]

Against this background, we can see that—if the hinge epistemologist opts for MacFarlane's way of thinking about what makes a philosophical position relativistic in an interesting sense, there is at least some scope for her to embrace *both* legs of the epistemic pluralist thesis while at the same time purporting to avoid relativism (proper) by simply embracing (along with epistemic pluralism) the epistemic relationism thesis that MacFarlane takes to disqualify Boghossian's epistemic relativist from being a relativist, proper (MacFarlane (2014, p. 33, fn. 5) reads Boghossian's relativist as a contextualist). On the resulting view, the Wittgensteinian can insist that rational assessment is essentially local in the sense that knowledge/justification attributions do not get truth values independent of local hinges that supply the relevant standards, but that once these standards are supplied, knowledge/justification attributions get their truth-values absolutely.

MacFarlane's characterisation of relativism in terms of assessment sensitivity is controversial, and it's unclear whether the hinge epistemologist would be satisfied by this characterisation of the landscape in the first place (or, satisfied enough to advert to it in justifying how a hinge epistemologist's commitment to both legs of the epistemic pluralist

thesis might be taxonomised as contextualism rather than relativism). I simply register this point as a means of showing the full scope of the hinge epistemologist's options in the face of the original dilemma. The original dilemma for the hinge epistemologist—viz. either reject (as Pritchard has) some version of descriptive epistemic pluralism or embrace the most provocative leg of the pluralist thesis—remains I think one of the most vexing issues facing hinge epistemologists who are convinced that rational assessment is a local affair while at the same time put off by the prospect of full-blown relativism.

Notes

1. (Boghossian 2006, p. 73).
2. For a notable expression of the more idea that descriptive forms of relativism, do not entail relativism in a philosophically interesting sense, see Rachels (2003, e.g. pp. 16−23); see also Boghossian (2006, Chap. 5–6). Cf., Baghramian and Carter (2015, §2.1) for an overview of various discussions of descriptive relativism in epistemology.
3. I will be using 'hinge epistemology' to refer to the epistemological theses clustered around the notion of a 'hinge proposition' as advanced in Wittgenstein's posthumous *On Certainty* (1969). See also Coliva (2010b) and Pritchard (2015) for some notable contemporary expositions of this approach.
4. At least provided we are cognitivists, epistemic principles can be true or false. By contrast, epistemic *rules* are prescriptive; for example, *In conditions C, belief in way W.* Epistemic principles, if true, imply epistemic rules or norms which, qua prescriptive, are not truth-apt. For example, if the principle *Beliefs formed on the basis of wishful thinking are unjustified* is true, then its truth gives rise to a non-truth-apt prescriptive epistemic rule to the effect that: *If your basis for B is wishful thinking, don't believe B.*
5. I borrow this expression from Quassim Cassam (2016).
6. https://www.theflatearthsociety.org/home/index.php/about-the-society.
7. See, for example, Sinha et al. (2011). For some examples of similar cases of epistemic diversity that include the unavailability of certain

concepts for the purposes of reasoning and planning, see (Hacking 1982).

8. http://news.bbc.co.uk/1/hi/programmes/from_our_own_corre-spondent/6734469.stm.

9. (Boghossian 2006, p. 73).

10. For a detailed historical account of this dispute, see Finocchiaro (2009).

11. For a detailed discussion of this case, see Carter (2016, Chap. 4). See also Boghossian (2006, Ch. 5–6) and Siegel (2011).

12. I am using 'deep disagreement' here in a way that is consonant with Hales (2014).

13. The most notable sceptical strategy for blocking the move from MDEP to EI is the Pyrrhonian strategy, which recommends withholding of judgment in the face of disagreement, or equipollence. For discussion of this strategy, see Carter (2016, Ch. 3).

14. For a recent overview of contemporary work on Moore's proof, see Carter (2012).

15. Just as hinge propositions cannot be rationally supported, neither can they be rationally doubted (e.g. OC §317, §322, §342); to doubt a hinge proposition would be to at the same time call into question the wider belief system (OC §115, §186). As Wittgenstein notes, 'If some-one doubted whether the earth had existed a hundred years ago, I should not understand, for this reason: I would not know what such a person would still allow to be counted as evidence and what not' (OC §231).

16. For related discussion on this point, see Coliva (2010a, pp. 188–190) and (2010b, pp. 1–3).

17. See, for example, Rorty 1980; Boghossian 2006; Trigg 1973.

18. See, however, Williams (e.g. 1996, 2007) for the opposite reading of Wittgenstein's epistemology, as a kind of 'antidote' to epistemic relativ-ism. Coliva (2010b) and Pritchard (e.g. 2015, *passim*) also resist the prevailing relativistic reading of Wittgenstein.

19. Pritchard is advancing this claim in criticism of Michael Williams' (1996, 2007) brand of Wittgensteinian contextualism, which Pritchard takes to overlook this point.

20. See Carter (2016, Chap. 4).

21. See Carter (2016, §4.3) for a more detailed presentation of this point.

22. See, for example, Richard Rorty (1980, 1989), Ian Hacking (1982), Paul Feyerabend (1975) and Thomas Kuhn (1962) among others.

23. See Carter (2016, Chap. 7; 2017) for further discussion of these differences.

References

Baghramian, Maria, and J. Adam Carter. 2015. Relativism. In *Stanford Encyclopaedia of Philosophy*, ed. Edward N. Zalta, Fall 2015, 1–46. http://plato.stanford.edu/entries/relativism/.

Boghossian, Paul. 2001. How Are Objective Epistemic Reasons Possible? *Philosophical Studies* 106 (1): 1–40.

———. 2006. *Fear of Knowledge: Against Relativism and Constructivism*. Oxford: Oxford University Press.

Carter, J. Adam. 2012. Recent Work on Moore's Proof. *International Journal for the Study of Skepticism* 2 (2): 115–144.

———. 2016. *Metaepistemology and Relativism*. London: Palgrave Macmillan.

———. 2017. Epistemological Implications of Relativism. In *Routledge Handbook of Epistemic*, ed. J.J. Ichikawa. London: Routledge.

Cassam, Quassim. 2016. Persuasion and Unlearning.

Coliva, Annalisa. 2010a. *Moore and Wittgenstein: Scepticism, Certainty, and Common Sense*. London: Palgrave Macmillan.

———. 2010b. Was Wittgenstein an Epistemic Relativist? Philosophical Investigations 33 (1): 1–23.

Davidson, Donald. 1983. A Coherence Theory of Truth and Knowledge. In *Kant Oder Hegel?*, ed. D. Henrich. Stuttgart: Klett-Cotta.

DeRose, Keith. 1992. Contextualism and Knowledge Attributions. *Philosophy and Phenomenological Research* 52 (4): 913–929.

———. 2009. *The Case for Contextualism: Knowledge, Skepticism, and Context*, vol. 1. Oxford: Oxford University Press.

Feyerabend, Paul. 1975. *Against Method*. London: Verso.

Finocchiaro, Maurice A. 2009. The Galileo Affair. *Physics World* 84.

Hacking, Ian. 1982. Language, Truth and Reason. In *Rationality and Relativism*, eds. M. Hollis and S. Lukes, 48–66. Cambridge: MIT Press.

Hales, Steven D. 2014. Motivations for Relativism as a Solution to Disagreements. *Philosophy* 89 (01): 63–82.

Kuhn, Thomas. 1962. *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.

———. 2014. *Assessment Sensitivity: Relative Truth and Its Applications*. Oxford: Oxford University Press.

Moore, George Edward. 1925. A Defence of Common Sense. In *Contemporary British Philosophy*, ed. J.H. Muirhead. London: Allen & Unwin.

———. 1939. Proof of an External World. *Proceedings of the British Academy* 25 (5): 273–300.

Pritchard, Duncan. 2009. Defusing Epistemic Relativism. *Synthese* 166 (2): 397–412.

———. 2010. Epistemic Relativism, Epistemic Incommensurability and Wittgensteinian Epistemology. In *A Companion to Relativism*, ed. Steven D Hales, 266–285. Malden: Wiley Blackwell.

———. 2015. *Epistemic Angst*. Princeton: Princeton University Press.

Rachels, James. 2003. *The Elements of Moral Philosophy*, 4th ed. New York: McGraw-Hill.

Rorty, Richard. 1980. *Philosophy and the Mirror of Nature*. Princeton: Princeton University Press.

———. 1989. *Contingency, Irony and Solidarity*. Cambridge: Cambridge University Press.

Siegel, Harvey. 2011. Epistemological Relativism: Arguments Pro and Con. In *A Companion to Relativism*, ed. Steven Hales, 199–218. Wiley Online Library.

Sinha, Chris, Vera Da Silva, Jörg Zinken Sinha, and Wany Sampaio. 2011. When Time Is Not Space: The Social and Linguistic Construction of Time Intervals and Temporal Event Relations in an Amazonian Culture. *Language and Cognition* 3 (1): 137–169.

Trigg, Roger. 1973. *Reason and Commitment*. Cambridge: Cambridge University Press.

Williams, Michael. 1996. *Unnatural Doubts: Epistemological Realism and the Basis of Scepticism*. Princeton: Princeton University Press.

———. 2007. Why (Wittgensteinian) Contextualism Is Not Relativism. *Episteme* 4 (1), 93–114.

Wittgenstein, Ludwig. 1969. *On Certainty*. Oxford: Blackwell Oxford.

Wright, Crispin. 2008. Fear of Relativism? *Philosophical Studies* 141 (3): 379–390.

———. 2012. Replies Part IV: Warrant Transmission and Entitlement. In *Themes from the Philosophy of Crispin Wright* ed. Annalisa Coliva 451–486. Oxford: Oxford University Press.

Part IV

Epistemic Pluralism: Some Applications

How to Be a Pluralist About Self-Knowledge

Annalisa Coliva

1   Introduction

What does it mean to be a pluralist about self-knowledge? Since this view and its name have not been appropriated before, there is still considerable room for maneuver. In this chapter, I am going to present my own preferred way of characterizing this position. Before delving into that, let me briefly connect pluralism about self-knowledge, as I understand it, with epistemic pluralism in general.

A. Coliva (✉)
Department of Philosophy, University of California, Irvine,
CA 92697, USA
e-mail: a.coliva@uci.edu

My own version of pluralism about self-knowledge is characterized by adherence to the idea that there is an asymmetry between what may be called first-personal and third-personal self-knowledge. In both cases, we do have true judgments of the form "I am F," where *F* ranges over psychological predicates. Yet, while in the third-personal case, these self-ascriptions do enjoy an epistemic support, in the first-personal one, they actually don't. Yet, we can still call them instances of self-*knowledge*, once we are hospitable to the idea that that very term may be used to refer to different states or properties of a subject. Moreover, it is part of my brand of pluralism that, when the relevant true psychological self-ascriptions enjoy an epistemic support and thus qualify as instances of third-personal self-knowledge, the methods by means of which such a support is obtained are many and diverse.

Hence, pluralism about self-knowledge, as I conceive of it, entails both pluralism of states and of methods. That is to say, it is not the case that in all instances of self-knowledge, subjects are in the same epistemic situation—that is, roughly, the situation of having a true and justified belief about their own mental states. Nor is it the case that when they actually do have a true and justified belief about their own mental states, the way they have acquired that justification goes through the same kind of method.

An important aspect of the kind of pluralism I will be presenting in the following is its meta-epistemological import. For, by embracing it, several disputes about *the* (allegedly) correct account of self-knowledge will dissolve. In particular, as we will see, various so far competing accounts will be reconciled with one another by showing how each of them offers a suitable account of at least some instances of self-knowledge, or of at least some of the ways in which third-personal self-knowledge can be achieved.

That, however, does not mean that, once epistemic pluralism about self-knowledge is embraced, anything goes. Pluralism about self-knowledge, as I conceive of it, does not entail relativism. Hence, I am not committed to the view that several allegedly competing accounts of self-knowledge would in fact all be on a par, the choice between them being only based on pragmatic factors. As we shall see, there is still considerable room for philosophical dispute regarding the correct account of

some, appropriately specified, specimens of self-knowledge. Indeed, the kind of pluralism about self-knowledge I will be presenting is intimately connected to first-order epistemological views, which I have developed and defended elsewhere.[1] Hence, it actually depends on taking a firm stance on certain first-order epistemological issues.

2   Pluralism About Self-Knowledge: An Outline

The main claim at the heart of pluralism about self-knowledge, as I conceive of it, is an *existential* one. Namely, that there is an *asymmetry* between first- and third-personal cases of self-knowledge. Examples of the former kind of self-knowledge are one's immediate judgment "I am in pain" issued after having just banged one's knee against a table, or "I intend to take my son out for lunch tomorrow," after weighing various reasons, such as the fact that your son seems to need some quality time just with you and that the first available option would be lunch tomorrow.

A classic example of the latter kind of self-knowledge is the one described by Jane Austen, in her novel *Emma*, when the protagonist realizes her love for her long-lasting friend Mr. Knightley.

> Emma's eyes were instantly withdrawn; and she sat silently meditating in a fixed attitude, for a few minutes. A few minutes were sufficient for making her acquainted with her own heart. A mind like hers, once opening to suspicion, made rapid progress. She touched – she admitted – she acknowledged the whole truth. Why was it so much the worse that Harriet should be in love with Mr. Knightley than with Mr. Churchill? Why was the evil so dreadfully increased by Harriet's having some hope of return? It darted through her, with the speed of an arrow, that Mr. Knightley must marry no-one but herself.

Other examples may be one's realization of not being brave, or of being intimidating, or of being biased, based on various epistemic methods we will review in the following.

Roughly, examples of first-personal self-knowledge are immediate manifestations of one's *ongoing* sensations, and intentions (as well as other commissive propositional attitudes), but also basic emotions—like "I'm scared of that dog," based on one's ongoing fear at the sight of an aggressively looking barking dog; ongoing perceptions—"I'm seeing a dog"; passing thoughts—"I'm thinking of my mother"; sudden recollections—"I remember where I left my keys yesterday"; and *cogito*-like thoughts—"I'm hereby thinking that it's sunny today."

Examples of third-personal self-knowledge, in contrast, are characteristically self-ascriptions of *dispositional* psychological properties, like being in love or jealous, being brave or intimidating or biased. The ways in which they are arrived at closely resemble the methods we usually employ to make knowledgeable judgments regarding *other* people's mental states, whence the label "third-personal self-knowledge." For just as we need to make inferences to the best explanation in order to attribute psychological properties to other people, based on the observation of their verbal and non-verbal behavior, so it seems that (at least in many cases, as we shall see) we arrive at those *self*-directed ascriptions by applying the same kind of method.

A corollary of this main existential claim is a *meta-epistemological* one. Namely, that a complete account of self-knowledge needs to address *both* first- *and* third-personal cases. I also couple this meta-epistemological claim with an *axiological* one. Not only are both kinds of self-knowledge in need of explanation, but, in my view, they are equally philosophically interesting and existentially important. That is to say, while different philosophers have favored one kind of self-knowledge over the other and have defended their take by downplaying the existential significance or the philosophical interest of the kind of self-knowledge their theories did not account for, I think both kinds of self-knowledge are existentially significant and philosophically interesting. That is, they are both extremely valuable to our lives and trying to philosophically account for either of them will raise equally interesting epistemological issues.

A more idiosyncratic aspect of my brand of pluralism regarding self-knowledge is that while third-personal self-knowledge is, for me, the result of an epistemic achievement, which takes a subject from

first-order mental states to true and justified beliefs about them, I hold that first-personal self-knowledge is not the result of any epistemic achievement. Rather, what goes by the name of "first-personal self-knowledge" is, in my view, just a *conceptual truth*, which can be variously redeemed. This entails that, properly speaking, there is no such thing as an epistemology of first-personal self-knowledge. Yet, for reasons we will see, it is not necessarily a misnomer to call that kind of self-*knowledge* thus.

Moreover, I hold that third-personal self-knowledge does not come about just through inference to the best explanation, based on the observation of one's verbal and non-verbal behavior. In fact, third-personal self-knowledge can be achieved through a variety of epistemic methods, which, so far, have received little attention by the philosophical community that has been working at the interface between philosophy of mind and epistemology. It is my conviction that a closer look at such a variety of methods will repay the effort and will bring out several *epistemically* interesting aspects of third-personal self-knowledge.

Finally, I also hold that the ways in which we can actually redeem the conceptual truth first-personal self-knowledge consists in are subtly different depending on the kind of mental state at issue. It is my conviction that paying attention to these subtleties will reveal important differences and will make the overall approach to first-personal self-knowledge, generally known as "constitutivism," more plausible. Let us now look at each of these main claims in more detail.

3   The Asymmetry Between First- and Third-Personal Self-Knowledge I: Against Deniers

Instances of first-personal self-knowledge are characterized by the following features. First, they are *groundless*. That is to say, the relevant psychological self-ascriptions are neither based on the observation of one's occurrent mental state, nor on inference to the best explanation based on the observation of one's overt behavior (or on any other epistemic method). This does not mean that the relevant mental state plays

no role in the relevant self-ascriptions. It simply means that the role of the first-order mental state is not that of offering an epistemic support or some kind of independent reason for its self-ascription.[2] Second, instances of first-personal self-knowledge are characterized by transparency: If one is conceptually competent, cognitively well-functioning, attentive, alert, and sincere, if one has a given mental state M, one will so judge. Third, instances of first-personal self-knowledge are characterized by authority. That is to say, if one is conceptually competent, cognitively well-functioning, attentive, alert, and sincere, if one judges to have a given mental state M, one does have it.

To exemplify, when I judge "I'm in pain," while feeling pain at my knee after having just banged it against a table, I do not do so by observing my mental state, by recognizing it as the kind of mental state it is and by judging that I myself instantiate it. Nor do I so judge by observing my moaning or even saying out loud "Ouch" or "It hurts" and inferring that I am in pain. Rather, the judgment is of a piece with the occurrence of that mental state. Indeed, it is (part of) its very manifestation. For we often avow our immediate sensations. The avowal voices or manifests the relevant mental state. Yet, it is made through a sentence of English and it expresses a proposition which has truth-conditions, which in the relevant cases are *manifested* as obtaining and not just *said* to obtain.[3] Conversely, if I so judge, while satisfying the conditions listed above, I am in pain.

Similarly, when I judge "I intend to take my son for lunch tomorrow," I do not do so by observing my intention, by recognizing it for what it is and by forming that judgment. Nor do I arrive at that judgment by observing my overt behavior, or my saying out loud "I'll take my son out for lunch tomorrow." The very self-ascription is of a piece with the occurrence of the intention, it is part of its manifestation and, as we shall see, it can actually bring it about, at least in some cases. Conversely, if I make the relevant self-ascription, under the previously mentioned conditions, that entails that I do have that intention.

By contrast, instances of third-personal self-knowledge do not exemplify any of these features. They are never groundless, because they are always based on inference to the best explanation or on other epistemic methods we shall review in the following. Nor are they ever transparent

or authoritative. For it is not enough to be possessed of the relevant concepts, be cognitively well-functioning, attentive, alert, and sincere (while satisfying also other conditions, we will review in Sect. 7), to judge that one has the relevant mental state. Indeed, for various reasons, one can satisfy all these conditions and remain "blind" to oneself. Conversely, one's judging to have a given mental state, while also satisfying the relevant conditions, does not guarantee that one's judgment be correct. One could, after all, have gone astray in the application of the various epistemic methods we normally employ in order to make these self-ascriptions.

The asymmetry between first- and third-personal self-knowledge has been denied by several philosophers. Behaviorists have denied the existence of genuinely first-personal self-knowledge. In their view, all psychological self-ascriptions are instances of third-personal self-knowledge, reached through inference to the best explanation starting with the observation of one's own behavior.[4] Contemporary deniers of the asymmetry have taken findings in cognitive psychology to show that we are very often mistaken about our own mental states and blind to their occurrence. Consider the following case, discussed by Eric Schwitzgebel (2008, p. 252):

> My wife mentions that I seem to be angry about being stuck with the dishes again (despite the fact that doing the dishes makes me happy?). I deny it. I reflect; I sincerely attempt to discover whether I'm angry—I don't just reflexively defend myself but try to be the good self-psychologist my wife would like me to be—and still I don't see it. I don't think I'm angry. But I'm wrong, of course, as I usually am in such situations: My wife reads my face better than I introspect. Maybe I'm not quite boiling inside, but there's plenty of angry phenomenology to be discovered if I knew better how to look. Or do you think that every time we're wrong about our emotions, those emotions must be nonconscious, dispositional, not genuinely felt? Or felt and perfectly apprehended phenomenologically but somehow nonetheless mislabeled? Can't I also err more directly? Surely my "no anger" judgment is colored by a particular self-conception and emotional involvement. To that extent, it's less than ideal as a test of my claim that, even in the most favorable circumstances of quiet reflection, we are prone to err about our experience. However, as long as we

focus on judgments about emotional phenomenology, such distortive factors will probably be in play. If that's enough consistently to undermine the reliability of our judgments, that rather better supports my thesis than defeats it, I think. Infallible judges of our emotional experience? I'm baffled. How could anyone believe that? Do you believe that? What am I missing?

Now, it is clearly preposterous to hold that we know we are in pain, while feeling it, because we observe our overt behavior and infer to its likely psychological cause. Yet, the case proposed by Schwitzgebel is subtler. Still, I think we can easily accommodate it. For there are cases of "cold" anger. That is, cases in which there is no feeling of anger while one's behavior clearly manifests it. What that shows is that a lot of emotions have both a phenomenological component to them and a dispositional one.[5] Yet, sometimes the former might not be present, while the latter is. No wonder, then, that with respect to the dispositional aspects of our emotions, we can be blind or not authoritative. For knowledge of the dispositional aspects of our emotions is just another case of third-personal self-knowledge.

One might protest that this is a simplistic understanding of the situation. For there could be feelings one would be aware of and yet one could have a hard time figuring out what they are feelings *of*. Again, I do not wish to dispute the possibility of being in such a predicament. Yet, this case would simply show that we could be bad at figuring out what these feelings are *symptoms* of—that is, of what kind of dispositional mental states they are a manifestation of.[6] Imagine for a moment, a less perceptive Emma, who, presented with the same inner phenomenology as the one described by Austen, had trouble realizing that it is a symptom of her being in love with Mr. Knightly. This self-blind Emma would manifestly be a poor self-interpreter, who would be incapable of acquiring third-personal self-knowledge. Yet, this does not entail that she would have no first-personal self-knowledge. After all, she could perfectly well judge "I am feeling bad," after realizing that Harriet might have some hope of return.

More generally, what findings in cognitive psychology show is that a lot of self-knowledge is third-personal, or, equivalently, that the scope of

first-personal self-knowledge is more limited than philosophers have traditionally thought. That does not mean, however, that there is no such thing as first-personal self-knowledge altogether.

4   The Asymmetry Between First- and Third-Personal Self-Knowledge II: Against Chauvinists

As mentioned in Sect. 2, a corollary of the claim that there is a genuine asymmetry between first- and third-personal self-knowledge is that a complete account of self-knowledge needs to address *both* first- *and* third-personal cases. To repeat, not only are both kinds of self-knowledge in need of explanation, but, in my view, they are equally philosophically interesting and existentially important. Various philosophers have denied that. For, although they may have acknowledged the existence of the asymmetry between first- and third-personal self-knowledge, they have tended to favor only one side of it. Characteristically, those who are all in favor of third-personal self-knowledge have pointed out that this is the only kind of self-knowledge which is worth obtaining, for it reveals important aspects of our character. Indeed, it is the kind of knowledge people are prepared to pay large amounts of money for, whenever they go for psychoanalytic therapy and counseling.[7] By contrast, first-personal self-knowledge is not that interesting as it just affords knowledge of humdrum truths such as "I am in pain" or "I believe it is sunny today" or "I intend to take my son out for lunch tomorrow."

Now, I do agree that knowledge of our character is existentially relevant. After all, the *motto* "Know thyself," inscribed on the front of the Temple of Apollo at Delphi, was meant to remind us of that. For it is only by knowing our character and our limitations that we will maximize chances of achieving a sober conception of ourselves and of making sensible lifetime decisions. What I do not agree with is inferring that since, existentially, this kind of self-knowledge is much more interesting than humdrum first-personal self-knowledge, it is the only

existentially interesting one. Obviously, knowing you are in pain right now as you are grabbing a hot pan with your bare hands may spare you a lot of further trouble. Moreover, only by knowing your own intentions and further propositional attitudes will you be able to plan action. Finally, knowledge of your inner feelings, emotions, sensations, etc. will very often provide you with the necessary data to gain third-personal self-knowledge, as Emma's case made vivid, and we shall presently discuss in more detail. Hence, first-personal self-knowledge is existentially relevant in its own right and actually a fundamental element to gain third-personal self-knowledge in very many cases.

Even less do I find it convincing that just because people are prepared to go to great lengths to get third-personal self-knowledge, that kind of self-knowledge is the only *philosophically* interesting one; or else, the only one on which philosophers should concentrate their future efforts.[8] On the contrary, accounting for first-personal self-knowledge has turned out to be one of the most difficult challenges in philosophy of mind to which virtually all major philosophers in the Western tradition have devoted considerable efforts. Despite their problems, these efforts have often produced interesting philosophical insights along the way, which, as we shall see, can be fruitfully integrated within a pluralistic approach.

That being said, let me hasten to add that I do not agree with those philosophers who have favored first-personal self-knowledge over its third-personal counterpart, precisely on the ground that only the former would pose an interesting philosophical challenge.[9] As we shall see, third-personal self-knowledge is not only existentially relevant, but also *epistemically* interesting, once one realizes that the ways in which we can obtain it are many and diverse and such as to connect the topic of self-knowledge with very important issues in epistemology, individual, and—as surprising as that might be—social.

5   The Common Route to Denial and Chauvinism: Monism

Despite their differences, deniers and chauvinists are united by a common assumption—that is, adherence to *monism*. For they tend to concentrate on one kind of mental state and offer what seems a (*prima facie*, at least) plausible explanation of how we come to know that kind of mental state. They then try to *generalize* that explanation to (almost) all other mental states, or else to *downplay* the extent and philosophical significance of those mental states whose self-knowledge they cannot account for. The attempt at generalization is the move that eventually gets all of them into trouble, assuming, for the sake of argument, that their accounts are satisfactory at least when targeted at a specific subset of cases of self-knowledge.

A couple of examples will suffice to make the point vivid. Consider behaviorism or Cassam's inferentialism. Clearly, inference to the best explanation is the obvious explanation of our knowledge of our Freudian unconscious mental states,[10] or of those mental states, like Emma's love for Mr. Knightly, which may be revealed to us by paying careful attention to our reactions and behavior. Indeed, it is the obvious candidate explanation for a lot of knowledge we may gain of our deep-seated dispositions. For instance, one may realize one is somewhat biased toward students of a certain gender, or race, or attitude by reflecting on one's performance to date in grading, or in selecting students from a pool of applicants, etc. But it is not plausibly extended to knowledge of our occurrent sensation of pain or of our deliberation to take our son out for lunch tomorrow.

Now consider Moran's deliberative account of self-knowledge, according to which it is only when we deliberate what to believe that we have genuinely first-personal self-knowledge. For, on his account, "I believe that" is merely appended to the content P one has judged to be the case, after weighing reasons for and against it. I take Moran's is, at least *prima facie*, a plausible explanation of first-personal self-knowledge of (a certain kind of) beliefs.[11] Yet, to hold, as he does, that in all other cases, there is nothing distinctively first-personal, not even when we

know our occurrent sensations or basic emotions, and that all third-personal cases of self-knowledge would reveal some kind of alienation from ourselves, is clearly an overstatement.

It seems to go against the grain to say that when I self-ascribe an occurrent excruciating pain, I am not doing that in a way which is distinctively first-personal just because it is not an instance of the deliberative model Moran is keen to support.[12] Equally, it is implausible to hold that all cases of third-personal self-knowledge would reveal a kind of psychological malfunctioning on our part. Surely not all unconscious mental states and mental dispositions we have are noble (nor despicable for that matter), but they are part of our psychological makeup and indeed determine our personalities and, in some sense, who we are. They are all, by definition, alien to us, in the sense that we are very often unaware of them. But that does not mean that third-personal self-knowledge—which is the only kind of knowledge we can have of them—is indicative of some psychological malfunctioning on our part. Indeed, given the very nature of these mental states, we can only know them by means of third-personal methods. And it would be preposterous to hold that well-functioning human beings could somehow know them in a first-personal way—that is, as a result of a deliberation! I cannot deliberate not to be (dispositionally) jealous of my more successful sister. I can only realize I am so jealous, conclude that there is no basis for such jealousy and try and behave differently in the future, in hope to be able to live up to my newly formed commitment. That is to say, in hope to be able to bring my dispositions in line with my (newly formed) commitments, and, characteristically, the attempt may be thwarted or successful only to some degree.

Let us take stock. We have seen that both deniers and chauvinists with respect to the asymmetry between first- and third-personal self-knowledge are ultimately united in embracing monism. We have seen that, quite independently of the details of their respective proposals, it is adherence to monism that eventually leads all of them into trouble. This, therefore, paves the way to a liberating proposal: to forsake monism and embrace pluralism with respect to self-knowledge.[13] That move will allow us to offer a comprehensive and hopefully correct account of self-knowledge as a whole. Thus, it will have certain "first-order" benefits.

But it will also have "second-order" —or meta-epistemological—payoffs. For it will allow us to throw into sharp relief the eventual merits of various positions, once taken to apply within their proper boundaries. That is, as is often the case, when one opens up to pluralism, several divergent accounts will no longer be seen as in competition with one another, but rather as valuable proposals with respect to a proper subset of phenomena under investigation. To put it yet another way, several *prima facie* competing accounts of self-knowledge can actually be reconciled with one another, once one adopts a pluralistic approach, which allows one to recognize their respective merits, once these accounts are taken to apply only to some, appropriately specified cases of self-knowledge. This does not mean that, once we open up to pluralism, anything goes. There is still considerable room for debate regarding the correct account of a properly specified class of instances of self-knowledge. For example, one may still sensibly debate whether constitutivism or Moran's deliberative account offer the correct explanation of first-personal self-knowledge of belief. However, what one cannot argue for, once pluralism about self-knowledge is embraced, is that, given the numerous failures at self-knowledge cognitive science has made us aware of, both constitutivism and Moran's deliberative account are clearly off the track. That is, one can acknowledge that they are unsuitable as accounts of our knowledge of our dispositional mental states, but that does not make them *ipso facto* incorrect as accounts of other specimens of self-knowledge. If they are incorrect, they are so only if it can be shown that they do badly at explaining the kind of self-knowledge to which they are actually applicable—that is, first-personal self-knowledge of some suitably specified class of first-order mental states.

6   Third-Personal Self-Knowledge: Methodological Pluralism

There are a variety of methods we employ to gain third-personal self-knowledge. Inference to the best explanation is only one of them, even though—to the best of my knowledge —it is the one theorists

of self-knowledge mostly focus on. Moreover, although it is often invoked, it is not fully understood either. For it is seldom noticed that the kind of "theory" we use in order to provide an explanation of overt behavior may vary considerably, from one case to the other. It may be a "scientific" one like the Freudian conception of the unconscious or a folk psychological one, we may have acquired in a variety of ways. Furthermore, it is rarely remarked that, in self-knowledge and self-knowledge only, the kind of data the theory aims to provide an explanation of are not only instances of overt behavior, but also a lot of inner promptings, like feelings, emotions, and so forth.[14] The quote from Jane Austen's *Emma* we introduced previously helps us see this point. As the reader will recall, Emma infers to her being in love with Mr. Knightley by attending to her own feelings and emotions at the prospect that he might reciprocate another woman's feelings. What that entails, as we remarked, is that third-personal self-knowledge, when provided by inference to the best explanation, often presupposes first-personal self-knowledge. This, in turn, shows that there is often an asymmetric dependence of third-personal self-knowledge on first-personal self-knowledge.

Another method we may use to obtain third-personal self-knowledge is simulation. Suppose I want to figure out whether I am brave or coward. I can proceed as follows. I could imagine being in a potentially scary situation and attend to how I would react. For instance, I could think of being alone, at night, in a house in the middle of nowhere and hear a creeping sound coming from the basement. I could then immerse myself deeper into the simulated scene and see what kind of reaction I would have. That is, I would attend to my feelings and behavior, while simulating, and conclude that I would be brave (or otherwise). This method too would quite clearly involve first-personal self-knowledge of one's feelings in the context of the simulation. Moreover, it would involve actively bringing one's imagination to bear on the task of determining whether one is brave or not. That involves a practical ability—that is, knowing *how* to imagine being in the simulated context—as well as knowing *that* what one is thereby experiencing is indeed one's own imagination of being in a context thus and so. Again,

third-personal self-knowledge obtained through simulation presupposes first-personal self-knowledge.

Another possible method to gain third-personal self-knowledge goes via the identification with the character of a novel or a movie. Suppose you read Austen's *Emma* and that you too have a close friend like Mr. Knightley. Suppose further that, up to that point, you have thought of that relationship just as friendship. Yet, by reading the novel and immersing yourself in it, you are struck by Emma's realization of being in love with Mr. Knightley and that prompts you to engage in similar speculations, which may reveal to you your actual feelings toward your friend (assuming, for the sake of argument, that your feelings are deeper than just friendship). This method seems to resemble simulation (and to involve inference to the best explanation), but certainly its application requires something over and above simulation. Namely, an identification with the character of the novel, in a way in which ordinary simulation does not. That is to say, it requires a projection of oneself into the overall situation depicted and, in particular, a projection of oneself into someone else. Such an identification is usually considered a characteristic trait of empathy, rather than of straightforward simulation. Hence, one way we can gain third-personal self-knowledge goes through empathizing with characters in fictional contexts.[15] Clearly, however, it could also depend on empathizing with real people to the point of realizing that what is true of them, *vis-à-vis* their dispositional psychological properties, may well be true of you. Notice, once again, that once empathy is operative, in the way of a projection of oneself into someone else's situation, it could also give rise to specific feelings and emotions, which one would then know in a first-personal way. These, in turn, could be used as data to make sense of oneself through inference to the best explanation; or else, just be the very objects of one's first-personal self-knowledge.

It is important to note that if all the previous discussion is broadly on the right track, it allows us to bring several debates to date to bear on the topic of self-knowledge in a radically new way. The long-lasting debate over theory-theory can now be reinterpreted not just (or even so much) as a debate about the acquisition of psychological concepts, in particular of propositional attitudes, with the attendant implausible

view that the relevant psychological self-ascriptions would always be inferential,[16] but as showing one way in which we do gain third-personal self-knowledge.[17] Its main competitor, the simulation model,[18] can now be seen as one of the methods we can use to gain third-personal self-knowledge too, rather than just a way of gaining knowledge of other people's minds or, implausibly, as a way of gaining first-personal self-knowledge. The two models would not be in opposition to one another, as accounts of third-personal self-knowledge, though. Rather, they can now be seen as complementing each other in a profitable way. Finally, empathy too can now be seen as one further method that allows us to gain third-personal self-knowledge, rather than simply a way of acquiring knowledge of other people's minds; or else, as just a means of bringing about feelings and emotions in us, which we would then know in a first-personal way.

Another method we may use to gain third-personal self-knowledge is induction. By reflecting on our performance to date, but also on how we have felt in potentially dangerous situations we have experienced in the past, we may come to a conclusion about such dispositional psychological properties like being brave (or cowards). Or again, by reflecting on our performance to date in grading and selecting candidates for a given position, we may realize we are (likely to be) biased (regarding gender, or race, or whatever have you).[19] Once more, only in one's own case, can induction be based on data regarding our inner phenomenology. Furthermore, the evidence on which the induction is based can take the form of past psychological self-ascriptions. Thus, once again, we can see how third-personal self-knowledge, while utilizing methods we could apply to gain knowledge of other people's mental states, is unique in allowing for psychological inputs, which are known in a first-personal way.

Two further methods that, so far, have received no attention, but which are indeed crucial to self-knowledge, are testimony and what I call, following Miranda Fricker, "hermeneutics." The former provides us with what might be called *second-personal* self-knowledge. Suppose you talk to a person you trust and she tells you that you are intimidating. Trusting that person, having no reason to believe she may be deceiving you or be incompetent with respect to the topic at issue, you may then form a justified and, for the sake of argument, true belief that you

are an intimidating person.[20] You would therefore acquire self-knowledge through the interaction with another person and, by relying on her judgment, you would then gain testimonial justification for your psychological self-ascription.

Notice that, on a notion of objectivity understood as inter-subjective validity, this is actually the most objective method one may employ, in order to gain third-personal self-knowledge. For it constitutively depends on the interaction with others and on taking them to their word. Furthermore, it may (even though it need not) disclose to us psychological dispositional properties, which constitutively depend on the interaction with other people, and that may well lack any inner phenomenology. Consider being intimidating, one is only intimidating because other people feel scared, or threatened, or challenged when dealing with that person. Did they not so feel, then one would never be intimidating, no matter how hard one tried or felt so.[21] Thus, through testimony, we can actually gain knowledge of some psychological properties of ours we would otherwise be precluded from knowing.

Testimony, furthermore, is at work whenever we acquire self-knowledge through the interaction with a psychoanalyst or a psychiatrist—that is, with an expert. We need to trust their judgment in order to acquire knowledge about our own minds—that is, we must have no reason to think they may want to deceive us, as well as no reason to think they may be incompetent in their profession. While the clinical context introduces complexities of its own,[22] it is important to stress the role of the interaction with others—mostly family members, partners, and friends—as an important source of self-knowledge. I would certainly not consider it the sole or the main value of having close relationships with other human beings, but it is certainly one of the payoffs of having meaningful interpersonal relations, which can actually help us make better sense of ourselves and give us knowledge and motivation to try to improve our own character or to lead a better life.

Finally, hermeneutics consists in bringing new concepts to bear onto one's own behavior and inner promptings, in such a way as to afford a new piece of self-knowledge. As I conceive of it, it can work in two rather different ways. First, it can work abductively, by allowing for redescriptions of already observed overt behavior and inner promptings.

Suppose you acquire a new concept, like *egotistical*. You can then use it to re-conceptualize your already observed behavior and inner promptings. This way, the newly acquired concept can cause a "switch of aspects." Up to a certain moment, you thought of yourself as just a bit self-centered and now you see yourself as egotistical. These aspects may alternate, as you weigh more one or the other feature of your behavior and inner promptings to date. The alternation would thus give rise to the switch of aspects characteristic of all instances of seeing something as.[23]

Hermeneutics, however, can also work in a different way. That is, by going through the characteristic notes of a concept and by taking oneself to instantiate them. Thus, instead of reasoning like "If you are egotistical, you display such-and-such a behavior; I display such-and-such a behavior; therefore, I am egotistical," you could reason as follows. "I exemplify $trait_1$, I exemplify $trait_2$…I exemplify $trait_n$. Since $traits_{1-n}$ are the characteristic traits of being egotistical, I am egotistical."[24]

Be that as it may, the important point is that newly acquired concepts can make a dramatic difference to self-knowledge. For the more concepts we have, the finer-grained distinctions we can make and the better understanding of our psychological dispositional properties we can achieve. Again, I would not consider this the only or the main value of pursuing conceptual competence and sophistication, but it is certainly one of its payoffs, as it can reveal important aspects of our personalities. Once more, the kind of evidence on which hermeneutics works, in one or the other of the two ways described, may often comprise inner promptings and past psychological self-ascriptions. It therefore heavily relies on first-personal self-knowledge.

This list of methods is not meant to be exhaustive, but it makes clear how diverse our routes to third-personal self-knowledge may be and how likely they are to depend (at least in many cases) on first-personal self-knowledge. Not only are they many and diverse, but, in real-life scenarios, they can combine with one another in several ways, leading to a high degree of methodological complexity. Thus, third-personal self-knowledge ain't easy!

Yet, it isn't epistemologically special either. For, first, none of these methods is sure-fired. Second, they all involve substantial cognitive and

epistemic work on our side. Finally, we may fail to apply these methods and thus end up in states of self-blindness, which in turn can lead to mistaken psychological self-ascriptions.[25] That is, no matter how it is achieved, third-personal self-knowledge is never groundless, transparent, or authoritative.

7   First-Personal Self-Knowledge: Property Pluralism

As already hinted at, I embrace a form of constitutivism regarding first-personal self-knowledge. At the heart of all constitutive positions lies the following schematic Constitutive Thesis (CT):

> Given *C*, *S* is in *M* iff *S* judges "I am in *M*" (where *C* ranges over a set of appropriately specified conditions, *S* is a subject and *M* ranges over an appropriately specified set of mental states).

The important point is that for all constitutive theorists the Constitutive Thesis holds a priori. They do therefore hold that what goes by the name of first-personal self-knowledge is not the result of any epistemic achievement. The overall idea is that it is constitutive of being a subject, who fulfills certain conditions *C*, and who is capable of enjoying the relevant kinds of mental state *M*, that she will judge that she is in *M* if and only if she is.

As I have claimed elsewhere,[26] the conditions that have to be fulfilled in order to be granted first-personal self-knowledge vary considerably regarding the kind of mental states at issue. What remains constant are the C-conditions, which can be generally specified as comprising the possession of the relevant concepts—that is, the first-person concept, the relevant psychological concepts,[27] and those concepts needed to specify the intentional content of one's mental states when they have one. That immediately excludes animals and infants from the range of suitable subjects for which the Constitutive Thesis is supposed to hold. For either they won't be able to enjoy at least some relevant class of

mental states; or else, while capable of enjoying them, they will not have the resources for making the corresponding self-ascriptions. Hence, they would actually fall out of the range of subjects capable of self-knowledge (no matter whether first- or third-personal). Furthermore, a subject will have to be lucid, attentive, and alert. While it may be tricky fully to specify this condition, we can clearly exclude cases in which a subject is under the effect of substances, like drugs and alcohol, on in certain environmental conditions, such as thin air and lack of oxygen, which would alter her reactions and attention, as well as her ability to apply the relevant concepts. Similarly, we can exclude cases in which the subject is under extreme fatigue or emotional distress, which would inhibit her proper functioning. Finally, we will have to impose a sincerity condition.[28] For in order for one's self-ascription to be an expression of first-personal self-knowledge—thus enjoying a kind of transparency and authority granted by default—a subject will have to take her share of responsibility and not deliberately misguide people around her.

The idea, then, is that depending on the kind of mental state at issue, the explanation of the Constitutive Thesis will vary and will bring to light features that can be considered to be constitutive of what it is to be a subject capable of enjoying each of these mental states. In particular, the characteristic groundlessness, transparency, and authority of the relevant psychological self-ascriptions can be seen as the fall out of having different, partly overlapping sets of properties that underlie the appropriate abilities.

Thus, for instance, when we consider basic sensations, such as purely phenomenal pain, cold, heat, and pleasure, or basic emotions,[29] like fear, joy, or anger, we can say that it is constitutive of being a subject, who, besides satisfying the above-mentioned C-conditions, is capable of enjoying those sensations and basic emotions,[30] of experiencing them as one's own,[31] and is rational,[32] that if one enjoys one of these mental states, one will judge that one does and if one so judges one will have them.[33]

Regarding perceptions, the Constitutive Thesis can be redeemed by noticing how it is constitutive of being a subject who satisfies the relevant C-conditions (further specified as excluding cases of massive deception, like in skeptical scenarios), is capable of enjoying them and

of experiencing them as one's own, and who is rational, that one will judge to have them iff one does have them. Notice, moreover, that since it is possible to have unconscious perceptions,[34] a further condition that will have to be imposed is that we are here dealing with perceptions that can enter an explanation of a subject's outer actions, for which she can legitimately be held responsible.[35]

Finally, with respect to beliefs, desires, intentions, and possibly other propositional attitudes, it is important to acknowledge a crucial distinction between these states as dispositions and as commitments.[36] That is, we are here considering propositional attitudes and intentions that are intrinsically normative[37] and dependent on a subject's conscious deliberation, based on reasons, as to whether P obtains, or would be good to have, or as to whether a given course of action is something to go for. Once such a distinction is taken on board and supplemented with the appropriate account of concepts' acquisition, such that "I believe" (or "I want" or "I intend") are simply taught to replace the expression of P/P would be good to have/Φ is worth pursuing, then there won't be any epistemic work to be done and transparency and authority will be guaranteed to obtain.[38] Hence, it is constitutive of being a subject capable of propositional attitudes and intentions as commitments that one will judge to have them iff one does.[39]

Three observations are in order. First, this kind of pluralism allows one to incorporate suggestions coming from both the expressivist camp and from fans of the so-called transparency method, such as Moran, as useful hints regarding the acquisition and deployment of the relevant psychological concepts, while remaining critical of both the expressivist and the deliberative account as accounts of first-personal self-knowledge.[40] Second, given the expressivist story regarding the acquisition and deployment of the relevant psychological concepts, it is easy to see how, as Sydney Shoemaker has it, the first-order mental state and the corresponding self-ascription "have the same core realization," while the latter "enables the core-realization of the first-order belief to play a more encompassing role."[41] For the self-ascription is just an alternative way of expressing the first-order mental state one is in. Finally, precisely in virtue of the role of the self-ascription, in at least some cases, the very self-ascription can bring about the first-order mental state. That is not

the case with sensations, perceptions, and emotions, but it may be the case with the self-ascription of propositional attitudes and intentions as commitments. That is, we sometimes deliberate what to do and thus form an intention (as a commitment) to Φ, say, by judging "I intend to Φ." It is only in these latter cases that constitutivism can be seen as endorsing a robust metaphysical attitude with respect to the fact that first-order mental states can be constituted through or by their very self-ascription.[42]

8   Conclusion

Where does all this leave us with respect to pluralism in general? A good way of approaching this issue is by considering a kind of criticism often raised against constitutive positions. Namely, that they do not explain what they are supposed to explain. That is, how we can indeed *know* our own mental states, in such a way that the resulting psychological self-ascriptions will have the features traditionally associated with first-personal self-knowledge—i.e., groundlessness, transparency, and authority.

Notice, however, that this quite radical objection depends, in its turn, on an implicit assumption. Namely, that if something is called "knowledge," it must pick out an epistemic property—roughly, the property of being a justified true belief. On this view, the term "knowledge" picks out one (and only one) epistemic kind whose exact nature is for epistemologists (or even for metaphysicians of epistemology) to get clear about.

Indeed, in the grip of monism, if one agreed that first-personal self-knowledge exists, one would be hard pressed to show how an epistemic method could actually guarantee groundlessness, transparency, and authority.[43] A good antidote against such a dogmatic position, besides realizing the problems one would incur by trying to abide by it, is to remind ourselves of the fact that the word "knowledge" is actually used in a plurality of ways, where there is no real expectation that it should always pick out or refer to the same epistemic property. For instance, we say that Jane knows how to play the violin or that Sally knows Jim, where in both cases the onus is entirely on the monist's shoulders to

show that knowledge-how and knowledge-who (or which) could actually be reduced to knowledge-that. For, at least *prima facie*, it seems quite plausible to say that knowledge-how actually consists in having an ability and that knowledge-who (or which) ultimately consists in being acquainted with the relevant individual in such a way as to be able to single it out from other ones in, for instance, a given perceptual scene. Hence, a monist would be hard pressed to conclusively show that either kind of knowledge is propositional in nature.[44]

Of course, one might acknowledge these differences and yet insist that, in the particular case at hand, we are not really concerned with knowledge-how or knowledge-who (or which). Rather, we are talking about propositions of the form "I am in M" (where *M* ranges over appropriately specified mental states). Thus, we would need to explain how it is possible for a subject to arrive at such knowledgeable self-ascriptions.

Luckily, there is yet another sense of "knowledge" we may bring to bear on the case at hand. Namely, the one in which we use "I know," for instance, to signal the fact that a doubt is excluded. Consider the following scenario: A subject says "I have a headache." An interlocutor challenges her by saying "Really? Are you sure?" to which the subject may reply by saying "Of course." And if the impertinent interlocutor insisted "How do you know?" the subject would be entitled to respond "What do you mean? Of course I know; I have it (or: I feel it/I'm experiencing it)." Now, since the latter part of her reply—viz. "I have it/feel it/am experiencing it"—would just be a restatement of the original judgment, the subject would not really be engaging in the game of providing some kind of epistemic backing for her initial claim. To that end, at the very least, one would be required to produce some independent corroborating reason in support of one's claim "I have a headache." Just repeating that one does have a headache won't do.[45] The "I know," therefore, is meant simply to signal the fact that the question raised by the interlocutor is bordering nonsense. For we take it for granted that if a subject says she has a headache, in what appear to be entirely normal conditions, she has it.

Wittgenstein coined the term "grammatical" for these uses of "I know."[46] No matter how you call them, the point is that they are not

meant to express the obtaining of a specific, possibly peculiar epistemic relation between a subject and a proposition. Rather, they are meant to signal the fact that it is part of a linguistic practice, or even of a certain *Weltbild*, or—more humbly—of a way of organizing experience, that certain things are taken for granted. In a similar fashion, according to constitutivists, as we have seen, it is actually constitutive of being a subject capable of enjoying the relevant mental states, under suitably specified conditions, that one judges to enjoy one of these mental states iff one does have them.[47]

Hence, the contentious, no doubt, yet liberating[48] claim at the heart of constitutive positions is that, when we talk about first-personal self-knowledge, we are talking about certain features we take the relevant psychological self-ascriptions to possess *by default*, in virtue of being made by subjects who satisfy certain specifiable conditions. Still, we are not talking about a very peculiar, as appears to be, epistemic relationship,[49] which obtains between a subject and either the mental states she would be having or, more mildly, a certain proposition describing them.

To conclude, "self-knowledge" is actually an expression that can disguise a plurality of states or properties a subject may be in. When we are confronted with third-personal self-knowledge, it picks out a genuine epistemic relation holding between a subject and a proposition, underwritten by a plurality of methods, as detailed in the previous section. When, in contrast, we are considering first-personal self-knowledge, "self-knowledge" does not refer to any special epistemic relation obtaining between a subject and a proposition. Rather, it refers to a set of features the relevant self-ascriptions are granted to have,[50] once certain further conditions obtain.[51] It is indeed a feature of the account of first-personal knowledge I have defended here (and elsewhere more thoroughly) that the characteristic traits of the relevant self-ascriptions are, in their turn, seen as the fall out of the fact that they are made by a subject who has a set of further specifiable, partly overlapping properties that allow for the exercise of a complex set of abilities. This is not to say that the term "knowledge" in the locution "self-knowledge" is ambiguous, in the same way in which the term "bank" is. Rather, we may say that it is the name of a family resemblance concept. The point of resemblance between the fully epistemic sense and the grammatical

one would be to signal a certain kind of standing of the subject and her claim. In the epistemic sense, the subject is flagged as a good source of information regarding P, given her way of having reached the relevant belief.[52] In the non-epistemic sense, the subject is still flagged as a good source of information, because of her fulfillment of other conditions, which allow her authority over her own psychological self-ascriptions.

Notes

1. See Coliva (2016). Due to space limitations, I will often refer the reader to that work for the details of my first-order epistemological views.
2. This is not the place to defend these first-order epistemological claims in detail. I have done so in Coliva (2016, Chaps. 3, 7–8). In Coliva (2017a, *forthcoming*, see in particular Sect. 2), I criticize David Chalmers' (2003) attempt to vindicate what may be regarded as a Russellian indexical account of how occurrent sensations may justify their self-ascriptions.
3. I defend this view in Coliva (2016, especially Chaps. 7–8). See also Bar-On (2004).
4. See Ryle (1949, Chap. 6).
5. I present my "border-line" account of emotions in Coliva (2016, Chap. 2).
6. See also Gertler (2011, pp. 70–86).
7. See Cassam (2014).
8. Cassam (2014) proposes this kind of conversational reorientation, as it were. For a critique, see Coliva (2015).
9. See Moran (2001), Bilgrami (2006), and in general those philosophers who have solely concentrated on first-personal self-knowledge. The usual route to such an idea is that if one agrees that first-personal self-knowledge is characterized by groundlessness, transparency, and authority, self-knowledge has to be underwritten by an entirely *sui generis* kind of epistemology.
10. But even that case is often more complex than what philosophers generally make of it. After all, people usually get to know these mental states through psychoanalytic counseling, which involves interaction with a therapist. That in turn entails that they rely on what their therapist tells them, in order to get knowledge of their unconscious.

So testimony (and trust in the expert) is actually the base for a lot of knowledge we may gain regarding our unconscious mental states.

11. While I think there is much to applaud in Moran's discussion of first-personal self-knowledge, I am not convinced he actually gives the correct account of it, but that is not relevant for present purposes. For a discussion of the details of Moran's position, see Coliva (2016, pp. 122–128).

12. Boyle (2009) glosses Moran as holding the milder view that deliberative first-personal self-knowledge is more fundamental than other possible instances of first-personal self-knowledge, like our first-personal knowledge of our ongoing pain or further sensations. Boyle also defends that view, for he thinks the deliberative account is tied to our notion of belief. I do not take issue with that. Still, if that were the motivation, I would argue that self-knowledge of our ongoing pain, for instance, is fundamental too, because it is tied to our notion of pain. Indeed, it is constitutive of having that concept that whenever in pain, and if further suitable conditions obtain, one self-ascribes it and, conversely, that if one does self-ascribe pain, in the relevant conditions, then one has it.

13. This is just one more instance in which Wittgenstein's *memento*, in *Philosophical Investigations* (PI, 66), "Don't think, but look!", finds an illuminating application. As is well known, Wittgenstein thought of having *King Lear*'s verse "I'll teach you differences" as a *motto* for his book.

14. Apart from Coliva (2016), the only other theorist who makes the point explicitly is Cassam (2014).

15. This shows one (but by no means the only one!) dimension along which literature (and other artistic forms) can have a cognitive value and can actually teach us something.

16. See Gopnik (1993) and Cassam (2014).

17. I take up the issue of how we may acquire the relevant psychological concepts in Coliva (2016, pp. 188–197).

18. See, in particular, Goldman (1993) and Gordon (1995). For a discussion of its bearing on self-knowledge, see Coliva (2016, pp. 88–95).

19. The self-ascription of such a bias would clearly involve also inference to the best explanation. For induction, in and of itself, would simply allow one to predict that one will keep behaving as one has done in the past and would not provide an explanation of one's performance to date.

20. This account is meant to be neutral with respect to any epistemological account of testimonial knowledge and justification.
21. In fact, I think it would make little sense to say "I feel intimidating" (as opposed to "I feel intimidated"). Be that as it may, being intimidating is a property which constitutively depends on there being the relevant kind of interaction with other subjects.
22. As is well known, a patient may actually resist the diagnosis and therefore irrationally distrust the expert's veracity and/or competence.
23. This method can actually be at the origin of some mistaken psychological self-ascriptions, which are usually used to illustrate the phenomenon of self-deception. For if a given concept is particularly thick or morally loaded, a subject's reverting back to a thinner, less morally loaded one, can explain her eventual (mistaken) self-ascription. Indeed, it would also offer an explanation of why self-deception appears to be motivated. However, I do not think such a story can apply to all cases usually presented in the literature to illustrate the phenomenon of self-deception. For a discussion of cases that are better accounted otherwise, see Coliva (2016, pp. 197–200).
24. Of course, these forms of reasoning will often be implicit. Yet, making them explicit allows us to appreciate their difference at a theoretical level.
25. Again, this interplay between self-blindness and further psychological self-ascriptions can explain at least some cases usually appealed to in the literature to illustrate the phenomenon of self-deception. For, if you are blind to some of your dispositional mental states, you can actually deny having them, when it would be evident to a third party that you do have them. They would then have good reasons to think that your psychological self-ascription is mistaken.
26. Coliva (2016, Chaps. 6–7).
27. The story we tell about the possession of the relevant psychological concepts is important and it has to be such as to avoid falling back into either the Cartesian or the behaviorist/inferentialist trap. In the first case, subjects would be supposed to have a given mental state in view, attend to it and learn how to name it, or how to single it out in thought. In the latter case, subjects would have to infer from their behavior to its likely psychological cause, by application of a little folk psychological theory they should have already acquired. For a criticism of these views, see Coliva (2016, pp. 52–58, 84–88) and Coliva

(2017a, forthcoming). The expressivist story, proposed by Wittgenstein in *Philosophical Investigations* (1953) and further developed by Bar-On (2004) is arguably more promising. As to the concept of belief, the ascent routines presented by Evans (1982), Gordon (2007), and Moran (2001) are in fact entirely compatible with the expressivist account, as proposed in Bar-On (2004) and Coliva (2016, pp. 188–197).

28. Importantly, and *contra* Wright (1989), we don't have to exclude self-deception.

29. The main difference between basic sensations and basic emotions has only to do with the fact that emotions often have a relational object, besides an intentional one. I do allow for the possibility of error in the individuation of their relational object. Hence, I take the Constitutive Thesis to apply only to the very individuation of the kind of basic emotion one is undergoing and to its possible intentional content.

30. Contrary to subjects like the ones affected by congenital insensitivity to pain with anhidrosis (CIPA).

31. Contrary to subjects who, while feeling pain, would know of it only through inference to the best explanation, based on the observation of their own behavior. These are 'self-blind' subjects in Shoemaker's (1996) sense.

32. Contrary to subjects who had the relevant concepts, satisfied the other features of the C-conditions, were capable of experiencing them and of experiencing them as one's own, and yet were either uncertain about whether they would be having the relevant mental states or made the relevant self-ascriptions and yet behaved in ways which would run totally contrary to them.

33. For the details, see Coliva (2016, pp. 222–231).

34. As exemplified by cases of blind-sight.

35. Also called "conscious perceptions" for short. While I have no qualms with calling them thus for simplicity, I think it is important, within a comprehensive constitutivist project, to bring out exactly the kind of ties conscious perceptions have with outer actions. Accordingly, the Constitutive Thesis will be said to hold, with respect to perceptions, only for subjects who can be considered to be rationally responsible for their outer actions. For further discussion of this point, see Coliva (2016, pp. 231–232).

36. For the details of such a distinction see Coliva (2016, pp. 26–38).

37. In Coliva (2015), reprinted in Coliva (2016), I claim that intrinsic normativity is what distinguishes beliefs as commitments from beliefs as dispositions. The idea is that a belief as a commitment is intrinsically normative iff it would be impossible to have it while also being open-minded with respect to its content, or while assenting to the negation of the latter.

38. I have developed the details of this part of the story in Coliva (2016, pp. 188–193).

39. Usual counter-examples to authority coming from cases of self-deception are therefore handled by saying that they actually manifest a conflict between one's propositional attitudes as commitments, on which one retains authority, and as dispositions regarding which one is not authoritative. For the details, see Coliva (2016, pp. 197–200).

40. For a criticism of expressivism as a theory of self-knowledge and not as an account of how we acquire and canonically deploy several psychological concepts, see Coliva (2016, Chap. 6). For a criticism of Moran's deliberative account of self-knowledge, see Coliva (2016, pp. 122–128).

41. Shoemaker (1996, pp. 243–244).

42. A similar account holds for self-verifying self-ascriptions too, like "I am hereby thinking that P." In the case of sensations, basic emotions, and perceptions, in contrast, the self-ascription does not bring about the first-order mental state, but it individuates it for what it is. This view is less metaphysically robust than the one I presented for mental states as commitments, but it still stands opposed to those accounts according to which it is the intrinsic quale of the very first-order mental state that individuates it for what it is. Such a metaphysically less robust position also makes constitutivism more plausible, as it would not entail that one's self-ascriptions of sensations, basic emotions, and perceptions bring about the corresponding first-order mental state.

43. In Coliva (2016, Chaps. 4–6), I have examined the prominent epistemic accounts of first-personal self-knowledge and showed how they all fail to account for the constitutive features of first-personal self-knowledge.

44. As is well known, Stanley and Williamson (2001) have maintained that knowledge-how is a subspecies of knowledge-that. I myself am skeptical of their proposal (see Coliva 2017b, forthcoming). An important treatment of knowledge-which can be found in Evans (1982). In Coliva and Sacchi (2001), the conditions that have to be satisfied to individuate

objects in a perceptual scene are explored further. They do depend on perceptual discrimination, which, in turn, does not depend on the possession of the relevant concepts. To the extent that concepts are necessary to grasp a proposition, therefore, knowledge-which cannot be considered a form of propositional knowledge.

45. For the details of this diagnosis, see Coliva (2016, pp. 54–55) and Coliva (2017a, forthcoming).
46. Wittgenstein (1969, §58).
47. A historical note: It is often remarked that Wittgenstein's ideas in *On Certainty* on the role of "I know" in relation to Moore's truisms were elicited by his conversations with Norman Malcolm in Ithaca in 1949. That is correct as far as it goes. But actually Malcolm was applying Wittgenstein's views as the latter had developed them, still in opposition to Moore, in the 1930s, in relation to self-ascriptions of pain. For the details of this historical reconstruction, see Coliva (2010, pp. 29–30).
48. Liberating because, arguably, none of the extant epistemological theories of first-personal self-knowledge could really vindicate groundlessness, transparency, and authority and do so for all psychological self-ascriptions we have reason to consider to be genuine manifestations of that kind of self-knowledge. The most prominent epistemic accounts are examined and criticized in Coliva (2016, Chaps. 4–6).
49. If the relation were epistemic, it would be peculiar indeed as it would have traits which do not characterize any other kind of empirical knowledge we have—let it be of physical objects in our surroundings or of other people's mental states.
50. Namely, groundlessness, transparency, and authority.
51. It is important to note that also supporters of different accounts of first-personal self-knowledge would end up embracing a pluralism of properties or states. In particular, Burge (1996, 2011), Peacocke (1999, 2003, 2014), and Moran (2001, 2003), with their—albeit different—notion of entitlement would have to say that while in third-personal self-knowledge "knowledge" picks out an epistemic property which depends on having a discursive justification for one's psychological self-ascription, in first-personal self-knowledge, "knowledge" picks out an epistemic property which does not depend on having such a discursive justification.

52. This is entirely in keeping with Edward Craig's (1990, p. 11) well known claim that "the concept of knowledge is used to flag approved sources of information." Still, the reasons behind such flagging may be entirely different when first-personal self-knowledge and third-personal self-knowledge are concerned.

References

Bar-On, D. 2004. *Speaking My Mind*. Oxford: Oxford University Press.

Bilgrami, A. 2006. *Self-knowledge and Resentment*. Cambridge: Harvard University Press.

Boyle, M. 2009. Two Kinds of Self-Knowledge. *Philosophy and Phenomenological Research* 77 (1): 133–164.

Burge, T. 1996. Our Entitlement to Self-knowledge. *Proceedings of the Aristotelian Society* 96: 91–116.

———. 2011. Self and Self-understanding: The Dewey Lectures (2007-2011). In *Cognition through Understanding*. Oxford: Oxford University Press.

Cassam, Q. 2014. *Self-Knowledge for Humans*. Oxford: Oxford University Press.

Chalmers, D. 2003. The Content and Epistemology of Phenomenal Belief. In *Consciousness: New Philosophical Perspectives*, eds. Q. Smith, and A. Jokic, 220–272. Oxford: Oxford University Press.

Coliva, A. 2010. *Moore and Wittgenste in Scepticism, Certainty and Common Sense*. London: Palgrave.

———. 2015. How to Commit Moore's Paradox. *The Journal of Philosophy* 112 (4): 169–192.

———. 2016. *The Varieties of Self-Knowledge*. London: Palgrave.

———. 2017a. Wittgenstein, Self-Knowledge and Nature. In *Wittgenstein and Naturalism*, eds. K. Cahill and T. Raileigh. London and New York: Routledge.

———. 2017b. Practical and Linguistic Competence. In *Language, Language and Language: Essays in Honor of Eva Picardi*, eds. A. Coliva and P. Leonardi. Berlin: Springer, forthcoming.

Coliva, A., and E. Sacchi. 2001. *Singular Thoughts: Perceptual-Demonstrative and I-Thoughts*. Macerata: Quodlibet.

Craig, E. 1990. *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford: Oxford University Press.

Evans, G. 1982. *The Varieties of Reference*. Oxford: Clarendon Press.

Gertler, B. 2011. *Self-Knowledge*. London and New York: Routledge.

Goldman, A. 1993. The Psychology of Folk Psychology. *Behavioral and Brain Sciences* 16: 15–28.

Gopnik, A. 1993. How We Know Our Minds: The Illusion of First-person of Intentionality. *Behavioral and Brain Sciences* 16: 1–15.

Gordon, R. 1995. Simulation Without Introspection or Inference From Me to You. In *Mental Simulation: Evaluations and Applications*, eds. M. Davies, and T. Stone, 53–67. Oxford: Blackwell.

Gordon, R. 2007. Ascent Routines for Propositional Attitudes. *Synthese* 159: 151–165.

Moran, R. 2001. *Authority and Estrangement*. Princeton: Princeton University Press.

———. 2003. Responses to O'Brien and Shoemaker. *European Journal of Philosophy* 5 (2): 141–161.

Peacocke, C. 1999. *Being Known*. Oxford: Clarendon Press.

———. 2003. *The Realm of Reason*. Oxford: Clarendon Press.

———. 2014. *The Mirror of the World: Subjects, Consciousness, and Self-Consciousness*. Oxford: Oxford University Press.

Ryle, G. 1949. *The Concept of Mind*. Chicago: The University of Chicago Press.

Schwitzgebel, E. 2008. The Unreliability of Naïve Introspection. *Philosophical Review* 117: 245–273.

Shoemaker, S. 1996. *The First Person Perspective and Other Essays*. Cambridge: Cambridge University Press.

Stanley, J., and T. Williamson. 2001. Knowing How. *The Journal of Philosophy* 98 (8): 411–444.

Wittgenstein, L. 1953. *Philosophical Investigations*. Oxford: Blackwell.

———. 1969. *On Certainty*. Oxford: Blackwell.

Wright, C. 1989. Wittgenstein's Later Philosophy of Mind: Sensation, Privacy and Intention. *The Journal of Philosophy* 89 (11): 622–634.

How to Be a Pluralist About Disagreement

Michele Palmira

1   Introduction and Overview

The past decade has witnessed a surge of interest in the phenomenon of disagreement. Debates in the philosophy of language on the semantics of certain expressions, such as "tasty", "knows", "might", take disagreement to be a datum against which semantic theories must be assessed (see e.g. García-Carpintero and Kölbel 2008; MacFarlane 2014; Sundell 2011). By contrast, epistemologists have been focusing on the rational response to disagreements between acknowledged epistemic peers (see e.g. Christensen and Lackey 2013; Feldman and Warfield 2010).

Virtually all participants in current debates about disagreement agree on one issue: disagreement comes in many varieties. To mention but a few authors acknowledging this: Carter (2016), Huvenes (2012), López de Sa (2015), MacFarlane (2014), Marques (2014), Marques

M. Palmira (✉)
Department of Philosophy, University of Barcelona,
Barcelona 08001, Spain
e-mail: michelepalmira@gmail.com

and García-Carpintero ([2014](#)), Sundell ([2011](#)). We are variously told by these authors that disagreement can arise between individuals having incompatible doxastic or conative attitudes. These attitudes can be understood in a full or graded way. Disagreements can be about the current or desired states of affairs, or they can centre on how we should use words and concepts to describe how things are. Furthermore, groups can disagree in all these ways.

However, a closer look at the literature reveals that the consensus on there being various ways whereby disagreement manifests itself betrays a lack of theoretical depth. It is somewhat surprising to register that questions such as what it means that disagreement comes in varieties and what the relation—if any—between such varieties is, have been almost entirely ignored in the literature. I take this to be rather unfortunate, for it is legitimate to ask: Is there any sense in which philosophers talking about different varieties of disagreement manage to capture different aspects of the same phenomenon, or are they simply talking about different things? And why do we use the same term "disagreement" (or the same concept DISAGREEMENT) to refer to some incompatibility relation which can be instantiated no matter whether we are talking about individuals' and groups' beliefs, desires and so on?

I should hasten to emphasise that these questions have a bearing on current debates about disagreement. To make an example: if it turned out that the different varieties of disagreement captured by contextualism (see e.g. López de Sa [2015](#); Marques and García-Carpintero [2014](#)) and relativism (see e.g. MacFarlane [2014](#)) about the predicates of personal taste or knowledge ascriptions do not share any noteworthy feature, we would be bound to conclude that when contextualism and relativism claim that they are equipped to capture disagreement data they are talking about completely different phenomena. If this were the case, the contextualism *vs.* relativism dispute about disagreement should be regarded as merely verbal.

The broad aim of this chapter is to show that we can find a way out of the disagreement jam by opening up the *pluralism toolbox*, as it were. One might indeed think that a pluralist approach to disagreement imposes itself in the light of the varieties of ways whereby this phenomenon manifests itself. Yet, as has emerged from the literature on alethic

and logical pluralism, there are different ways of developing a pluralist theory which carry substantively different implications. For this reason, we should pay careful attention to the specific form a pluralist theory of disagreement has to take.

Since it is impossible to make justice to all the varieties of disagreement in the space of a single contribution, the more specific aim of the paper is to investigate *doxastic* disagreement, thereby deferring a full examination of *conative* disagreement—that is, disagreement involving incompatible conative attitudes such as desires, preferences and the like—to further works. (I shall henceforth take all the examples of disagreement to be doxastic in kind, unless otherwise stated and omit the "doxastic" qualification). However, in closing, I will get back to the question of the nature of conative disagreement, for I believe that the discussion of doxastic disagreement can be taken to teach us something about how to investigate conative disagreement.

The chapter is organised as follows. In Sect. 2, I single out four varieties of disagreement which have taken centre stage in current epistemological and semantic debates. In Sect. 3, I develop and criticise a form of *Disjunctive Pluralism* about disagreement that can be elicited from the works of John MacFarlane. In Sect. 4, I argue for what I shall call *Kinship Pluralism* about disagreement. In Sect. 5, I briefly comment on the possibility of extending Kinship Pluralism to yet other possible varieties of disagreement, such as *group* and *agnostic* disagreement. In Sect. 6, I discuss the wide-ranging implications of Kinship Pluralism.

2   Four Varieties of Disagreement

In this section, I pin down four different varieties of disagreement. I will present them as belonging to two distinct pairs: the descriptive/conceptual pair and the full/credal pair.

2.1   Descriptive/Conceptual Disagreement

Let us consider this case:

(BALD)

Mary says: "John is bald".

Albert says: "Nuh uh, John is not bald".

We can take Mary and Albert to be using the gradable adjective "big" in the same *descriptive* way. That is to say, they use it to convey information about John's hairs. Hence, they disagree about how to describe the state of affairs concerning John's hairs. Call this an instance of *descriptive* disagreement.

Suppose now that Mary and Albert share all the relevant information about John's hairs, e.g. they both know and agree on the exact number of his hairs. In such a context, Mary nonetheless utters "John is bald". In this case, Mary's utterance does not add new information about John's hairs, but it rather conveys what the relevant contextual standard of baldness is and how the adjective should be used in that context. John negates the same sentence uttered by Mary. By doing so, again, he does not add new information about John's hairs, but it rather conveys how the adjective should not be used in that context in the light of what he takes to be the relevant contextual standard of baldness. This use of adjectives and negation has been dubbed *metalinguistic* (Barker 2002).

If we take Mary and Albert to be using the adjective "bald" in a metalinguistic rather than descriptive way, they do not disagree about whether John is bald. That is, their disagreement is not descriptive. Rather, they disagree about how the adjective should be used in that context, thereby engaging themselves into a *metalinguistic negotiation* (see Plunkett and Sundell 2013) about how the adjective should be used in the context.

Since in order for two individuals to be in a state of disagreement, it is not necessary that they linguistically express their attitudes through utterances or inscriptions, and since we can take concepts to be the mental counterparts of words, call disagreements about what entities concepts should refer to—as opposed to what these concepts do refer to—*conceptual* disagreements.

2.2   Full/Credal Disagreement

This chapter deals with disagreements between individuals having incompatible—in a sense which will be specified below—doxastic attitudes. However, one of the most important theoretical divides in epistemology centres on how to conceive of such attitudes. The traditional Cartesian epistemological approach takes the attitude of *belief* to be the central doxastic attitude. Belief bears a privileged relation, both descriptive and normative, to truth. As far as the descriptive side of the relation is concerned, we often read in philosophical papers that *belief aims at truth*. For present purposes, we can rest content with a minimal understanding of this metaphor to the effect that beliefs are actually regulated for truth, that is, one believes that *p* just in case one regards *p* as true for the sake of getting *p*'s truth-value right, as it were. Truth has also been taken to provide the standard of *correctness* for belief. The so-called truth norm of belief is often formulated thusly: For all *p*, one ought to believe *p* only if *p*.

These observations naturally accord with what I shall call the *qualitative model of belief*. According to this model, belief is an all-or-nothing, black-and-white binary representational attitude one might take toward a proposition: either one believes that *p*, or one fails to believe that *p*. A question such as "Do you believe that it's raining in Montreal right now?" and an answer such "Yes I do" are perfectly felicitous. A proper answer to this question does not require specifying to what extent one believes the proposition, nor does the answer suggest the degree to which one believes it. The same holds, *mutatis mutandis*, for disbelief.

Let us call disagreements involving beliefs *full disagreements*. Here is an example of full disagreement:

(SUNNY)

Hichem at 4 pm on February 27th, 2017, believes *that it's sunny in Montreal*.

Sophie, at 5 pm on February 27th, 2017, believes *that it wasn't sunny in Montreal an hour ago*.

However, another model of belief has been embraced by formally inclined epistemologists to the effect that doxastic attitudes come in degrees. The intuition is that I might be more confident in the truth of the proposition *that right now I am in front of a screen* than in the truth of the proposition *that I have turned off the light before leaving my apartment this morning*. Thus, doxastic attitudes can be regarded as the levels of confidence—call them *degrees of belief* or *credences*—individuals invest in the truth of the targeted proposition *p*. These levels of confidence are usually modelled by real-valued functions which assign a real number between 0 and 1 (inclusive) to the propositions they take as arguments.[1] This means that instead of having a binary doxastic option to take towards a proposition, an individual has, at least in principle, infinitely many doxastic options to take. Roughly put, the more confident one is in *p*, the higher one's credence in *p*. Let us call this the *quantitative model of belief* and let us dub disagreements involving credences *credal disagreements*.[2] Here is an example of credal disagreement:

(RAINY)

John is .2 confident *that it's raining in Barcelona at 5 pm on February 27th, 2017*.

Sam is .8 confident *that it's raining in Barcelona at 5 pm on February 27th, 2017*.

There is an impressive body of literature that speaks to the metaphysical and epistemological relations between belief and credences. While I won't survey such literature here, it must be acknowledged that the question whether credal and disagreement call for different definitions is sensitive to whether we endorse some form *reductionism* about doxastic attitudes to the effect that full belief reduces to credences or vice versa; or else, whether we subscribe to a form of *nonreductionism* taking belief and credences to be non-reducible types of doxastic attitudes.

Be that as it may, the problem of how to define disagreement arises even if we are reductionists. One should indeed ask: if we adopt a credence-first (or a belief-first) view, which definition of disagreement

should we accept? The pluralist view of disagreement I will defend enables us to answer this question.

That being said, it seems safe to say that it is better to carry out a pluralist-oriented inquiry into disagreement by adopting the nonreductionist approach rather than the reductionist one, in that the former is distinctively pluralist since it acknowledges the existence of a variety of doxastic attitudes in a more substantive way than the latter does. So, I will henceforth assume that this version of nonreductionism about belief and credences is correct.[3]

I have now presented the four varieties of disagreement this chapter will mostly focus on. Before going on, let me notice that the descriptive/conceptual distinction can be stated in both full and credal terms.

3   Disjunctive Pluralism

In this section, I assess the prospects for a version of pluralism about disagreement which can be deployed to make sense of the four varieties of disagreement just outlined. The pluralist view I am about to lay out can be elicited by John MacFarlane's thought-provoking chapter on the disagreement of his 2014 book on relativism. MacFarlane proposes two definitions of disagreement which accomplish two different aims: on the one hand, he proposes a definition of disagreement which is meant to take care of the idea that in order to correctly respond to the question whether two individuals instantiate the disagreement relation, we should look at the content of their attitudes as well as the *context* in which they are entertained. This idea is illustrated by (SUNNY): if we paid attention to the content of Hichem and Sophie's beliefs only, we would incorrectly take them not to disagree. On the other hand, he proposes an another definition of disagreement which is equipped to capture credal disagreement.[4] Let us present and examine both definitions in turn.

Representational acts (such as utterances) or states (such as beliefs) with propositional content are *about* individuals and the properties predicated of them and *concern* a particular context in which they occur

and with respect which they are evaluated. Context plays a role in determining both the aboutness of our utterances, by contributing to fix the reference of context-sensitive expressions such as "I", "here" and "now", as well as the situation they concern. In the standard Kaplanian framework (see Kaplan 1989), this double function of context is systematised by the distinction between *contexts of use* and *circumstances of evaluation*. Contexts of use determine the reference of context-sensitive expressions; circumstances of evaluation are the circumstances against which our beliefs (or utterances) are evaluated.

To illustrate, suppose that one believes the proposition expressed by the sentence "He's running". The belief is about the demonstrated individual male and concerns the time and the possible world at which one entertains such a belief. So, the belief is true iff the demonstrated individual male runs at the time and possible world at which one entertains the belief. Importantly, this is so independently of whether sentences express time neutral or time specific propositions. That is to say, even if we accept *temporalism*, viz. the view that sentences express time neutral propositions whose truth-value varies across times (see Kaplan 1989), one's utterance (or belief) *concerns* a particular time, viz. the time at which one has uttered the sentence (or held the belief). So, the utterance (or belief) turns out to be correct only if the uttered (believed) proposition is true relative to the time of the linguistic utterance (or mental tokening) of the sentence.

The foregoing suggests that, by affecting the truth-conditions or the truth-value distribution of believed contents, context also affects the normative status of our representational acts and states. To take proper account of this fact, MacFarlane—as well as other authors such as Francén (2010), Rieppel (2011)—brings into the picture the notion of *accuracy*:

[Full Accuracy]

A belief that p held at a certain context of use c_u is accurate just in case p is true at c_u, at the relevant circumstance of evaluation c_e.

On the basis of [Full Accuracy], MacFarlane proposes the following definition of disagreement:

[Preclusion of Accuracy View]

To disagree with someone's attitude, in this sense, is to have attitudes the accuracy of which would *preclude* its accuracy.

(MacFarlane 2014, p. 126, my emphasis)

Unfortunately, MacFarlane leaves the notion of "preclusion" unspecified. On the face of it, this makes [Preclusion of Accuracy View] hard to parse.

In Belleri and Palmira (2013) and Palmira (2017) the "preclusion" talk is unpacked in counterfactual terms. The relevant relation that should be instantiated between doxastic attitudes' accuracy conditions is as follows: if the accuracy conditions of one's attitude were fulfilled, this would ipso facto make the other's attitude inaccurate. In the light of this, I propose to state an accuracy-based definition of disagreement as follows (see also Belleri and Palmira 2013; Palmira 2017):

[Full Accuracy View]

A and B are in disagreement[5] if and only if the accuracy conditions of A's doxastic attitude are such that, if they were fulfilled, this would *ipso facto* make B's doxastic attitude inaccurate, or vice versa.

Let me now show how the [Full Accuracy View] fares with respect to various disagreement data. Take the following widely discussed example, due to MacFarlane (2007, p. 23):

(MOON)

Jane (who inhabits this world, the actual world) and June, her counterpart in another possible world. Jane believes the proposition expressed by the sentence "Mars has two moons", while June disbelieves it.

The majority view has it that cases like (MOON) are not disagreement cases.[6] The Full Accuracy View issues this verdict. The accuracy of Jane's attitude, e.g. a belief *that Mars has two moons*, is fulfilled just in case the proposition is true at the actual world w; the fulfilment of this accuracy condition, however, does not *ipso facto* make June's attitude of disbelieving that very proposition inaccurate, since the accuracy of June's attitude depends on another possible world w^*. This is why Jane and June do not disagree.

We can also easily see that the Full Accuracy View yields the right verdict about (SUNNY): Hichem and Sophie disagree since if Hichem's belief were accurate (that is, if it were the case that it's sunny in Montreal on 27th February 2017 at 4 pm), this would *ipso facto* make Sophie's belief inaccurate (for it would be false that it's not sunny in Montreal on 27th February 2017 at 4 pm).

However, closer inspection reveals that the Full Accuracy View cannot account for credal disagreement. Take (RAINY). Since the accuracy conditions of doxastic attitudes are established on the basis of the truth of the relevant proposition, we should say that a credence .2 in the proposition *that it's raining in Barcelona at 5 pm on February 27th, 2017* is accurate just in case the proposition is true. And yet, this is false, for truth provides a standard of accuracy for on/off representational states in virtue of them representing something to be the case, while there is no intuitive sense in which credences do so. So, how to account for credal disagreement?

MacFarlane discusses some examples of credal disagreement (see MacFarlane 2014, p. 122) in connection with another notion of disagreement, whose key insight is encapsulated in the following passage:

> In one sense, I disagree with someone's attitude if I could not coherently adopt the same attitude […] without changing my mind [..]. In other words, I disagree with attitudes that are not cotenable with my current attitudes.
>
> (MacFarlane 2014, p. 121)

This passage suggests the following definition of disagreement:

[Noncotenability View]

A and B disagree if and only if A cannot coherently adopt B's doxastic attitude towards *p*, and vice versa.

MacFarlane's insight *seems* correct: the Noncotenability View is equipped to issue the intuitively correct predictions about credal disagreement. Take (RAINY): it is *prima facie* plausible to hold that what gives rise to John and Sam's disagreement is the fact John cannot assign the credence Sam assigns to the proposition *that it's raining in Barcelona at 5 pm on February 27th, 2017* without being incoherent. So, their attitudes are noncotenable.

Let us now take stock and ask: what kind of theory of disagreement emerges from the discussion pursued so far? I believe that the considerations offered by MacFarlane can be articulated into a disjunctive definition of disagreement. Roughly put: being in disagreement requires either having credences satisfying the relation specified by the Noncotenability View or having full beliefs satisfying the relation specified by the Accuracy View. To put this idea a bit more rigorously: Let "c" stand for cases in which two individuals entertain doxastic attitudes, "D" for disagreement, "nv" for Noncotenability View and "ac" for Full Accuracy View:

$$\Box \forall c \ (D(c) \leftrightarrow ((D_{nv}(c) \land \text{credal}(c)) \lor (D_{av}(c) \land \text{full}(c))))$$

We might therefore take disagreement (i.e. D) to be a single higher-order relation whose obtaining is determined by two numerically distinct lower-level relations, viz. the relations stated by the NonCotenability View (i.e. D_{nv}) and the Full Accuracy View (i.e. D_{av}). Call this *Disjunctive Pluralism* about disagreement.

A key feature of Disjunctive Pluralism is that the Full Accuracy View and the Noncotenability View define disagreement by looking at different properties of doxastic attitudes. Coherence has been taken to be the most uncontroversial examples of the requirements that *epistemic rationality* imposes upon our doxastic attitudes. Rationality requires of us to hold coherent doxastic attitudes. By contrast, accuracy is a *normative* property of our doxastic attitudes. Thanks to the works of Broome

(2000), Kolodny (2005) and Parfit (2001), there is widespread agreement that rational requirements and normative requirements should be kept distinct. Briefly put, rational requirements are requirements about the structural relations amongst an individual's mental states, while normative requirements specify what counts in favour of entertaining a certain mental state, e.g. a belief. It is a further substantive question whether rational requirements are also normative in the sense of there to be epistemic reasons, viz. something on the basis of which one determines whether to φ when one's goal is truth, to respect such requirements. That is to say, it is a further substantive question whether there are epistemic reasons to be coherent. Thus, since there is no obvious conceptual link between holding coherent beliefs and holding accurate beliefs, I contend that there is no obvious conceptual link between the Noncotenability View and the Accuracy View.

In my view, the key feature of Disjunctive Pluralism just highlighted makes such view subject to a worry that any disjunctive pluralist theory runs up against. To illustrate the worry, it must be kept in mind that the various cases of disagreement discussed so far are taken to be data about disagreement *as such*. It seems that we are ready to use the same concept, viz. DISAGREEMENT, to classify those data, no matter whether they involve credences or full beliefs and no matter whether we have reductionist inclinations or not. This suggests that, even if we acknowledge that disagreement comes in many varieties, we should nonetheless preserve the unity of the phenomenon. That is to say, we have to preserve the apparent fact that all these varieties are, to put it roughly, varieties of the same thing. If we are not able to do so, we have at least to explain this apparent fact away. Insofar as a Pluralist theory of disagreement does not tell us what is common to all varieties of such a phenomenon (or why these varieties mistakenly appear to have something in common), it seems safe to say that there is something amiss with her view.

To be sure, this is a structural challenge that pluralist theories about any subject matter (e.g. truth, logical consequence, justification, self-knowledge and so on) should meet in order to offer a plausible characterisation of our linguistic and conceptual practices that seem to point to the fact that there is *a* phenomenon, i.e. disagreement, which

can manifest itself in different ways. For labelling purposes, call this the *Unity Challenge*.

However, even if Disjunctive Pluralism managed to satisfactorily address the Unity Challenge, it would have to face other problems which makes it ultimately unappealing.

To begin with, consider the following case:

(MOON-Cr)

Jane (who inhabits this world, the actual world) and June, her counterpart in another possible world. Jane assigns .4 credence to the proposition *that Mars has two moons*. June assigns .8 credence to the same proposition.

Which definition of disagreement does Disjunctive Pluralism recommend adopting in order to deal with (MOON-Cr)? One might claim that since (MOON-Cr) involves credences, it falls within the remit of the Noncotenability View. This view, however, treats (MOON-Cr) as a genuine case of disagreement, for Jane and June have noncotenable attitudes towards the same proposition. However, the majority view maintains that in cases involving counterfactual scenarios the disagreement relation is not instantiated. This is the intuitive verdict about (MOON), and there's no reason why it should not carry over to (MOON-Cr). Since I believe that the majority view is correct, the Disjunctive Pluralist had better not treat (MOON-Cr) via the Noncotenability View.

If we are Disjunctive Pluralists, the other option available is the Full Accuracy View. And yet, since the Full Accuracy View is stated for full beliefs, such view is unable to yield any verdict about (MOON-Cr) whatsoever, in that this case involves subjects entertaining partial rather than full attitudes. So, it seems that Disjunctive Pluralism could not make sense of cases such as (MOON-Cr).

At this stage, however, one might suggest that the Noncotenability View could be modified so as to deliver the intuitively correct verdict about (MOON-Cr). For instance, one might deny one of the core assumptions of the semantic debate about disagreement, namely that truth is relative to possible worlds, by proposing the idea that possible worlds are built into the proposition, as it were. This view could

be developed either via an index-free semantics to the effect that truth is not relative to possible worlds,[7] or by reconciling possible worlds semantics with modally neutral contents.[8]

On this *new* Noncotenability View, Jane and June would not disagree because they would entertain different propositional contents: roughly, the proposition *that Mars has two moons in w*, and the proposition *that Mars has two moons in w**. Since these two propositions are perfectly cotenable, the new Noncotenability View yields the intuitively correct verdict about (MOON-Cr).

On reflection, however, the new Noncotenability View would commit us to a specific account of the relation between propositions and possible worlds in order to deliver the intuitively correct verdict about (MOON-Cr), in that it commits us to the view that propositions are modally specific. This is a significant theoretical cost which some—perhaps most (MacFarlane included)—of us are not ready to pay.

Finally, consider this other case:

(TYLER)

Graham assigns .35 credence to the proposition *that Steve Tyler will be the next Nobel literature laureate* and Julie assigns credence .36 to the same proposition.

Do Graham and Julie disagree? Some might answer the question in affirmative, others in the negative. It is far from obvious how to establish which intuition is the correct one, for both of them strike us as *prima facie* permissible. Such permissibility has to be compatible with our definition of disagreement; that is to say, a good definition of disagreement must be so flexible as to enable us to accommodate conflicting intuitions about (TYLER) and, more generally, about cases where the difference between degrees of belief is very little and intuitions are shaky.

However, both the new and old Noncotenability View commit us to taking (TYLER) to be a case of disagreement: as a matter of fact, Graham's and Julie's attitudes are noncotenable, for one cannot

coherently assign credence .35 and credence .36 to the same proposition. So, the Noncotenability View is not flexible enough.

Let us take stock. I have developed a version of Disjunctive Pluralism according to which two radically different definitions of disagreement are needed to account for credal and full disagreement. I have argued that Disjunctive Pluralism faces the Unity Challenge, it fails to be extensionally adequate without committing us to theoretically controversial assumptions about the nature of semantic content and it is not flexible enough to accommodate uncertainty about some disagreement-related data. This, to my mind, suffices to warrant examination of a new pluralist theory of disagreement.

4   Kinship Pluralism

I believe that we can draw four lessons from the discussion of Disjunctive Pluralism pursued so far:

(1)   A good pluralist account of disagreement has to make sense of the Unity Challenge.
(2)   Any such account has to be extensionally adequate, thereby complying with the intuitively correct verdicts about *clear* cases of disagreement (or lack thereof).
(3)   It has to make room for the permissibility of diverging intuitions about non-clear cases, such as (TYLER).
(4)   It has to be extensionally adequate and respectful of diverging intuitions about non-clear cases without burdening us with too much theoretical commitments (e.g. in terms of the nature of the semantic contents we believe).

(1)–(4) can be plausibly seen as the *desiderata* for a good pluralist account of disagreement.

Thus, my criticism of Disjunctive Pluralism is that it is not clear how it can meet (1), it fails to meet (3) and either it doesn't meet (2) or it doesn't meet (4).

I turn now to articulate a different version of pluralism about disagreement, which I will dub *Kinship Pluralism* that meets the *desiderata* in a more satisfactory way than Disjunctive Pluralism does.

The discussion of cases such as (BALD), (SUNNY) and (MOON) teaches us that that in order to establish whether two individuals hold doxastic attitudes which give rise to a disagreement, we have to pay attention not only to the content of such attitudes, but also to the circumstances they concern. The notion of full accuracy defined above enables us to preserve this idea, and the Full Accuracy View defines the kind of relation that must be in place between full beliefs' accuracy conditions in order for such beliefs to give rise to a disagreement. In Belleri and Palmira (2013), we have argued at length that the Full Accuracy View enables us to satisfy *desiderata* (2) and (4) insofar as we focus on full belief.

The key insight behind Kinship Pluralism is that we can offer a plausible account of both credences' accuracy conditions, and what it takes for them to be fulfilled such that we can state a counterfactual relation between credences' accuracy conditions that mirrors the counterfactual relation required for full disagreement. This enables us to see that credal and full disagreement share a significant degree of kinship, as it were, or so I shall contend. Let me develop this strategy in some detail.

In recent times, thanks to the work of authors such as Joyce (1998), Leitgeb and Pettigrew (2010), Pettigrew (2016), formal epistemologists have developed a new way of linking credences to truth. If we take credences to determine certain epistemically relevant quantities and if we take these quantities to be propositions, then we can take one's credence in p to be one's estimate of the truth-value of p. This provides us with an intuitive way of defining the accuracy conditions of credences which mirrors the case of full belief: one's credence is accurate to certain a degree d depending on how good one's estimate of p's truth-value is. The more one's credence is closer to p's actual truth-value, the more accurate it is.

We can implement these ideas formally by using *scoring rules* (see Joyce 1998). For a proposition p, a credence Cr and a truth-value v, a scoring rule assigns a real number ≥ 0 which measures the inaccuracy of holding Cr when the truth-value of p in w is as given in v. The best

score achievable is 0, in that the distance from p's actual truth-value is minimised to 0. If Cr's score is higher than Cr^*'s is, Cr^* is closer to the truth than Cr is; so, it is less inaccurate (and therefore more accurate).

This way of measuring the (in)accuracy of credences can be extended so as to take the truth-values distribution of propositions to vary across both possible worlds and times. That is to say, such a measure can be given independently of what coordinates constitute circumstances of evaluation. Thus, we can state the following definition of *gradational accuracy*:

[Gradational Accuracy]

A credence Cr in p held at a certain context of use c_u is accurate to degree d just in case Cr is close to degree m to p's truth-value at c_u, at the relevant circumstance of evaluation c_e.

Let me now clarify the notion of fulfilment of credences' accuracy conditions. As has emerged previously, the fulfilment of full beliefs' accuracy conditions is categorical: since either the believed proposition is true or it is not, and since accuracy in the full case is defined on the basis of truth, it follows that either the accuracy conditions are fulfilled or they are not. However, since the degree to which credences are in accurate depends on their closeness to the truth, credences' accuracy conditions should be seen as fulfilled to a certain degree, where this degree is given by their absolute (in)accuracy score. Thus, it is natural to maintain that credences' accuracy conditions are fulfilled to the maximal degree only when one entertains a true proposition with the highest degree of belief; they are fulfilled to the minimal degree only when one entertains a true proposition with the lowest degree of belief; and they are fulfilled to a certain degree when only when one entertains a true proposition with some other degree of belief.

In the light of this understanding of what it takes for credences to be accurate and what it takes for their accuracy conditions to be fulfilled, I propose the following definition of credal disagreement:

[Partial Accuracy View]

A and B disagree[9] if and only if the accuracy conditions of A's doxastic attitude are such that, for any degree c, if they were fulfilled to degree c, this would *ipso facto* make B's doxastic attitude inaccurate to degree b, or vice versa.

I turn now to assess the Partial Accuracy View against *desiderata* (2)–(4).[10]

First, it must me noticed that the Partial Accuracy View takes credal disagreement to come in degrees in virtue of the fact that its *definiendum* is a counterfactual relation between credences' accuracy conditions which has built into it the idea that such accuracy conditions can be fulfilled to a certain degree.

I believe that taking credal disagreement to come in degrees is not only acceptable, but also desirable. For one thing, we should bear in mind that credal disagreement takes place against a purely quantitative doxastic background; so, it would be surprising if our definition of a certain relation between credences (and their accuracy conditions) turned out to be of a different kind—qualitative as opposed to quantitative—than the *relata* themselves. For another, regarding credal disagreement as being itself graded enables the Partial Accuracy View not to take any definite stance about cases such as (TYLER). The counterfactual relation stated by the Partial Accuracy View has it that if A's credence's accuracy conditions were fulfilled to a certain degree c, this would make ipso facto B's credence inaccurate to a certain degree b. The Partial Accuracy View has it that, in (TYLER), such a counterfactual relation is instantiated to a certain (very minimal) degree d, which is the measure of the absolute distance between the inaccuracy scores of credences .35 and .36. The question whether such a distance is enough to give rise to a *real* disagreement is independent of the Partial Accuracy View. That is to say, the Partial Accuracy View is compatible with both an affirmative and a negative answer to this question. Thus, the Partial Accuracy View enables us to comply with *desideratum* (3).

Secondly, I submit that the Partial Accuracy View enables us to satisfy *desideratum* (2). While a full vindication of the extensional

adequacy of Partial Accuracy View comes down to a case-by-case scrutiny that cannot be undertaken here, there is a good reason to think that the view yields the intuitively correct verdicts about various credal disagreement data. To illustrate this point, let us take the problematic (MOON-Cr). The Partial Accuracy View yields the verdict that Jane and June do not disagree. Since credences' accuracy is established relative to possible worlds, and since Jane's and June's attitudes are held in different worlds, it turns out that Jane's attitude has a certain degree of closeness to the truth relative to w, whereas June's credence has a different degree of closeness to the truth relative to another possible world w^*. So, it turns out that however fulfilled the accuracy conditions of Jane's credence might be, such fulfilment can never make June's credence inaccurate to any degree, for the accuracy of June's credence depends on how things in w^* are whereas the accuracy of Jane's credence depends on how things in a different possible world w, are. This means that the relevant relation between the accuracy conditions of Jane's and June's attitudes does not obtain. Therefore, Jane and June do not disagree.

Thirdly, the Partial Accuracy View yields the correct verdict about (MOON-Cr)—as well as about (RAINY)—quite independently of how we conceive of the relation between propositions and possible world semantics. More specifically, and unlike the Noncotenability View, the Partial Accuracy View captures the absence of disagreement in (MOON-Cr) even if we accept the standard idea that the contents of beliefs and sentences are sets of possible worlds. Thus, the Partial Accuracy View enables us to comply with (4).

Let us take stock. The discussion pursued in this section, together with the outline of the Full Accuracy View offered Sect. 3, leads us to the following intermediate conclusion: There are two different definitions of disagreement, i.e. the Full Accuracy View and the Partial Accuracy View, which comply with *desiderata* (2)–(4) in a satisfactory way.

We should not rest content with this intermediate conclusion, though. For a good pluralist account of disagreement has to take up the Unity Challenge. That is to say, it has to explain (or explain away) the appearance that credal and full disagreement are varieties of the same

phenomenon. The next step is therefore to bring out the common traits of these two definitions, thereby showing that the Full and the Partial Accuracy View enjoy a significant degree of kinship, as it were.

Let me start off with noticing that both the Full Accuracy View and the Partial Accuracy View take disagreement to arise in virtue of the instantiation of a counterfactual relation between the attitudes' accuracy conditions such that the fulfilment (or partial fulfilment) of one's attitude's accuracy conditions *impacts negatively* on the other's attitude's accuracy conditions. This "negative impact" relation between attitudes' accuracy conditions is the first shared trait of the Full Accuracy View and the Partial Accuracy View.

Of course, what it takes for the fulfilment of a full belief's accuracy conditions to negatively impact on the accuracy of another full belief differs from what it takes for the (partial) fulfilment of credence's accuracy conditions to negatively impact on the (partial) accuracy of another credence. This is so since gradational and full accuracy are two distinct properties: the former comes in degrees, therefore making it the case that the fulfilment of accuracy conditions is partial, whereas the latter is absolute, therefore making it the case that the fulfilment of accuracy conditions is categorical. However, it must be noticed that both gradational and full accuracy provide a *truth-directed* standard of evaluation for doxastic attitudes. In virtue of their being doxastic—as opposed to conative—the accuracy of such attitudes can't but be involve truth, for even if gradational and full accuracy are two distinct properties, they can both be traced back to what I take to be the key concept we deploy while reflecting on whether a doxastic attitude is accurate, that of *truth-directedness*. I maintain that TRUTH-DIRECTEDNESS is instantiated by different properties depending on the kind of doxastic attitude, i.e. full or partial, we focus on. In the case of full belief, the concept is instantiated by the property of full truth (relative to the relevant circumstance of evaluation c_e); in the case of partial belief, it is instantiated by the property of approximation to the truth (always relative to the relevant circumstance of evaluation c_e).

As far as I can see, the fact that both the Full Accuracy View and the Partial Accuracy View involve the claim that disagreement requires that the fulfilment of one's attitude's accuracy conditions has a negative

impact on the accuracy conditions of the other's, together with the fact that gradational and full accuracy are different ways of instantiating the *same* concept whereby we characterise the very idea of a doxastic attitude being accurate, shows that the Full Accuracy View and the Partial Accuracy View display a significant degree of kinship. No matter whether credal or full, disagreement must involve a counterfactual relation between attitudes' accuracy conditions, and such accuracy conditions are determined on the basis of the truth-directedness of the relevant attitudes. This, I contend, enables us to preserve the apparent unity of the phenomenon of doxastic disagreement across its credal and full manifestations. Thus, I regard the Unity Challenge met.[11]

Let us take stock. I have argued that Kinship Pluralism vindicates desiderata (1)–(4) in a satisfactory way. Or, to say the least, it does so in a better way than Disjunctive Pluralism. Thus, I contend that Kinship Pluralism offers a good pluralist account of credal and full disagreement. I wish to conclude this section by turning to the descriptive/conceptual pair singling out the other two central varieties of disagreement that have been mostly discussed in the literature. Let us therefore get back to (BALD) and suppose that Mary and Albert are using the adjective "bald" metalinguistically to convey a disagreement about the correct usage of that term in that context.

To begin with, notice that the question of what kinds of proposition is expressed—or perhaps better, pragmatically conveyed—by sentences featuring expressions used in a metalinguistic way is a hotly debated one. Fortunately, however, Kinship Pluralism does not commit us to answering this question in one way or another. Insofar as conceptual disagreement is conceived along the lines of doxastic disagreement, we should picture Mary and Albert as entertaining conflicting doxastic attitudes towards the relevant pragmatically conveyed propositions of the metalinguistic negotiation. For any kind of specification of the propositions pragmatically conveyed by metalinguistic uses of sentences, all we need to establish is whether the doxastic attitudes held by individuals towards such propositions are such that their accuracy conditions instantiate the relation defined either by the Full Accuracy View or by the Partial Accuracy View. So, nothing in principle prevents Kinship

Pluralism from accounting for the conceptual disagreement in the same way it accounts for descriptive disagreement.

This is good news, for we tend to take both conceptual and descriptive disagreements to be varieties of the same phenomenon. Kinship Pluralism explains this fact by subsuming both descriptive and conceptual disagreement under the view that they both require the instantiation of a given counterfactual relation between attitudes' accuracy conditions.

5   Possible Extensions: Agnostic and Group Disagreement

I have so far dealt with the four varieties of doxastic disagreement which have been mostly invoked in current epistemological and semantic debates about disagreement. However, one might wonder whether disagreement reaches even further than these four varieties and ask how Kinship Pluralism would fare with respect to such comparatively less studied varieties. More specifically, in this section, I briefly assess the possibility of there to be disagreements in which one of the parties is agnostic about the matter at hand and the possibility of there to be disagreements amongst groups.

Traditionally, philosophers have distinguished between three distinct kinds of doxastic attitude one might have towards p: one can either believe it, or disbelieve it or being agnostic (or being suspended) about it. So, to take the most familiar conflict involving agnosticism:

(GOD)

Athos says: "God does not exist";

Agnos says: "I've made up my mind to suspend judgment on whether God exists or not".

Surely, Athos and Agnos are in some sort of conflict or opposition regarding God's existence. But do they disagree? Do we have any strong intuition about the kind of conflict they're having?

I must confess that, speaking for myself, I am at loss for an intuitive response here. I doubt that there is a clear intuitive agnostic disagreement datum to be captured in the first place without delving a little more into the question of what it means to be agnostic (or suspended) about a given matter.

Characterising agnosticism, however, is no easy matter. The safest thing we can say is that being agnostic amounts to being in a state of committed epistemic neutrality or indecision towards the matter at hand. Such an idea has been articulated in different ways:

First-order reductionism: agnosticism is either assigning .5 credence to *p* and its negation (Kelly 2010), or assigning some imprecise credence to them (Hájek 1998; Van Frasseen 1998), e.g. credences taking values from the [1/3, 2/3] interval.

Higher-order reductionism (e.g. Bergmann 2005): agnosticism is holding the belief that one does not/cannot have evidence for either *p* or not-*p*.

Nonpropositionalism (e.g. Friedman 2017): agnosticism is a *question-directed* (viz. non-propositional) *sui generis* doxastic attitude.

These proposals take agnosticism to be reducible to a mental state. However, it might well be the case that a proper characterisation of agnosticism requires something more than identifying it with a mental state. For instance, Rosenkranz (2007) argues that in order to make sense of the idea that agnosticism stands in opposition to belief (or endorsement) and disbelief (or denial), we have to characterise it as a *stance*. Stances are individuated by means of assertions and their constitutive commitments regarding a defence of one's stance in a debate with one's opponents. Rosenkranz calls "True Agnosticism" the stance to the effect that we are not in a position to know the truth-value of a proposition and will continue not to know it relative to all the states of information, we can reach through our current methods of investigation and cognitive powers.

In the light of this overview of the different approaches to agnosticism, it seems that the only answer to question whether an individual being agnostic about *p* and another believing that *p* disagree is a slightly

frustrating: "It depends on which view about agnosticism one subscribes to". Let me unpack this a little.

If first-order reductionism were correct, we could say that Athos and Agnos entertain different credences giving rise to credal disagreement. In such a case, we could deploy the Partial Accuracy View and establish that Athos and Agnos disagree to some degree. This is so since credences are attitudes which can stand in a disagreement relation. However, things are much more complex with higher-order reductionism, Friedman's non-propositionalism and Rosenkranz's stance-based approach.

Take higher-order reductionism: the impossibility of having evidence in favour of p or its negation does not make one's belief that p inaccurate (unless one accepts some antirealist-oriented epistemic constraint on truth). The converse also holds: one's belief that p being accurate would not make the attitude to the effect that one cannot have evidence in favour that p and its negation ipso facto inaccurate. This shows if higher-order reductionism about agnosticism were correct, Athos and Agnos would not be in disagreement according to Kinship Pluralism.

As for non-propositionalism, let me notice that Friedman does not explicitly formulate a norm governing suspended judgement, even though she takes such an attitude to be normatively incompatible with knowledge. Since the Kinship Pluralist takes disagreement to be a normative phenomenon, not much can be said about whether (GOD) comes out as a case of disagreement or not under this interpretation.

Finally, Rosenkranz devotes a great deal of attention to the question whether agnosticism is a genuine third stance that stands in opposition to both positive (i.e. endorsement) and negative (i.e. denial) stances. However, if agnosticism requires such a rich characterisation in terms of epistemic statues and commitments, and it cannot be reduced to a more basic cognition, it might seem reasonable not to talk of disagreement and speak instead—as Rosenkranz does—of "opposition". As far as I can see, individuals can be in a disagreement relation merely in virtue of the attitudes they entertain and quite independently of the individuals' epistemic statuses and commitments to defend their views in a debate. This seems to be a basic fact about disagreement which is held

fixed across all its. This is not to say, of course, that agnostics are not in conflict with believers (or endorsers). It is quite in the spirit of epistemic pluralism to recognise a variety of ways whereby two individuals can stand in an epistemically relevant conflict which need not necessarily reduce to disagreement.

Let us turn now to group disagreement. The question whether two groups disagree and how to account for such a disagreement is intertwined with one of the central debates in collective epistemology centring on the types of attitude that can be held by groups.

Let us begin with the following piece of linguistic evidence: we usually ascribe beliefs to groups. It indeed seems perfectly appropriate to say, for instance, that F.C. Barcelona's fans believe that Lionel Messi is better than Cristiano Ronaldo, and it is equally appropriate to say that Real Madrid's fans disagree. However, there is much debate about whether group belief ascriptions have to be taken at face value and, even if so, there are various ways of explaining what it means for groups to believe propositions (see Gilbert and Pilchman 2014 for a recent opinionated overview). So, the question is: should we take these pieces of linguistic evidence at face value?

According to so-called *believers*, we should: groups do believe propositions. According to so-called *rejectionists*, we should not: groups do not believe but rather *accept* propositions, where *acceptance* and belief differ on a number of counts (see for instance Engel 1998; Proust 2012). Furthermore, within the believers camp, we should distinguish between *summativists* and *non-summativists*: the former hold that a group G believes that p iff all or most of the members of G believe that p, whereas the latter maintain that G's belief that p does not reduce to G's members individual beliefs. Let me say something on each of these options.

It seems safe to contend that the summativist believer is naturally committed to the idea that groups can disagree. A moment of reflection shows that the Kinship Pluralist View I defend is compatible with this option. On a simple summativist view, a group's belief reduces to, say, the belief that most of the members of G have towards p. Thus, we can say that the accuracy conditions of the group's belief that p reduce to the accuracy conditions of the belief that is held by most of the members

of G. So, in order for two groups to disagree, most of their respective members must hold beliefs whose accuracy conditions instantiate the counterfactual relation specified by the Full Accuracy View (and the same holds, *mutatis mutandis*, if we ascribe credences to groups).

The non-summativist believer is also committed to the possibility of there being disagreement between groups. However, on the non-summativist view, the accuracy conditions of the group's belief that *p* cannot be reduced to the accuracy conditions of the belief that is held by most of the members of the group. I do not take this to be a fatal problem, though. Insofar as the non-summativist believer provides us with a way of stating the accuracy conditions of the group's belief, we can still take disagreement to be defined either by the Full Accuracy View or by the Partial Accuracy View. Thus, Kinship Pluralism is also compatible with non-summativism.

Let us finally take the rejectionist option. Rejectionists maintain that groups hold true propositions by *accepting*—as opposed to believing— them. I can't examine here in detail the differences between acceptance and belief, but let me stress that in order for us to define group disagreement by adopting Kinship Pluralism within a rejectionist framework, it must be possible to state accuracy conditions for acceptances such that they could instantiate the counterfactual relation specified by the Full (or Partial) Accuracy View. Authors such as Kaplan (1981) and Proust (2012) maintain that acceptances have accuracy conditions while, at the same time, advocating a pluralist stance to the effect that accepting *p* is subject to two (Kaplan) or more (Proust) norms depending on the goals one is pursuing while considering whether to accept *p*. The possibility that the relevant dimension of normative assessment of acceptances is not accuracy but something else leaves us in the following situation with respect to group disagreement. Insofar as we focus on the accuracy-related normative dimension of acceptances, it is possible to deploy Kinship Pluralism within the rejectionist approach to group belief and take F.C. Barcelona's fans to disagree with Real Madrid's fans about whether Lionel Messi is better than Cristiano Ronaldo: if the accuracy conditions of accepting the proposition that Lionel Messi is better than Cristiano Ronaldo were satisfied, this would make *ipso facto* the acceptance of the negation of this proposition inaccurate. However, since two

acceptances can be assessed against other non-truth-directed normative dimensions, it might well be the case that sometimes groups have conflicts that cannot be captured by the counterfactual relation specified by the Full (or Partial) Accuracy View. I submit that such conflicts should therefore not be regarded as genuine disagreements. For they are best seen—in a fairly pluralist spirit—as other varieties of normative conflict.

6   Conclusions

I want to conclude by highlighting some of the wide-ranging implications of the discussion pursued so far. First of all, Kinship Pluralism enables us to establish that the relativism *vs.* contextualism debate and the peer disagreement debate can be said to be illuminating different aspects *of the same phenomenon*, in that whereas the former focuses on full disagreement and the latter on credal disagreement, these varieties of disagreement are accounted for by two definitions, i.e. the Partial Accuracy View and the Full Accuracy View, which share a significant degree of kinship. This point is far from trivial, for epistemologists and philosophers of language have been working on disagreement in an almost complete isolation from one another. So, having established that they are, after all, targeting different features of the same phenomenon should be welcomed as good news and motivate us to develop transversal semantic-epistemological approaches to disagreement.

Secondly, the degree of semantic flexibility exhibited by both the Full and the Partial Accuracy View, and their ability to capture both descriptive and conceptual disagreement, is such that it supports the contention that at least some of the varieties of disagreement relativists and contextualists claim to be in a position to redeem can, after all, be unified. I take this to have a twofold implication. On the one hand, parties in this debate are not talking past to one another, and they have been having a genuine dispute around disagreement. On the other hand, though, since these views can all—at least in principle—subscribe to Kinship Pluralism, it seems that there is no real point in debating about which theory captures disagreement data better than any other. So, the arguments proposed in this chapter strengthen the point, already made

in previous works (see Belleri and Palmira 2013; Palmira 2015), that we won't find any support for relativism (or contextualism, for that matter) from disagreement data.

Thirdly, and perhaps most surprisingly, the present investigation can be taken to have a bearing on the question of how to define conative disagreement. To see how, let me begin with pointing out a neglected asymmetry between doxastic and conative disagreement. I submit that while it is relatively easy to provide intuitive data about when two individuals holding doxastic attitudes are in disagreement, it is much harder and less common to offer extensional adequacy considerations in favour of our preferred view of it than it is in the doxastic case. If we asked two laymen whether their respective preferences—say, one prefers playing basketball over football whereas the other prefers playing the latter over the former—give rise to a disagreement, they would simply be at loss. Similarly, I believe that no clear intuition about whether two individuals desiring the course of action in the same circumstances give rise to something we would be ready to call "disagreement".

This should not come as a surprise, for the idea of there to be a disagreement relation between conative attitudes has been, first and foremost, introduced by meta-ethical expressivists—originally by Stevenson (1963)—to show that expressivism can somehow accommodate the intuition of there being genuine disputes about moral matters. Recently, contextualists about predicates of personal taste (see López de Sa 2015; Marques 2015; Marques and García-Carpintero 2014) have invoked conative disagreement to reply to the lost disagreement challenge raised by relativists (see MacFarlane 2014).

Thus, insofar as it is hard to be guided by intuitions about conative disagreement, and since this phenomenon has been mostly invoked as a means to a theoretical end, I take it that one important way of constraining our inquiry into the nature of conative disagreement is to rely on our independently well-understood notion of doxastic disagreement. Let me explain. Since we have a somewhat stable take on the central cases of doxastic attitudes giving rise to disagreement, and since I have now provided a full-fledged characterisation of doxastic disagreement which respects several *desiderata*, I contend that if there is any interesting sense in which conative attitudes can stand in relation

which we would be ready to call "disagreement" and classify under DISA-GREEMENT, this relation should bear at least some similarity to the one defined for the doxastic case. What I have specifically in mind here is the fact that the definition of doxastic disagreement I advocate is such that it makes this phenomenon *normative* in kind. While conative attitudes, by definition, do not have accuracy (i.e. truth-related) conditions, it appears fair to say that they are subject to normative evaluation. So, since it is possible to define normative properties for conative attitudes, and since the variety of disagreement we understand better, i.e. doxastic disagreement, is normative in kind, we should expect from an account of conative disagreement to define this relation in broadly normative terms in order for it to preserve the apparent unity of the phenomenon of disagreement across its doxastic and conative varieties.

Notes

1. Supporters of the so-called Bayesian approach to epistemic rationality maintain that in order for agents to be rational, their credence functions should obey the axioms of probability calculus and evolve by conditionalisation. However, my discussion need not assume the Bayesian approach.
2. Henceforth, I will interchangeably use "full belief(s)", and "full attitude(s)" to refer to doxastic attitudes as conceived within the qualitative model of belief. I will interchangeably use "partial belief(s)", "degree of belief(s)" "partial attitude(s)" and "credence(s)" to refer to doxastic attitudes as conceived within the quantitative model of belief.
3. See Palmira (2017) for more details on this issue.
4. To forestall misunderstandings: MacFarlane's main aim in his work on disagreement is not to provide a careful account of credal and full disagreement in general but rather to show that in order to capture a certain class of disagreement data, namely those involving the use of predicates of personal taste, epistemic modals and so on, we should adopt his brand of truth-relativism. Nevertheless, it must be acknowledged that MacFarlane is one of the few authors who explicitly address both credal and full disagreements data and treat them via two different definitions of disagreement.

5. This abbreviates: *A* disagrees with *B*'s φ-ing in context *c*, where "φ" can be replaced here with the verb "believe" (or "disbelieve"). See MacFarlane (2014).

6. See e.g. MacFarlane (2007), (2014), Marques (2014), Schaffer (2011). The consensus is not universal: for instance, Cappelen and Hawthorne (2009) have tried to resist the majority view. It is not my aim here to discuss in detail this issue.

7. See e.g. Schaffer (2012).

8. See e.g. Ninan (2010).

9. This abbreviates: *A* disagrees with *B*'s φ-ing in context *c* to degree *d*, where 'φ' can be replaced here with the verb 'assigning a given credence to *p*'.

10. I spell out the formal aspects of the Partial Accuracy View in detail in Palmira (2017).

11. This way of understanding the relation between truth-directedness on the one hand, and full truth and approximation to truth on the other, is obviously inspired to Crispin Wright's version of alethic pluralism (see Wright 2012). Wright's idea is that there is a single concept of truth and different truth properties that satisfy the concept in different domains of discourse. My claim is that there is a single concept of accuracy, viz. truth-directedness, and different accuracy-relevant properties that satisfy the concept depending on whether we take attitudes to be full or partial.

References

Barker, C. 2002. The Dynamics of Vagueness. *Linguistics and Philosophy* 25: 1–36.

Belleri, D., and M. Palmira. 2013. Towards a Unified Notion of Disagreement. *Grazer Philosophische Studien* 188: 139–159.

Bergmann, M. 2005. Defeaters and Higher-Level Requirements. *The Philosophical Quarterly* 55: 419–436.

Broome, J. 2000. Normative Requirements. *Ratio* 12: 398–419.

Cappelen, H., and J. Hawthorne. 2009. *Relativism and Monadic Truth*. Oxford: Oxford University Press.

Carter, J.A. 2016. Group Peer Disagreement. *Ratio* 29 (1): 11–28.

Christensen, D., and J. Lackey (eds.). 2013. *The Epistemology of Disagreement: New Essays*. Oxford: Oxford University Press.

Engel, P. 1998. Believing, Holding True, and Accepting. *Philosophical Explorations* 1 (2): 140–151.

Feldman, R., and T. Warfield (eds.). 2010. *Disagreement*. Oxford: Oxford University Press.

Francén, R. 2010. No Deep Disagreement for New Relativists. *Philosophical Studies* 151 (1): 19–37.

Friedman, J. 2017. Why Suspend Judging. *Noûs* 51 (2): 302–326.

García-Carpintero, M., and M. Kölbel (eds.). 2008. *Relative Truth*. Oxford: Oxford University Press.

Gilbert, M., and D. Pilchman. 2014. Belief, Acceptance, and What Happens in Groups. In *Essays in Collective Epistemology*, ed. J. Lackey, 189–212. Oxford: Oxford University Press.

Hájek, A. 1998. Agnosticism Meets Bayesianism. *Analysis* 58: 199–206.

Huvenes, T. 2012. Varieties of Disagreement and Predicates of Personal Taste. *Australasian Journal of Philosophy* 90: 167–181.

Joyce, J.M. 1998. A Nonpragmatic Vindication of Probabilism. *Philosophy of Science* 65 (4): 575–603.

Kaplan, D. 1989. Demonstratives. In *Themes From Kaplan*, eds. J. Almog, J. Perry, and H. Wettstein, 481–563. New York: Oxford University Press.

Kaplan, M. 1981. Rational Acceptance. *Philosophical Studies* 40 (2): 129–145.

Kelly, T. 2010. Peer Disagreement and Higher-Order Evidence. In *Disagreement*, eds. R. Feldman, and T. Warfield, 111–174. Oxford: Oxford University Press.

Kolodny, N. 2005. Why Be Rational? *Mind* 114 (455): 509–563.

Leitgeb, H., and R. Pettigrew. 2010. An Objective Justification of Bayesianism I: Measuring Inaccuracy. *Philosophy of Science* 77 (2): 201–235.

López de Sa, D. 2015. Expressing Disagreement: A Presuppositional Indexical Contextualist Relativist Account. *Erkenntnis* 80 (1): 153–165.

MacFarlane, J. 2007. Relativism and Disagreement. *Philosophical Studies* 132: 17–31.

———. 2014. *Assessment Sensitivity. Relative Truth and Its Applications*. Oxford: Oxford University Press.

Marques, T. 2014. Doxastic Disagreement. *Erkenntnis* 79 (1): 121–142.

———. 2015. Disagreeing in Context. *Frontiers in Psychology* 6 (257): 1–12.

Marques, T., and M. García-Carpintero. 2014. Disagreement about Taste: Commonality Presuppositions and Coordination. *Australasian Journal of Philosophy* 92 (4): 701–723.

Ninan, D. 2010. Semantics and the Objects of Assertion. *Linguistics and Philosophy* 33: 355–380.

Palmira, M. 2015. The Semantic Significance of Faultless Disagreement. *Pacific Philosophical Quarterly* 96 (3): 349–371.

———. 2017. Disagreement, Credences, and Outright Belief. *Ratio* 1–18. online first, doi: 10.1111/rati.12163.

Parfit, D. 2001. Reasons and Rationality. In *Exploring Practical Rationality*, eds. D. Egonsson et al., 17–39. Aldershot: Ashgate.

Pettigrew, R. 2016. *Accuracy and the Laws of Credence*. Oxford: Oxford University Press.

Plunkett, D., and T. Sundell. 2013. Disagreement and the Semantics of Normative and Evaluative Terms. *Philosophers' Imprint* 13 (23): 1–37.

Proust, J. 2012. The Norms of Acceptance. *Philosophical Issues* 22: 316–333.

Rieppel, M. 2011. Stoic Disagreement and Belief Retention. *Pacific Philosophical Quarterly* 92: 243–262.

Rosenkranz, S. 2007. Agnosticism as a Third Stance. *Mind* 116 (461): 55–104.

Schaffer, J. 2011. Perspective in Taste Predicates and Epistemic Modals. In *Epistemic Modality*, eds. A. Egan and B. Weatherson, 179–226. Oxford: Oxford University Press.

———. 2012. Necessitarian Propositions. *Synthese* 189 (1): 119–162.

Stevenson, C.L. 1963. *Facts and Values*. New Haven: Yale University Press.

Sundell, T. 2011. Disagreements about Taste. *Philosophical Studies* 155 (2): 267–288.

Van Fraassen, B.C. 1998. The Agnostic Subtly Probabilified. *Analysis* 58: 212–220.

Wright, C. 2012. A Plurality of Pluralisms. In *Truth and Pluralism: Current Debates*, eds. N.J.L.L. Pedersen and C.D. Wright, 123–153. Oxford: Oxford University Press.

A Pluralistic Way Out of Epistemic Deflationism About Ontological Disputes

Delia Belleri

1   The Composition Debate

One of the main questions that ontologists today debate is whether there are composite objects. Theorists who maintain that there are no composites, but only simples, are commonly known as *Nihilists*. And let us call those who believe that there are composite objects, *Non-Nihilists*. Among the non-nihilists, let us say that Commonsensicalists believe that the domain of composite objects coincides with the domain of the objects recognised by common sense; while Universalists believe that the domain of composite objects goes much beyond that of common-sense entities, because composition always occurs, that is to say, it is sufficient that two objects exist for them to compose something.[1]

What motivates adoption of each of these accounts? Reviewing the arguments that have been advanced by the defenders of each position,[2]

D. Belleri (✉)
Department of Philosophy, University of Vienna, Vienna 1010, Austria
e-mail: deliabelleri@gmail.com

it is possible to identify the general motivation behind each view in the kind of theoretical features it favours and seeks to vindicate.

Defenders of Nihilism seem to be concerned with at least three things. First, a response to the question whether there are composite objects should *avoid certain paradoxical implications.* One of these is the sorites paradox of decomposition, whereby if we subtract an atom at a time from a table, the latter does not cease to be a table; and yet, once the last atom has been removed, we still have to say that a table is there. According to the nihilist, we get out of the paradox by denying that there are any tables—or ordinary objects—at all (Unger 1979, p. 120ff; Wheeler 1979, pp. 164–168). Secondly, the nihilist approach values *ontological parsimony*, which emerges from the conviction of some nihilists that "for *macrophysical objects*, to be is to have causal powers" (Merricks 2001, p. 81). Merricks argues that composite objects would exist as causally irrelevant epiphenomenal entities, since their alleged effects are exhausted by the effects produced by the joint action of their parts. If we accepted their existence over and above the existence of their parts, we would be violating the requirement of ontological parsimony—among others. Finally, Nihilism is *opposed to a methodology based on common sense*, where ontological conclusions are drawn from ordinary perceptual experience (Merricks 2001, pp. 8–9, 73). This common-sense methodology is usually rejected through so-called debunking arguments (Korman 2016), according to which our ordinary way of experiencing and of dividing up the world into objects is contingent; this implies that if our ordinary-objects beliefs were true, it would be a mere coincidence; hence, we should not believe there are ordinary objects.

Commonsensicalism defends the existence of ordinary objects, so it clearly prioritises acceptability of a theory from the folk's point of view, with the aim of achieving accord with the everyday picture of the world—a kind of "mereological sanity" (Markosian 2014, p. 96). It strives to maintain a charitable approach towards the intuitions of ordinary competent speakers of English. Indeed, Hirsch (2010, pp. 98–99) has argued that we should adopt an ontology that honours charity by making ordinary existence statements true. Furthermore, Commonsensicalism aims at resisting charges of untenability, by

proving itself a viable alternative to its philosophical rivals (Carmichael 2015, p. 476; Korman 2015, p. 3).

Finally, proponents of Universalism are at pains to *avoid ontic vagueness*. Composition is characterised as unrestricted because this would prevent that statements containing the existential quantifier or the notion of identity are vague (Lewis 1986, pp. 212–224; cf. also Sider 1997, 2001). Furthermore, proponents of Universalism are concerned with *avoiding the arbitrariness* that underlies the common-sense criteria for composition. Once we admit certain objects, like artefacts or scattered objects, it seems arbitrary to leave out others just because they look more "exotic" to us. This ultimately leads us to embrace Universalism (Rea 1998, pp. 354–355; Van Cleve 1986, pp. 144–146).

2   Bennett's Epistemicism About the Composition Debate

In a recent paper, Karen Bennett (2009) formulates an epistemological critique of the composition debate, arguing that neither theory in the debate enjoys sufficient evidential support. From this, it follows that there is little justification for each party to advocate their respective views. Bennett calls this deflationary position "Epistemicism".

Bennett's critique takes as its point of departure a specific reconstruction of the composition debate. First of all, she distinguishes between the nihilists and non-nihilists, as we do above. She then depicts the non-nihilists as trying to downplay their ontological commitment to ordinary objects; as Armstrong remarks, "Mereological wholes are not ontologically additional to all their parts" (1997, p. 12). Conversely, nihilists are portrayed as up-playing their expressive power, by arguing that talk of composite objects is still acceptable despite the fact that these objects do not exist. For example, the nihilist will allow the use of sentences like "There is a table over there" to the extent that this locution is (or should be) short for the plurally quantified sentence "There are simples arranged table-wise over there" (Merricks 2001, p. 4). Provided nihilists and non-nihilists are really engaged in this kind of

dialectic, Bennett notes that there is no significant difference between the two in a number of respects.

First, Bennett says, neither of the two approaches guarantees more *simplicity* than the other. Non-nihilists are committed to composite objects and thus offer a less simple ontology. However, although nihilists may avoid commitment about objects, they will have to admit increasingly complicated properties. For instance, if they wish to capture the content expressed by a sentence like "There are tables arranged in rows", they will have to admit complex properties like (roughly) *being-arranged-table-wise-being-arranged-row-wise*. And even if they are nominalists, they will not avoid complications, because they will bloat their *ideology,* rather than their ontology, with overly complex predicates like "being-arranged-table-wise-being-arranged-row-wise".

Second, the hard problems and puzzles that beset Non-Nihilism will reappear for Nihilism. For instance, while the non-nihilist has to deal with the question "When do atoms compose a table?" from which a number of puzzles arise, the nihilist will also have to deal with the highly problematic question "When are atoms arranged table-wise?" Similar considerations hold for the Problem of the Many, causal overdetermination and colocation puzzles (2009, pp. 66–70).

In Bennett's view, the upshot is that we have no conclusive evidence to favour either Nihilism or Non-Nihilism. The debate reaches an epistemic stalemate, which in turn motivates deflation of the dispute as one where, regardless of whether there is a fact of the matter as to which theory is the best one, we have *insufficient grounds* for declaring one account superior to the other (2009, p. 71).

Bennett's considerations should prompt us seriously to think about the epistemic status of the composition debate. However, I wish to argue that giving in to Epistemicism would be premature. Indeed, I wish to resist the epistemicist conclusion by pointing to an alternative way of looking at justification in the composition debate. In what follows, I will outline a notion of "internal justification", which will help more properly to illuminate the epistemic situation of the composition theorists.

3   External and Internal Justification

It seems that an evaluation of the epistemic status of the material composition debate can be approached from two different angles. The first is an "external" angle, coinciding with the neutral stance of somebody who acts as outside observer. The second is an "internal" angle, identified with the point of view of somebody who is engaged in the dispute and comes to the table with some clear preferences. Having traced this distinction (to be further refined in what follows), the first thesis I wish to defend is the following:

> [External-Internal] There are insufficient grounds to adopt any of the competing theories if the dispute is considered from the *external* angle. However, if an *internal* angle is adopted, there might be conclusive grounds to endorse any of the views.

It is extremely natural to read Bennett as voicing a demand for grounds which comes from an *external* angle, that is, from somebody who has no pre-determined position[3] and wishes to learn which view is best supported by the available evidence.[4]

3.1   The External Angle Is Not a "View from Nowhere"

Note that reading Bennett's argument as voicing a demand from the "external" angle does not mean ascribing to her the implausible presupposition that there is a perfectly objective perspective, a "view from nowhere" from which one can say that the evidence supports at most one position. All external evidential support requires is that there be a maximally shared "epistemic system" between the disputants, consisting of the background standards, principles and requirements that they all acknowledge and presuppose, relative to which epistemic support can be assessed. In the case of the material composition dispute, this epistemic system could be thought of as including a variety of "basic" principles (see Boghossian 2006) as well as principles that are specific to the field of philosophy and even of ontology. The basic principles

could sanction the rationality of beliefs that are acquired through methods of deduction, induction or inference to the best explanation. The philosophy-specific principles would state things like: "*Ceteris paribus*, it is rational to believe a philosophical theory T if T explains our intuitions"; or "*Ceteris paribus*, it is rational to believe a philosophical theory T if T is simple". Ontology-specific principles may include Razor-style norms, such as: "*Ceteris paribus*, it is rational to believe an ontological theory T if T is ontologically parsimonious".

We could suppose that this epistemic system is shared by the participants to the material composition dispute. If—as Bennett contends—philosophical reflection offers grounds to think that evidence does not support, say, Nihilism (or any other position) *relative to this shared epistemic system*, we may say that this licenses the conclusion that Nihilism is not "externally" justified. The same goes for the other positions in the composition controversy. In sum, the external notion of justification does not imply the attribution of an implausible and nowadays unpopular strong objectivism about evidential support. All it requires is the presence of a maximally shared epistemic system like the one I have just outlined. Having qualified my interpretation of her argument, I will henceforth assume that Bennett's objections are offered on behalf of an assessor seeking external justification in the sense just characterised.

3.2   Making Room for Internal Justification

In this subsection, I wish to resist this demand of external justification. A first step in this direction consists in showing that the justification that the nihilist or non-nihilist has for adopting her view can proceed from *internal* considerations. These internal considerations will typically specify: (i) which theoretical features should be valued the most: for instance, avoiding certain paradoxes or avoiding arbitrariness; and (ii) which theoretical virtues should have priority: for instance, parsimony, explanatory power or adequacy to intuitions. Internal considerations have an *evaluative aspect* that is not present in the shared epistemic system: they set priorities, apply an order and establish privileges between principles that, by the lights of the broader epistemic system, stand on the same footing.

Consider Bennett's point that the nihilist is taking up additional commitments, and therefore loses simplicity, when it comes to properties like *being arranged table-wise-being-arranged-row-wise*. This is supposed to make the view look as essentially on a par with Non-Nihilism, inferring from this that evidence favours none of the alternatives.

The nihilist might acknowledge this. However, as Bennett mentions, the nihilist could be a nominalist about properties, which would imply that any complication merely affects her ideology. This might be regarded as bearable cost, if the nihilist ranks ontological simplicity higher than ideological simplicity. Given this ranking of theoretical virtues, the increase in ontological simplicity might provide a conclusive reason (and therefore a justification) for the nihilist to prefer her own account, at least from the *internal* point of view that results from the particular ranking of theoretical virtues she subscribes to.

Similar considerations might hold for the non-nihilists. Some of them believe that composition occurs only in some cases, when we are faced with ordinary objects. Here it seems that a conclusive internal reason (and therefore justification) for adopting this version of Non-Nihilism is its accord with common sense. Thus, if accord with common sense receives the highest ranking among theoretical virtues, the fact that an account honours common sense provides an internal justification for embracing that account.

These considerations support the thesis that I have called [Internal-External], and license the conclusion that even though, from the external perspective, one might judge that there are insufficient grounds to adopt any of the competing theories, once each of the positions is regarded *from within* a perspective where some clear preferences and priorities are set, there may be sufficient grounds to adopt each view.

3.3　Internal and External Justification, Admissibility and Pluralism

Let us delve further into the distinction between external and internal justification. As we have already noted, all is required of external justification is that epistemic support is assessed according to a *maximally*

shared epistemic system between the disputants. Conversely, it is suffi-
cient for internal justification that epistemic support obtains relative to
a particular, *non-maximally shared* perspective. Note incidentally that,
given how the notion of "maximally shared" has been introduced, what
is maximally or non-maximally shared will depend on *contextual ele-
ments*: for instance, on the principles, standards and aims that are typi-
cal of the context where ontological debates take place.[5]

The contrast between "maximally shared epistemic system" and "par-
ticular (non-maximally shared) perspective" should not suggest that the
broader epistemic system defines what is rational in the debate, while
the particular perspectives merely identify a preference or taste—which
may be thought of as a-rational or even irrational. For it may be per-
fectly acceptable by the lights of the shared standards of rationality to
endorse the ranking sanctioned by a certain perspective. Thus, for
instance, it is deemed rationally acceptable in the ontology community
to privilege ontological parsimony over ideological simplicity, as nihilists
do. But it is also deemed rationally acceptable to prioritise adherence
to common sense, as commonsensicalists do. Or, to rank the avoidance
of ontic vagueness and arbitrariness highest, as universalists do. These
rankings of theoretical features and virtues are considered *admissible* rel-
ative to the broader rationality standards of the ontology community.
By contrast, some perspectives may express evaluations that are consid-
ered *non-admissible* by the lights of said standards, maybe because they
are viewed as idiosyncratic: for example, prioritising the funniness of a
theory. Thus, the external-internal contrast does *not* reflect a contrast
between the rational and the irrational (or a-rational); for even the eval-
uations expressed by particular perspectives may count as admissible by
the lights of the shared standards of rationality, thus enjoying at least
some degree of rationality themselves.[6]

It is important to stress that the external-internal contrast does not
give rise to a form of pluralism about the *types* of justification, in the
sense that it does not lead to the claim that justification may correspond
to substantially metaphysically different states or properties. Indeed,
on closer inspection, the view currently presented has *monistic* traits.
For it conceives both external justification and internal justification as
complying with the same model, namely one where epistemic support

obtains *relative to a certain parameter*. In the case of external justification, epistemic support is relativised to the maximally shared epistemic system where the disputants are operating; in the case of internal justification, epistemic support is relativised to a particular perspective expressing a ranking of theoretical features. Perspectives may contain an evaluative element, but they do not differ substantially from the broader epistemic system; they contain the same principles it contains (although not necessarily all of them), only they *order* them. This difference cannot amount to a substantial metaphysical distinction: whence the present account's independence of a view whereby substantially metaphysically different states or properties bring about justification.[7]

Of course, the external–internal distinction may come with a pluralistic attitude *towards perspectives*. One form of pluralism in this regard would state that there is a plurality of non-maximally shared perspectives. In the debate about material composite objects, this plurality is exemplified by the perspectives of the nihilist, of the commonsensicalist, of the universalist and so on. I will come back to this pluralistic statement later on, in Sect. 6.

To sum up, in this section, I have argued that even though there may be insufficient reasons to adopt Nihilism, Commonsensicalism or Universalism, if the dispute is considered from the *external* angle, one can envisage considerations *internal* to each view relative to which one might be justified to endorse them. The external-internal distinction has been clarified through the contrast between justification relative to a maximally shared "epistemic system" and justification relative to a perspective expressing a particular evaluative stance (e.g. a ranking of theoretical features).

4 A Closer Look at Internal Justification

I have argued that the views competing in the material composition debate can be internally justified, in the sense that there can be conclusive reasons to adopt them at least relative to the perspectives assumed by their advocates. This section takes a closer look at the notion of internal justification and its application in the material composition debate.

4.1   Defining the Scope of the Proposal

First of all, the internal conception of justification is, for the purposes of this paper, only applied to the composition debate. To be sure, relativistic forms of epistemic support have been put forward for philosophy debates by other authors. Van Inwagen (2010, pp. 27–28), for instance, suggests a form of epistemic relativism about the debate between compatibilists and incompatibilists about free will. Also, Field (2009, pp. 251, 286–289) shows how relativistic justification can dispel problems related to the rationality of induction, the possibility of justifying logic a priori or the possibility to justify specific logical laws (such as the excluded middle). Comparing these proposals with the present one would take us too far afield, though.

Secondly, the present account is linked with, but ultimately rather distant from, relativistic views of justification applied to scientific beliefs and empirical beliefs in general, which have been defended by many influential authors in the last decades.[8] Besides its application being limited to the composition debate only, the present account seems to have a narrower conception of the parameter to which justification is relativised: for it simply describes perspectives as rankings of theoretical features. By contrast, epistemic relativism about science relativises the justification of beliefs to much more complex systems, for instance methodological paradigms including principles concerning empirical observation, experimental techniques and practices, training methods and schools of thought (Kuhn 1962, pp. 10–11; Bloor 1976, p. 32). Also, epistemic relativism about empirical beliefs in general relativises justification to, for instance, socially established practices and evidence samples (Kusch 2002, pp. 154–155).

Third, it has been implicitly assumed so far that whether a belief is or not justified, in the present ontological debate, depends on it being *supported by evidence*. To gain a firmer grip on internal justification, we should turn to the relation of evidential support and to the possibility that it be cashed out as relative to a perspective, too. This does not in principle constitute a challenge, for it is nowadays widely accepted that evidential support is relative to, for instance, background beliefs or

auxiliary assumptions.[9] I propose to characterise the relation between internal justification and evidence as follows:

> [Internal Justification and Evidence] Internal justification for a theory T about material composition obtains if there is a perspective P (containing, for instance, rankings of theoretical features and virtues) against which the evidence E (consisting mainly of first-order arguments and intuitions) supports T over other alternatives.

To illustrate this thesis, consider Nihilism. The evidence E the nihilist has at her disposal consists mainly of: (1) arguments supposedly favouring her theory, like the sorites argument from decomposition, the argument from causal overdetermination, debunking arguments, etc. (briefly rehearsed in Sect. 1); and (2) intuitions supporting certain premises of her arguments. One example would be the intuition driving the decomposition sorites, which can be expressed via the claim: "A table does not stop being a table if one of its atoms is subtracted from it"; or the intuition driving causal overdetermination considerations, to be spelled out roughly as: "If the causal action of O can be equally well explained by the joint causal action of its parts, then O is causally superfluous"; or the modal intuition supporting debunking arguments: "Had we been relevantly different creatures, our experience could have divided up the world into different objects to those into which it actually divides it up". Given this evidence, [Internal Justification and Evidence] says that the nihilist is internally justified to believe that Nihilism is true if she adopts a perspective P against which E favours Nihilism over the alternatives. Let us assume that in her perspective P, the highest ranked principle is: "It is rational to believe a theory T if T is ontologically parsimonious" and that various other assumptions inform her ranking, like: "Ontological parsimony should be prioritized over ideological simplicity"; or "Avoidance of sorites paradoxes should be prioritized over adherence to common sense" and so on. Provided this perspective is in place, we may say that evidence E favours Nihilism if, in the specific case at hand, *the conclusions of the nihilist's arguments comply with the ranking she adopts* in P. The conclusion of the sorites argument from decomposition is that we should drop the assumption that there are composite

objects at all (Unger 1979, p. 121); this of course ensures ontological parsimony, complying with the nihilist's highest ranked theoretical virtue. A similar conclusion is reached by the argument from causal overdetermination, adding evidential support to the position relative to P and so on.

In this picture, we can say that if the evidence E supports T relative to P, the theorist is internally justified to believe T. It is thus possible to forge a clearly intelligible link between internal justification and "perspectival" evidential support.

4.2   The Epistemic Import of Internal Justification

Having brought internal justification into sharper focus, an important worry should be addressed, concerning whether or not this type of justification has a genuine *epistemic status*. For one might suspect that the only type of justification deserving this name is external justification, i.e. justification attained relative to the maximally shared epistemic system. The principles pertaining to the broader epistemic system may be viewed as the only *truth-conducive* ones. Particular perspectives may deviate from the shared epistemic system, thus not offering a sufficient guarantee of truth-conduciveness. So, for instance, suppose that Nihilism was true; there are no composite objects. The nihilist would then be fortunate enough that her perspective has led her to believe the truth. But if justification is internal, then the commonsensicalist and the universalist will be internally justified to believe their main theses too, despite these being false. Most importantly, they would be entitled to their particular perspectives and the beliefs issued by them, but these perspectives would *systematically lead them to believe falsehoods.* How could beliefs thus obtained be genuinely justified?

There are two possible lines of response to this objection, and therefore two ways of clarifying the epistemic import of internal justification. Which response should be chosen will depend on one's stance about the overall aim of the ontological inquiry presently under discussion. For the purposes of this paper, the choice can be left open.

Here is the first line. In the specific case of ontological disputes, it may be questioned whether the aim of the ontologist is to entertain a theory that is true, rather than a theory that is simply *adequate*. Inspired by Van Fraassen's constructive empiricism, Kriegel indeed proposes that "(a) ontology aims to give us theories which are adequate to the data of ontological inquiry", where the data of ontological inquiry are (suitably characterised) intuitions; and that "(b) acceptance of an ontological theory involves as belief *only that it is adequate to those data*". (Kriegel 2011, p. 185, my italics). In the picture currently being defended, one may argue that adequacy to the data will be judged differently according to different rankings of theoretical features and virtues, because different sets of data will be given different degrees of prominence. For example, the commonsensicalist will give absolute prominence to folk-intuitions as to what constitutes a composite object; nihilists, by contrast, will downplay the importance of such intuitions insofar as they stand in the way of ontological parsimony, and they will privilege the intuitions underlying sorites, overdetermination and debunking arguments; finally, universalists will also disregard folk-intuitions in their pursuit of non-arbitrary criteria of composition. Since, according to this proposal, the only aim of ontological inquiry is to work out data-adequate theories, then as long as adequacy obtains relative to a theorist's perspective, she will be internally justified to adopt her favoured theory *by believing that it is data-adequate (relative to her perspective)*. In this scenario, internal justification for one's beliefs would obtain because acceptance of a theory simply amounts to the belief that the theory is data-adequate (relative to a perspective), a belief for which the methods adopted by each of the competing parties may well be truth-conducive.

Here is the second potential line of response. It could be argued that in the specific realm of ontology, the aim is to believe the truth, however, there is no method of belief acquisition that offers a strong guarantee in this sense. All methods adopted by the contenders in the debate comply with a number of general, truth-conducive epistemic norms, for instance, concerning inference to the best explanation, or the relationship between theory and intuitions. Yet, besides these general epistemic norms, each position adopts some norms of its own, for example, because it prioritises different theoretical features or virtues.

These specific norms may lead the theorist to systematically believe a falsity; however, her belief would still enjoy a sufficient degree of genuine justification insofar as the norms she follows are deemed acceptable in the light of the principles of the broader epistemic system, which are (by assumption) truth-conducive. The theorist's belief would therefore enjoy what Goldman calls a *second-order* type of justification which is weaker than regular (first-order) justification, but still relevantly linked to truth-conducive methods, and such that "it makes a contribution toward attainment of an overall level of positive non-culpability or reasonableness" (Goldman 2009, p. 205).

Both lines of response show that it is possible to account for the epistemic import of internal justification, and specifically for its compatibility with (some degree of) truth-conduciveness. To be sure, both responses imply some theoretical loss: loss of the assumption that the point of ontology is getting at the truth (in the first response) and loss of a robust sense of justification (in the second response). These theoretical losses are nevertheless arguably acceptable if compared with the main theoretical gain of the proposal: rehabilitating the epistemic status of the beliefs entertained by the ontologists involved in the composition debate, as opposed to dismissing their beliefs as insufficiently justified. I will say more about this theoretical gain in Sect. 5.

4.3   Rival Accounts: Conservatism and Private Seemings

As I have been at pains to show, the model I have proposed manages to accommodate the epistemic respectability of the competing positions in the debate. Yet, this is not the only model available: Conservatism (Chisholm 1981, Kvanvig 1989) might provide an alternative account. The conservatist would say that each party to the dispute is justified in continuing to believe their favoured theory in virtue of the fact that they believe it (and until they were presented with positive grounds to doubt it). Importantly, Conservatism would need *no distinction* between external and internal justification; the competing views would be enjoying the regular type of justification sanctioned by the maximally shared

epistemic system, only acquired through an *additional method* of justification for beliefs, consisting in the (undefeated) believing of that proposition. The account would therefore appear simpler.

Despite its apparent greater simplicity, I am not entirely sure that Conservatism is a serious competitor of internal justification here. Conservatism seems ideally suited to account for the justification of certain well-entrenched beliefs we live by, such as the belief that there is an external world (see McCain 2008), or that there are medium-sized objects. These beliefs are typically with us since the very first phases of our intellectual development, and all our other beliefs cohere with them.[10] When a proposition p is so well-entrenched in our cognitive set-up, the fact that we believe p seems to justify us to continue to believe p. But what should we say of the nihilist's belief that there are no ordinary objects? Or of the universalist's belief that there are all sorts of exotic mereological sums? These beliefs are certainly not well-entrenched in the cognitive set-up of their proponents, for many of these subjects' beliefs may not even cohere with them. So, the fact that these propositions are believed hardly provides any justification for continuing to believe them. Rather, it is the fact that they are supported by good enough arguments that provides a justification (if any) to hold on to them. In a picture where justification is internal, we can state that it is because there are *internally compelling arguments* to believe a certain theory that their advocates are internally justified to believe it. This comports better with the way we think about the epistemic situation of the disputants involved in the composition debate.

Another possible alternative account would explain the justification enjoyed by nihilists, universalists, commonsensicalists, etc., by claiming that they possess different pieces of private evidence—perhaps in the form of *private seemings* (Rosen 2001), which immediately justify the attendant beliefs. So for instance, the nihilist would be justified to believe that there are no composite objects in virtue of her being struck by the insight that there only are aggregates of atoms and no tables, chairs or mountains. Analogous considerations would hold for the commonsensicalist and the universalist. Appeal to private seemings would render the distinction between external and internal justification unnecessary, for the disputants would now enjoy regular justification simply

acquired from an *additional source*—namely the mental faculty responsible for these inner insights. Plus, they would now enjoy different bodies of evidence which support different propositions.

The first reservation I have with regard to the private seemings view is that it gives a misguided reconstruction of how belief and justification for *some* of the competing theses in the debate are acquired. It simply does not seem like the nihilist comes to (justifiedly) believe that Nihilism is true *just*, or even *mainly*, because she is struck by a certain inner appearance. The belief is rather predominantly the result of argumentative efforts. My second reservation has to do with the notion of private evidence. Of course, what counts as evidence in a debate is an open question, and philosophers have considerable space for manoeuvre when it comes to settling this matter. If the material composition debate is one where private evidence has any bearing, though, this claim should be appropriately motivated. Otherwise, we should stick to the less esoteric conception according to which evidence is public and shareable through language. Internal justification goes well with this assumption; indeed, it presupposes that all ontologists share the same evidence— mainly consisting of arguments and the (linguistically articulable) intuitions that supposedly drive them—but that this evidence supports different theories relative to different perspectives. This view remains faithful to the more plausible *public* conception of evidence, while delivering an account of how epistemic justification can obtain for the different competing views.

To sum up, in this section, I have delved into the notion of internal justification, sharpening the scope of the proposal (it applies to the composition debate only), linking internal justification to the idea of "perspectival" evidential support, defending the genuine epistemic import of internal justification and opposing two alternatives to the view.

5   The Primacy of Internal Justification

So far, I have argued that internal justification can be invoked for the purpose of rehabilitating the epistemic status of beliefs about material composition. However, for all that has been said so far, an assessment of the debate in terms of internal justification may still be *as good* as an assessment in terms of external justification. What I set out to argue in the next paragraphs is that the former assessment is superior to the latter in an important respect, on the grounds that it offers a *more charitable picture* of the debate.

Performing an assessment of the debate in terms of external justification is likely to result in a rather bleak picture of the composition dispute. Suppose it is true—as Bennett argues—that Nihilism and Non-Nihilism are on a par in terms of simplicity and their treatment of problem cases. Then, the putative external observer lacks sufficient reasons to believe each of the theories; but also, most importantly, the nihilist, commonsensicalist and universalist cannot claim any reasons which are sufficient from their *own* point of view either.

However, in contrast to this reasoning, it seems possible to hold that the participants to the dispute do have sufficient reasons at least from their *own* point of view, because their theory promotes the theoretical virtues they value the most. This is the gist of an approach to the debate in terms of internal justification. Citing the conformity of the theory to their favoured requirements explains a lot; for instance, it explains why they seem to be entitled to believing that theory; or why it is not completely unreasonable for them to go at lengths to defend it in journals and conferences. It renders their doxastic and practical conduct not rationally defective, but rationally acceptable. In other words, it provides a more charitable portrayal of the epistemic situation of these subjects.

In general, a view which maximises a charitable interpretation of the beliefs, assertions and actions of the subjects under assessment should be favoured over a view which fails to maximise charity. Since an assessment of the composition debate in terms of internal justification ensures more charity in this respect, it should be preferred to the

approach in terms of external justification that has been attributed to Bennett. This does not mean that the second approach results in a false view; only that that view can be legitimately disregarded[11] based on charity-related considerations.

To be completely fair to Bennett, it should be noted that she does consider the idea that "there may be some broader theoretical grounds that can justify our choice" between Nihilism and Non-Nihilism (p. 73). Indeed, she grants that we might be justified to believe Nihilism if we managed to argue that: (a) Nihilism is compatible with strict nominalism about properties like *being-arranged-table-wise-being-arranged-row-wise*; (b) ontological simplicity is a guide to truth; and (c) ideological simplicity is not a guide to truth. However, notice that here Bennett is still presumably talking about justifying Nihilism "externally", i.e. by the lights of the maximally shared epistemic system. Yet, it is very difficult to make (b) and (c) part of the maximally shared epistemic system that the ontologists share, mainly because it is not clear to what unanimously recognised evidence one should appeal in order to show that ontological simplicity, but not ideological simplicity, is a guide to truth. Bennett's seeming concession places on the opponent the burden of establishing these "broadly theoretical grounds", and especially (b) and (c), where no convincing proof is however forthcoming. Her critique is therefore left unharmed by this concession.

The opponent of Bennett's critique can, however, be discharged of this burden once the notion of internal justification is introduced. For she can argue that we do not need to "externally" show either (b) or (c) to hold; it is sufficient for us to acknowledge that each participant to the debate holds a perspective which sets priorities between theoretical features and virtues, and it is *in light* of these set priorities that they are internally justified to believe their favoured theories.

The upshot of this section is that it is preferable to assess the material composition debate in terms of the internal justification enjoyed by its participants, on account of the greater charity ensured by this approach.

6   Pluralism about Perspectives and Methodological Relativism

The positions competing in the debate about material composition are undeniably a plurality: they include Nihilism, Universalism and intermediate accounts, spanning versions of Commonsensicalism and Organicism. As I have argued, all these positions constitute particular perspectives to which justification is relativised. These observations license what may be called a *pluralism about perspectives*.

Pluralism about perspectives is a purely descriptive claim, merely reflecting the fact that more than one theory, and more than one attendant perspective, is involved in the composition dispute. This purely descriptive claim may be combined with an *evaluative* claim concerning whether any of these perspectives is epistemically superior to the others, in the sense of being the (maximally) truth-conducive one. However, it can be shown that any attempt to establish such an evaluative claim would soon run into problems.

To illustrate the point, let us focus on the different rankings of theoretical features and virtues that each party in the composition debate arguably adopts. Let us ask: *Is any ranking epistemically better than the others*? This question would specifically concern whether any of these rankings is (maximally) truth-conducive. Suppose we set out to argue that one specific ranking is the best—for instance, the ranking that puts preservation of common sense first. We would have to invoke a further, second-order ranking (or equivalent set of considerations) in order to explain why one should rank common-sense preservation first. However, in order to explain why this second-order ranking (or equivalent set of considerations) is the best, we would have to appeal to a further, third-order ranking (or equivalent set of considerations) and so on. A regress would ensue, preventing us adequately to support the original claim of superiority.

Presuming we would not concede that adoption of any of the rankings is totally ungrounded, the only alternative left seems to be admitting that each first-order ranking circularly justifies its being the best. For instance, we should admit that the ranking where preservation of common sense is placed first appeals to the value of preserving common

sense in order to justify itself, and *mutatis mutandis* for the other rankings. From this, it follows that no ranking is better justified than the others, for all of them can be circularly justified. In turn, this conclusion seems to supply reasons to think that no ranking is epistemically better than the others.[12]

Now, I do not think that the argument just presented should be read as supporting the metaphysical claim that no ranking is epistemically better than any other. It rather supports a more modest claim, namely that given the extreme difficulty of proving one ranking epistemically better than the others, the only alternative is that of *considering* each ranking as epistemically on a par with the others. In other words, I do not see the argument just given as supporting a *metaphysical* epistemic relativism about perspectives, but rather as supporting a *methodological* epistemic relativism about perspectives, according to which we should *proceed as if* all perspective were on a par.

Combining the idea of internal justification with the methodological relativism about perspectives that I have just characterised, we obtain a form of epistemic relativism which may be stated as follows:

> [Epistemic Methodological Relativism] Each position in the composition debate is *internally* justified, i.e., it is justified relative to its perspective (a ranking of theoretical features and virtues); plus, we should proceed as if no perspective were epistemically better than the others.

[Epistemic Methodological Relativism] expresses a *metaphilosophical* view, and particularly a view about the epistemic status certain first-order philosophical beliefs. It is easy to embrace this metaphilosophical view if one is an external spectator of the debate, who has no belief about the relationship between the different perspectives or is even antecedently convinced that none of them is better justified than the others. Yet, could a theorist who was endorsing a particular perspective also embrace such a form of metaphilosophical relativism? One might think that if a nihilist, a commonsensicalist or a universalist were persuaded by [Epistemic Methodological Relativism], they should just *stop believing* their favourite views, for they would have to proceed as if they themselves were in no better epistemic position than their opponents

are. [Epistemic Methodological Relativism] would therefore seem to have the same implications of Epistemicism at least for the thinkers directly engaged in the debate.

This reasoning is flawed, however. For [Epistemic Methodological Relativism] should be read as stating the following: we should proceed as if no perspective were better justified than the others in the *external* sense, that is, by the lights of what we may identify as the maximally shared epistemic system. So, if a participant to the composition debate were to embrace [Epistemic Methodological Relativism], they would simply have to treat their ranking as not more justified than their opponents' ranking in the external sense. Arguably, though, this does *not* imply that they should stop believing their favourite view. Quite to the contrary, they may continue to believe their favourite view precisely because, according to [Epistemic Methodological Relativism], each of them would be *internally* justified to hold them, relative to the specific ranking they are adopting. For a participant to the composition debate, there would be no conflict, *either semantic or normative*, in saying: "Although I should treat no perspective in the debate as better justified than the others in the external sense, I should not stop believing my favourite theory, for I am internally justified to do so".

7   Rethinking Disagreement in the Light of Internal Justification

If justification is internal to perspectives, it seems like subjects arguing from different perspectives will not be in a position to have a genuine dispute, for the reasons invoked by one will never count as conclusive for the other. The reasons adduced by the nihilist from her perspective, call it P_n, will never suffice for the non-nihilist in order to regard the nihilist position as justified, because the non-nihilist will be adopting a different perspective, call it P_{nn}, by the lights of which the nihilist's reasons will not appear as conclusive.

The only way for the non-nihilist to appreciate the nihilist's reasons would be to *switch* perspective, and adopt criteria relative to which the opponent's reasons count as conclusive. Yet, if no perspective is

externally justified more than the others are (at least if one accepts that they are all self-justifying, as argued in Sect. 6), switching from one to the other will be a sheer matter of personal preference. Thus: (i) either the reasons invoked by each party cannot be appreciated by the other or (ii) if they can, that is because the opponent has somehow idiosyncratically switched perspective.

This picture is not as bleak as the critic depicts it. For, despite appearances, in the framework that has been canvassed so far, it is possible to abandon one account and embrace another by appealing (at least partially) to external considerations that are extraneous to idiosyncratic perspective-switching.

For starters, suppose that switching perspective was the consequence of temporally or hypothetically projecting oneself in the opponent's framework and of appreciating the opponent's reasons "from the inside". First, the process of projection leads one to identify oneself with the opponent's preferences, thus temporarily acting "as if" one endorsed a certain ranking of theoretical features and virtues. Secondly, while one is in this temporary cognitive situation, one is brought to appreciate the variety of arguments and considerations that bear on the opponent's position. One can appreciate how the development of the theory honours these background requirements, thus appreciating the *internal* link between the former and the latter.

But that is not all. Appealing to the broader, more widely shared epistemic system, the theorist can in principle also compare how a certain ontological theory T_1 performs vis-a-vis its ranking of features and virtues with how another theory T_2 performs vis-a-vis its ranking of features and virtues. That is to say, one can establish a *higher-order comparison* between the different theories. This comparison could in principle be performed in an "external" setting, that is, only assuming the epistemic principles, requirements and standards that are maximally shared between the disputants. At least in principle, nothing forbids that such an external comparison will lead the subject to conclude that one account is superior to another, and therefore to be compelled to abandon one account based (at least partly) on "external" considerations which have nothing to do with a purely idiosyncratic perspective-switching. For instance, one may in principle externally

justify abandoning Commonsensicalism for Nihilism because, all things considered, one reaches the conclusion that Nihilism better honours the key requirements of its own framework. For instance, it clearly gets rid of sorites-like paradoxes for composites; it respects the causal relevance clause and it avoids debunking considerations. By contrast, one could observe that Commonsensicalism honours ordinary intuitions, but gives rise to cumbersome composition principles.

In this scenario, it is possible to say that one has been led to adopt the opponent's view with the help of internal and, most importantly, *external* considerations. For one does not embrace the rival theory only in virtue of an idiosyncratic perspective-switching, but rather through a reflective comparison at the "external", meta-level, between the own theory's inner workings and the inner workings of its competitors.

8   Conclusion

I have argued that philosophers engaged in a dispute like the one about the existence of composites are "internally justified" to their positions: the available evidence—consisting mainly of arguments and the intuitions sustaining them—favours their view *relative to the ranking of theoretical features and virtues* that these theorists are adopting, that is to say, relative to the specific *perspective* they, but not their opponents in the dispute, are occupying. The proposed view counters the idea that none of the parties enjoys sufficient justification (Bennett 2009): this idea has been associated with an approach to the epistemic assessment of the composition debate in terms of "external" justification, that is, justification relative to the epistemic system maximally shared by the disputants. I have argued that an assessment in terms of internal justification should be favoured over an assessment in terms of external justification, on account of its *greater charity* towards the epistemic status of each party's position.[13]

Notes

1. For the sake of simplicity, I will ignore intermediate views other than Commonsensicalism, for example Organicism (van Inwagen 1990).
2. See Korman (2016) for an updated and detailed overview.
3. This person could be a philosopher who is not invested in the debate, but simply wishes to study it and adjudicate the best view "from the outside", as it were; or it could be a first-year student who has barely done any philosophy before. Details are not crucial here.
4. This reading reflects a vastly shared approach to the epistemic assessment of controversies, expressed in the so-called *Principle of Uniqueness*, whereby "a body of evidence justifies at most one proposition out of a competing set of propositions" (Feldman 2011, p. 148; White 2005). Bennett's argument plausibly endorses this idea as a background assumption with regard to the theories competing in the ontological debate about composite material objects.
5. This context-sensitivity is compatible with there being an epistemic system which is universally (or near-universally) shared by all rational beings. For the purposes of this paper, though, no appeal will be made to such a universal system. All considerations will be restricted to the context of the ontologists' discourse.
6. These remarks are merely preliminary and leave two important issues unaddressed: first, *which considerations* legitimise the particular perspectives by the lights of the broader epistemic system? And second, does one particular perspective enjoy a *greater degree* of rationality than the others? I will take up these questions again in Sects. 5 and 6.
7. One may read such a metaphysical commitment in the epistemic pluralism set forth by William Alston (2005, 1993), whereby a belief can be justified in many ways: by being true, being based on good grounds, being reliably formed, allowing to fruitfully explain the world and so on.
8. Laudan (1990) provides an accessible overview of the main issues raised by relativism in the philosophy of science. With regard to epistemic relativism about scientific beliefs, the key texts are doubtlessly Kuhn (1962), Bloor (1976), Barnes and Bloor (1982), Rorty (1981). The idea championed by all these authors is that scientific knowledge and justification are relative to sociological factors pertaining to the science community: for instance, mainstream methodological approaches or

training schools. The idea that knowledge and justification of empirical beliefs generally are relative to social factors is defended by, among others, Kusch (2002), Brandom (1994) and Williams (2007).

9. Quine (1951) is the *locus classicus* where this relativistic insight is expressed. Plus, the contemporary dominant approach to confirmation, namely Bayesianism, is relativistic as well (Douven 2011).

10. Wittgenstein would have called them "hinges". See Coliva (2010, 2015) for a detailed study of the role of hinge propositions in Wittgenstein's work, as well as for a systematic defence of an epistemology of hinges.

11. As already said, disregarding Epistemicism should not imply believing that it is false, for no argument has been given to that effect. What one can do is, rather, disregard what seems to be a straightforward *normative consequence* of the position, namely that the parties involved in the debate should refrain from believing their respective views. If the internal justification view is shown superior, it is epistemically permissible (in a relevant sense) for the parties to believe their views.

12. This argumentative structure is reminiscent of the so-called Problem of the Criterion, according to which it is not possible non-circularly to justify the choice of a certain method of inquiry (in the present case, a ranking of theoretical features). Sankey (2011, 2012) has argued that this typology of argument is intimately connected with epistemic relativism. The Problem of the Criterion in turn can be seen as a version of Agrippa's Trilemma whereby either (i) we enter an endless chain of justification, or (ii) we arbitrarily stop at some point, or (iii) we adopt circular justification. The latter option can lead to a form of relativism, as noted also by Seidel (2014, pp. 153-163).

13. Work on this article was possible, thanks to the support of the Humboldt Foundation. For helpful comments, I am grateful to the audiences at the University of Uppsala, at the University of California (Irvine) and at the University of Hamburg. Special thanks to Matti Eklund, Richard Woodward, Annalisa Coliva, Karl Schafer, Ralph Wedgwood, Yuval Avnur and Moritz Schulz.

References

Alston, W.P. 2005. *Beyond Justification: Dimensions of Epistemic Evaluation*. Ithaca: Cornell University Press.

Alston, W.P. 1993. Epistemic Desiderata. *Philosophy and Phenomenological Research* 53 (3): 527–551.

Armstrong, D.M. 1997. A World of States of Affairs. *Philosophical Perspectives* 7 (3): 429–440.

Barnes, B., and D. Bloor. 1982. Relativism, Rationalism and the Sociology of Knowledge. In *Rationality and Relativism*, ed. Martin Hollis, and Steven Lukes. Cambridge: MIT Press.

Bennett, K. 2009. Composition, Colocation, and Metaontology. In *Metametaphysics: New Essays on the Foundations of Ontology*, ed. David John Chalmers, David Manley, and Ryan Wasserman. Oxford: Oxford University Press.

Bloor, D. 1976. *Knowledge and Social Imagery*. Chicago: University of Chicago Press.

Boghossian, P.A. 2006. *Fear of Knowledge: Against Relativism and Constructivism*. Oxford: Oxford University Press.

Brandom, R.B. 1994. *Making It Explicit: Reasoning, Representing, and Discursive Commitment*. Cambridge: Harvard University Press.

Carmichael, C. 2015. Toward a Commonsense Answer to the Special Composition Question. *Australasian Journal of Philosophy* 93 (3): 475–490.

Chisholm, R. 1981. *Foundations of Knowing*. Minneapolis: University of Minnesota Press.

Coliva, A. 2010. *Moore and Wittgenstein: Scepticism, Certainty, and Common Sense*. London: Palgrave-Macmillan.

———. 2015. *Extended Rationality: A Hinge Epistemology*. London: Palgrave-Macmillan.

Douven, I. 2011. Relativism and Confirmation Theory. In *A Companion to Relativism*, ed. Steven Hales. Malden: Wiley-Blackwell.

Feldman, R. 2011. Reasonable Religious Disagreement. In *Social Epistemology: Essential Readings*, ed. Alvin Goldman, and Dennis Whitcomb. New York: Oxford University Press.

Field, H. 2009. Epistemology Without Metaphysics. *Philosophical Studies* 143 (2): 249–290.

Goldman, A. 2009. Epistemic Relativism and Reasonable Disagreement. In *Disagreement*, ed. Richard Feldman, and Ted Warfield. Oxford: Oxford University Press.

Hirsch, E. 2010. *Quantifier Variance and Realism: Essays in Metaontology*. Oxford: Oxford University Press.

Korman, D.Z. 2015. *Objects: Nothing Out of the Ordinary*. Oxford: Oxford University Press.

———. 2016. Ordinary Objects. *Stanford Encyclopedia of Philosophy.*

Kriegel, U. 2011. Two Defenses of Common-Sense Ontology. *Dialectica* 65 (2): 177–204.

Kuhn, T.S. 1962. *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.

Kusch, M. 2002. *Knowledge by Agreement: The Programme of Communitarian Epistemology*. Oxford: Oxford University Press.

Kvanvig, J. 1989. Conservatism and its Virtues. *Synthese* 79: 143–163.

Laudan, L. 1990. *Science and Relativism: Some Key Controversies in the Philosophy of Science*. Chicago: University of Chicago Press.

Lewis, D.K. 1986. *On the Plurality of Worlds*. Oxford: Blackwell Publishers.

Markosian, N. 2014. A Spatial Approach to Mereology. In *Mereology and Location*, ed. Shieva Keinschmidt. Oxford: Oxford University Press.

McCain, K. 2008. The Virtues of Epistemic Conservatism. *Synthese* 164: 185–200.

Merricks, T. 2001. *Objects and Persons*. Oxford: Oxford University Press.

Quine, W.V.O. 1951. Two Dogmas of Empiricism. *Philosophical Review* 60 (1): 20–43.

Rea, M.C. 1998. In Defense of Mereological Universalism. *Philosophy and Phenomenological Research* 58 (2): 347–360.

Rorty, R. 1981. *Philosophy and the Mirror of Nature*. Princeton: Princeton University Press.

Rosen, G. 2001. Nominalism, Naturalism, Epistemic Relativism. *Noûs* 35 (s15): 69–91.

Sankey, H. 2011. Epistemic Relativism and the Problem of the Criterion. *Studies in History and Philosophy of Science Part A* 42 (4): 562–570.

Sankey, H. 2012. Scepticism, Relativism and the Argument From the Criterion. *Studies in History and Philosophy of Science Part A* 43 (1): 182–190.

Seidel, M. 2014. *Epistemic Relativism. A Constructive Critique*. London: Palgrave Macmillan.

Sider, T. 1997. Four-Dimensionalism. *Philosophical Review* 106 (2): 197–231.

———. 2001. *Four Dimensionalism: An Ontology of Persistence and Time*. Oxford: Oxford University Press.

Unger, P. 1979. There Are No Ordinary Things. *Synthese* 41 (2): 117–154.

Van Cleve, J. 1986. Mereological Essentialism, Mereological Conjunctivism, and Identity Through Time. *Midwest Studies in Philosophy* 11 (1): 141–156.

van Inwagen, P. 1990. *Material Beings*. Ithaca/London: Cornell University Press.

———. 2010. We're Right. They're Wrong. In *Disagreement*, ed. Richard Feldman and Ted Warfield. Oxford: Oxford University Press.

Wheeler, S.C. 1979. On That Which Is Not. *Synthese* 41 (2): 155–173.

White, R. 2005. Epistemic Permissiveness. *Philosophical Perspectives* 19 (1): 445–459.

Williams, M. 2007. *Groundless Belief: An Essay on the Possibility of Epistemology*, 2nd ed. Princeton: Princeton University Press.

Index