



Endre Begby

“From Pragmatist Anti-Metaphysics to Non-Ideal Epistemology?”

In: Yvonne Huetter–Almerigi and Robert Sinclair (Eds.) (2026).
Pragmatism, Metaphysics and Method—Essays for Bjørn Ramberg (pp. 51–73).
Nordic Studies in Pragmatism 5. Helsinki: Nordic Pragmatism Network.

ISSN-L 1799-3954

ISSN 1799-3954

ISBN 978-952-67497-4-7

Copyright © 2026 The Authors and the Nordic Pragmatism Network.

This work is licensed under a Creative Commons
Attribution-NonCommercial 3.0 Unported License.

 For more information, see
<http://creativecommons.org/licenses/by-nc/3.0/>



Nordic Pragmatism Network,
Helsinki 2026

www.nordprag.org

From Pragmatist Anti-Metaphysics to Non-Ideal Epistemology?

Endre Begby
Simon Fraser University

1. Introduction: a personal note

I believe I first met Bjørn at the beginning of the Fall term of 1998: I was a Master's student at the University of Oslo and Bjørn was a visiting professor on leave from his regular position at Simon Fraser University in Canada. But I already knew about him, of course: after all, he was a former student from the Oslo department who had since gone on to make a name for himself in the contemporary philosophical landscape. I had already taken out his book on Davidson from the library, and had copied every page, every chapter, on the second-hand photocopier that my roommate had recently purchased and kept in our flat. I think it would be fair to say, I was already a fan.

There is an old saying that one should never meet one's heroes. But meeting Bjørn in person did not disappoint. Everyone who has met him knows that he is exceptionally charismatic. But they may also know that he is unfailingly kind, generous, attentive, and supportive. As a student, I was lucky enough to be able to follow, I think, two of his courses that year: one on pragmatism, another on Gadamer's hermeneutics. I can't remember what I wrote for my term papers, but I suppose it was sufficiently competent that, when time came around for him to pack up and return to Canada, he approached me and said that if I were interested, he would be happy to supervise my Master's thesis from a distance. Naturally, I was thrilled by the idea, and gladly accepted the offer.¹

¹ I mention this because in the meantime, Bjørn is back at University of Oslo while I am a Professor at Simon Fraser University, teaching many of the same classes that I imagine he taught when he was here.

So began an incredibly rewarding relationship stretching to this day, during which Bjørn has served variously as supervisor, mentor, co-author, career counsellor, life coach, colleague, and finally—I would like to think—as a friend. I don't think there's any other academic mentor who has meant more to me personally and professionally than Bjørn, both in giving shape and direction to my own philosophical efforts but also in helping me gain a sense of the sort of Philosopher-person I should aspire to be, in particular, in terms of one's interactions with students and colleagues. I am immensely thankful to him for this. And so I was very happy to accept this invitation to pay homage to Bjørn and to put some thought into the question of how his philosophical contributions have shaped and inspired my own.

2. Bjørn's "golden period"

From the mid 90s and well into the 2000s, Bjørn published a remarkable series of papers. These papers helped consolidate my pre-existing philosophical preoccupations but also introduced me to new topics and perspectives. In addition, they showed me the way to a certain kind of constructive and forward-looking way of doing philosophy that appealed to me enormously and that I aspired to emulate in my own work.

In terms of presentation, however, these papers often seemed quite modest at first glance: they were often framed as commentaries on the works of older contemporaries—Rorty, Davidson, Gadamer, Dennett, in particular—offering constructive re-readings and friendly amendments along the way.²

Nonetheless, it was clear to me even then that their modest presentation belied their potential philosophical significance. In particular, they seemed to show how certain arguments, positions, ways of thinking, could be illuminating well beyond the debates in which these arguments were initially proffered. (Many of these papers were published in volumes where the target authors also provided responses: it was always clear from these responses that the authors in question, like me, found that Bjørn's unassuming commentary had revealed to them aspects and dimensions of their own thinking that they themselves had not yet fully realized. For instance, Davidson writes, in response to "The Significance of Charity" (Ramberg, 1999), that "this is the kind of criticism that makes doing

² A matter of "regional pixel adjustment," as he memorably puts it in Ramberg, 1997, p. 459.

philosophy worth the trouble [. . .]. [Ramberg] has shown me a better way to put together some of the strands in my own thinking, and gently suggested what I would do well to leave out” (Davidson, 1999, p. 619). A similarly effusive note is struck by Rorty, in response to “Post-Ontological Philosophy of Mind” (Ramberg, 2000): “Most of my responses in this volume are, at least to some extent, rebuttals. But in the case of Bjørn Ramberg’s paper, I find myself not only agreeing with it, but also very much enlightened by it” (Rorty, 2000, p. 370).³

Perhaps due to my educational background at the University of Oslo, I was particularly impressed with Bjørn’s willingness to engage productively with philosophical issues spanning what was then known as the “analytic-continental divide.”⁴ While Bjørn’s philosophical outlook seemed solidly grounded in analytic philosophy (though certainly tending toward a kind of “naturalism,” more about which later), it would have been hard to emerge from a philosophical education at the University of Oslo—its curriculum shaped in no small part by the legacy of Dagfinn Føllesdal—without feeling that at least *some* of the questions typically associated with the works of, say, Edmund Husserl, Martin Heidegger, Maurice Merleau-Ponty, or Hans-Georg Gadamer—questions about human socio-historical situatedness and cognitive finitude—were questions worthy of serious philosophical consideration.

Bjørn’s papers from this era seemed to me just the ticket for an aspiring philosopher in my predicament. They ask serious questions—not simply critical questions, but probing and constructive questions—about how we can best integrate these humanistic insights and perspectives into the starker and leaner vocabulary of analytic philosophy. They take for granted that these are philosophical concerns worthy of serious consideration, no matter what one’s philosophical starting point or preferred methodology. But they are not simply trying to graft one perspective onto the other or to point out similarities between competing philosophical traditions, unlike much of superficially similar work from that era, which seemed content to simply list a number of commonalities between the Analytic and Continental approaches and leave it at that. By contrast,

³ Subsequently, he concedes that “Ramberg has persuaded me to abandon two doctrines which I have been preaching for years: that the notion of ‘getting things right’ must be abandoned and that ‘true of’ and ‘refers to’ are not word-world relations” (Rorty, 2000, p. 375). Anyone who knows Rorty’s work from this period will know that these are hardly peripheral points in his overall philosophical worldview.

⁴ Apparent already in Ramberg, 1988 and Ramberg, 1989, chs. 9–10

Bjørn seemed to me to be probing for actual transcending insight, a true *Horizontverschmelzung*.⁵

These papers range across a variety of topics, but I take it that their central concern is to articulate a viable picture of human agency and its place in the natural world.⁶ But I don't think I'm overreaching in assuming that they were also animated by the idea that facts about human historicity and social embeddedness really shouldn't be considered alien to this more naturalistic approach: rather, a distinctive feature of the human situation is that our sociality is, as it were, part of our nature.

From the outset, though, it was clear that the works of Richard Rorty held a special place for Bjørn. The inspiration from Rorty was both substantive and methodological. Substantively, Rorty had already in *Philosophy and the Mirror of Nature* (Rorty, 1979) broached the ambition of bringing a "hermeneutic" perspective to bear on historical trends and developments in philosophy. Methodologically, Rorty proposed that "pragmatism" was the appropriate antidote to some of the counter-productive tendencies on display in modern analytic philosophy: i.e., its tendency to insulate itself and to prefer inquiry into the highly technical questions that supposedly constitute the domain of "proper philosophy." On Rorty's perspective, these questions are unrecognizable to people outside the discipline. They may well also turn out to be "pseudo-questions" that don't really permit satisfactory answers at all, however much we might admire the technical acumen of the various proposed solutions. Far from the "perennial questions" that philosophers like to think of themselves as occupied with, these problems—the mind–body problem, the problem of knowledge, etc.—may simply be artifacts of particular, historically situated philosophical framing.

Substantively, I have always thought Bjørn's body of work during this period constituted a considerable advance on Rorty's. Methodologically, however, I must admit that even back then, it was never entirely clear to me that what declaring oneself a pragmatist was supposed to amount to. In what follows, I will try to elaborate on both of these.

While Rorty's early work was certainly ground-breaking and inspirational, his subsequent efforts seemed to me to remain programmatic at best. After *Mirror*, Rorty rarely seemed motivated to engage in the "dirty

⁵ Gadamer, 1960. See, for instance, Ramberg, 2003, for warnings against "easy assimilations" of the two perspectives.

⁶ Bjørn would later amend his thinking to also include a pre-occupation with non-human agency. This was also very influential on me. On this, see, e.g., Ramberg, 2005.

work” of digging into texts, situating perspectives in ongoing research, fully articulating what we might draw from each to paint the kind of integrated picture of human agency that we—so-called “pragmatists”—might aspire to. Bjørn’s papers, by contrast, did all of that in spades: moreover, its manner was always that of deep, constructive engagement, rather different from the sort of sloganeering that Rorty’s work occasionally devolved into.

A central thread running throughout much of this period in Bjørn’s work is the question of how we can aspire to a broadly “naturalistic” account of human agency without going by way of the sort of metaphysical trappings to ineluctably lead to questions about reductionism or eliminativism. This was an issue of contention between Bjørn’s two primary philosophical lodestars during this period, Davidson and Rorty. As is well known, Davidson had argued that the kinds of explanations we offer when we talk about people’s beliefs, desires, intentions, and so on, are irreducible to the kinds of explanations that we offer when we talk in terms of physical events.⁷ Nor can there be any strict generalizations that span the two “vocabularies,” the mental and the physical. So while we can remain committed to some form of ontological monism—minds and bodies are not two sorts of things—the internal logic of the way that we think and talk about these categories remains irreducibly distinct.

While deeply appreciative of many strands of Davidson’s thinking, Rorty repeatedly expressed reservations about this distinction. At issue is not whether the mental might turn out to be reducible to the physical after all, but rather why Davidson remained so peculiarly committed to the vocabulary of the mental in the first place. On Rorty’s view, we should, as naturalistically inclined pragmatists, remain open to the possibility that our category of the mental is just another one of those philosophical fictions that we may eventually learn to live without.

Bjørn’s attempt to navigate and negotiate these disputes (in particular Ramberg, 1999; 2000) are, I think, deeply illuminating pieces of philosophy. In essence, he proposes that we should grant the following point to Rorty: while the basic argumentative strategy that Davidson has developed certainly *applies* to the relationship between the mental and the physical, it may also turn out to apply more widely, for instance to the relationship between evolutionary biology and “the physical.” In this sense, Rorty is right that the “mental” is not particularly special. As Bjørn summarizes

⁷ See, in particular, Davidson, 1970.

the point: “The quality of being nomologically irreducible to physics is not a particularly distinguished or distinguishing property” (1999, p. 604).

Bjørn urges Davidson to take this point to heart: one of Davidson’s “great achievements,” he argues, is to have “shown us how to dissociate reduction and ontology, thereby greatly reducing the philosophical significance of either pursuit. It would be a pity [...] if this liberation were confined to our talk of the propositional attitudes. Not with respect to the mental, nor any other vocabulary, is there a special, philosophical task which consists in tracing its conceptual relations to an ideal basic science, for then to pronounce upon its ontological legitimacy” (1999, p. 608).

Nonetheless, he likewise cautions Rorty against thinking that this shows that he was right all along that the mental is not “special” or in any noteworthy sense, a “fundamental” category. As I understand him, what Bjørn has in mind here is the sense in which Rorty’s pragmatism remains caught up—very much despite itself—in a recognizably metaphysical way of thinking about what it would take for a vocabulary to be “fundamental” in this sense:⁸ if it doesn’t point us toward ontologically distinctive categories, then it is not particularly “special.”

Against this, Bjørn presses, persuasively as I see it, for a naturalistic perspective on which the specialness of the mental vocabulary is due to facts about us, not facts about “the world” as seen from some more metaphysically austere point of view: quite simply, humans “are built to be sensitive to the sort of pattern that norm-tracking agency-predicates reveal.” Even as scientific inquiry probes ever deeper into the cognitive and neurological basis of this form of agency-attribution, the vocabulary of agency “is not one we could choose to lose.”⁹ This is a good thing, surmises Bjørn, and I agree: “Given the evident value of sensitivity to these patterns, perhaps we should, in a Humean spirit, be thankful that nature has not left it up to us to decide whether or not to preserve it” (1999, p. 615).

I am particularly taken with one strand of argument in support of this view of the irreducibility of the mental. Fundamentally, what gives content and structure to our thoughts in the first place is a causal triangle

⁸ Perhaps a hangover from his early days gleefully espousing a type of eliminativism: see Rorty, 1965; 1970.

⁹ See also Ramberg, 2000, p. 362: “The norms of agency [...] are not norms that we can hold up before ourselves or others as directives or guides to behavior. *That* we generally conform to them is what makes us language-users, and so thinkers and knowers. *Whether* in general to conform to them is not a question of subjective choice at all.” Differently put, perhaps: these are constitutive norms, not guidance-norms.

that connects us, the world, and other minds.¹⁰ What I like about this is that it attempts to locate a place for intersubjectivity—or more broadly, sociality—into our “natural” cognitive design. Some of the arguments in support of this view are clearly speculative, as they must be. But they are nonetheless very suggestive, and, I think, in large part borne out by subsequent research into cognitive psychology. There is no point in our cognitive development where, by way of some kind of intellectual feat, we individually come to realize that there are “other minds” out there similar to our own. To be sure, there are various developmental milestones along which we come to deepen our understanding of these other minds.¹¹ But from very early stages of development, it is clear that our conception of an objective world is one that already has other minds, other agents, in it. Bjørn writes: “The basis of knowledge, any form of knowledge, whether of self, others, or the shared world, is not a community of *minds*, in the sense of mutual knowledge of neighboring belief-systems ... Rather, it is a *community* of minds; that is, a plurality of creatures engaged in the project of describing their world and interpreting each other’s descriptions of it” (2000, p. 361–2).

I draw two deep inspirations from this way of thinking. First, we should feel no pressure to “validate” the vocabulary of agency in terms of some more “fundamental” metaphysical picture of the world, specifically a picture of the world that might not have a special place for human beings (or other sentient creatures) in it at all. What validates the mental vocabulary are simply facts about human nature. Second, we must recognize that human sociality is part of human nature: in other words, there are no “minds” without “other minds”; a baseline of social responsiveness is embedded in our picture of “the world” from the start. Bjørn’s “naturalism,” then, is divorced from any vestige of metaphysics: we are fully invited to view these various explanatory vocabularies as “artifacts of the bio-evolutionary and, eventually, cultural history of natural creatures with aims befitting their natures” (2001, p. 233).

I find this form of relaxed, anti-metaphysical naturalism tremendously appealing and indeed inspiring. But what does it have to do with “pragmatism”? I take it the answer will primarily point to pragmatism’s instrumental value in helping “breaking the spell” of the kind of metaphysical framing that questions of human agency (as a locus of subjective experience) had apparently become tangled up in at the time. Pragmatism, at

¹⁰ See, in particular, Ramberg, 2001, drawing on Davidson, 1991.

¹¹ See Begby, 2023a.

least as Rorty tells the story, does not set up to offer new answers to the old questions of whether human subjectivity and human cognitive capacities are “reducible” (in this or that sense) to some more “fundamental” physical descriptions of the world. To the contrary, the assumption that these are fundamental questions is just an artifact of a particular metaphysical cast of mind: it’s a byway that philosophers might—and will—get trapped in, once their preferred philosophical methodology becomes sufficiently insular and divorced from what goes on in other areas of inquiry and from questions that concern human lived experience. Pragmatism counsels us to simply move past these questions, suppressing our educationally inculcated impulses to think that “progress” in philosophy cannot be made until we have been given satisfactory answers to these “perennial questions.”

I like the overall arc of this narrative. It was certainly hugely influential on the early parts of my own philosophical career. But even at that time, I had significant concerns about the invocation of “pragmatism” (particularly in the guise of “American pragmatism,” i.e., Peirce, Dewey, and James to make the point). The more I read about these figures, the more pragmatism seemed to create its own set of distinctively theory-engendered philosophical problems (perhaps even pseudo-problems), and that perhaps “pragmatism” was, in this sense, really not fundamentally better than the metaphysically-tinged theories that it was meant to challenge. (In particular, I have never had any time for the endemic debates about what the “true” form of pragmatism is, the Peircean, the Deweyan, or the Jamesian: one would be hard pressed to find anything less “edifying,” as Rorty might say, anything less indicative of a “pragmatic mindset,” than self-declared pragmatist philosophers debating what should be the fundamental philosophical commitments of “Pragmatism”.)¹²

On the other hand, “being right about Pragmatism” never seemed to be particularly high on Bjørn’s philosophical agenda. And rightly so: neither methodologically nor substantively do his central philosophical concerns seem intrinsically or particularly indebted to Pragmatism as such, as contrasted, say, with that older tradition of science-inspired philosophical humanism that connects figures like David Hume, Thomas Reid, Auguste Comte, and John Stuart Mill (though plausibly *also* the American pragmatists).¹³

¹² Though see Andersen and Mitchell, eds. 2023, for a recent and constructive attempt to revive a “pragmatist metaphysics” in the context of philosophy of science.

¹³ Cf. Begby, 2014.

Accordingly, I consoled myself (as I suspect did Bjørn) with the idea that there was a workable sense of “pragmatism” on which pragmatism was not to be thought of as a superior “theory” or “philosophical methodology,” but simply as an outlook that ranked the urgency philosophical questions in terms of their ability to speak, in some broad and admittedly vague sense, to the kinds of problems that we, as human beings, take ourselves to be preoccupied with at a particular time, and moreover, which encouraged philosophers to give consideration to the widest possible range of source-materials to guide their inquiries. This would be a kind of outlook that permitted us to brush aside concerns, say, about how explanations of worldly events in terms of beliefs, desires, and intentions ultimately differ from explanations articulated in the standard vocabulary of natural science. We would be permitted to brush them aside not because we could claim to *know* that these questions were ultimately ill-grounded. Instead, “pragmatism” gave us the courage to resist the supposition that these should be the *first* questions on the philosopher’s docket, and that we had no right to engage in any other questions until we had answered them.

But even as I have eventually found some manner of peace with the “Pragmatism-wars,” I admit that there’s a certain sense of loss of context when revisiting these discussions now. For even this kind of minimal engagement with metaphysical questions about mind, meaning, and agency in the natural order now strikes me as a bit quaint. While Bjørn’s papers had a huge impact on the early to middle stages of my philosophical career, I admit that these framing questions have barely even crossed my mind since I found my own philosophical footing. It seems, then (at least from my own perhaps provincial point of view), that the philosophical landscape has changed in important ways: the kind of outlook that Bjørn’s views were defined in opposition to—e.g., various kinds of physicalism and reductionism—are just not central to the way that current philosophical inquiry is carried out anymore. Now of course, this is good news: something to celebrate rather than to mourn. But no one, least of all someone like Bjørn, should ever think that these developments have come about because professional philosophers were finally convinced by the strength of the arguments for pragmatist anti-metaphysics. Instead, I would like to think that they came about because more philosophers started *doing*—rather than just *talking about doing*—the kind of inspirational work in question: i.e., work that takes up questions of central human concern, but addresses them in an eclectic and inclusive manner, not

overly concerned with whether what they are doing is “really philosophy” (let alone whether it is “pragmatism”).¹⁴

As noted, these are happy developments. But even so, it seems fair to ask, given this loss of context, is Bjørn’s work from this period relevant anymore? Or is it more like the ladder we threw away to get to where we are now? I think it remains relevant, and I believe I can certainly trace its influence in my own current work.

3. From pragmatist anti-metaphysics to non-ideal epistemology

In the last decade or so, I’ve devoted most of my philosophical efforts to epistemology. In particular, I’ve been exploring questions about human information-gathering and information-processing through the lens of developments in social epistemology (branching into situated epistemology, standpoint epistemology, etc.), also drawing inspirations from psychology, sociology, and cognitive science. Most relevantly, perhaps, I’ve come to develop an approach to these issues which I have broached in terms of “non-ideal epistemology,” or more recently, “epistemology for finite minds.”¹⁵

Now, epistemology was never a particularly central concern in Bjørn’s writings, however much it might have come up in passing. Growing up as a philosopher, epistemology was never my central concern either. When I was in graduate school, the default framing for epistemology was still the Gettier-problem, along with more recent developments centered on concepts such as “safety” and “sensitivity,” all with its standard suite of increasingly esoteric examples and counterexamples. The influence of Timothy Williamson’s contrary “externalist” perspective—on which the concept of knowledge is primitive and irreducible, so that we don’t even have to answer the question of what needs to be added to “justified true belief” to have knowledge—was just starting to be felt.¹⁶ Neither seemed to me to be very interesting: the reason they didn’t seem interesting is precisely that they struck me as tending toward insularity. As the philosophical analyses on offer became more sophisticated, they also became more esoteric, more divorced not just from broader worldly concerns but also from the rest of philosophy, and from the rest of scientific inquiry. The ap-

¹⁴ One particularly important such figure is, of course, Daniel Dennett. See, e.g., Dennett, 1991, of which Bjørn was an early and avid reader.

¹⁵ Begby, 2021*b*; Begby, 2023*b*.

¹⁶ Williamson, 2000.

pearance of rigour was generally bought and paid for by distinctive kinds of idealizations, which invariably rendered the resulting analyses not particularly representative of the kinds of epistemic challenges that most of us are confronted with in our daily lives. The pragmatist leanings I had inherited from Bjørn counselled me to stay far away from these areas.

Shortly after I took up my current position at SFU, I was assigned to teach an undergraduate epistemology course. In preparing for this course, I was thrilled to discover a slew of more recent developments centering on social epistemology, relating to questions about testimony, trust, communication, about peer disagreement, and about the interrelation between the social status of “knower” and prestige hierarchies in society at large. This seemed to me to provide a much more interesting and relevant template for epistemology, connecting not only with other areas of philosophy, but also with sociology, psychology, linguistics and biology, as well as with social and political issues in the broader world. And so I was suddenly motivated to try to find my own way in this emerging field of study.

Under this guise, I believe epistemology has emerged from the doldrums of a theoretical framing which separated it from other, potentially relevant areas of concern, along with the metaphysical trappings that inevitably arise from that framing: what is knowledge *really*? Does knowledge confer distinct normative privileges from “mere” justified true belief? etc.

Instead, epistemology can now be seen as aspiring to provide a potentially integrating account of human agency at large: after all, it ties together questions of what we do to provide ourselves with information, how we process this information, how this processing feeds into decision making and action, how it is communicated to others in social learning and shared deliberation, etc.

I soon got into thinking that there was a systematic outlook worth articulating here. “Non-ideal epistemology” was already a coinage that had been floating around in the literature, but its moorings seemed to me unclear. (In particular, invocations of “non-ideal theory” in the broad vicinity of epistemology—often influenced by the pioneering work of Charles Mills—were highly politicized and seemed to be framed primarily as a criticism of the sort of “ideal theory” outlook famously espoused by John Rawls. To be clear, I was perfectly happy to think that epistemological theorizing should eventually also serve as input to an account of political agency (cf. Begby, 2022; forthcoming; in progress). But I also thought that many of the idealizations distinctively at play in epistemology were

of broader concern than that: they were the kinds of idealizations that would hamper *any* account of agency long before we would get to the political level.)

What I tried to do, then, was to articulate the rudiments of a systematic account, starting from the recognition of epistemology as a normative inquiry, of how giving proper due to human cognitive and situational constraints should shape the approach we take to questions pertaining to information-gathering and information-processing. I proposed to call these, respectively, “endogenous” and “exogenous non-ideality”: the former, roughly speaking, denotes the relatively stable limitations to the inherent information processing capacities of the human mind, as we know it; the latter, the highly contingent and fluctuating limitations to the information that is available to us at a particular time, given the social and historical situation we happen to find ourselves in (Begby, 2021 *b*). Both are fundamental constraints on human epistemic agency, and in fact, on any agency at all. By introducing idealizations that fail to reflect these constraints, traditional epistemology has needlessly set itself adrift from any connections it might have to, for instance, cognitive science, developmental psychology, or sociology. In doing so, it has effectively *made it the case* that the kinds of questions philosophers concern themselves with will be *esoteric questions*, of no apparent relevance outside our discipline.

Looking back while preparing this essay, it now dawns on me that my distinction between “endogenous” and “exogenous” non-ideality was really an attempt to tease apart the two major perspectives on human agency that I had first started thinking about after reading Bjørn’s work, namely the biological-naturalistic perspective and the socio-historical perspective. But not only that: I now also recognize the fundamental importance of an idea that resonates throughout Bjørn’s work. From one point of view, it is certainly natural to describe these (as I did above) as *constraints* on epistemic agency. But from another point of view, it seems more relevant to describe them as *enablers*. In brief: these facts constrain our theoretical accounts of human epistemic agency. But with respect to the epistemic activities that we are trying to theorize about, they are clearly enablers. Differently put, while our natural cognitive capacities and our socio-historical setting are inevitably limited, they nonetheless offer us rich interfaces with the world, essentially providing the parameters of our worldly embodiment and context, and thereby defining human agency as we know it.

One way to ease into this perspective is to start by considering what an “ideal epistemology” might even amount to. Traditional epistemology was never entirely transparent about its idealizations, but it was always clear that it was imposing idealizing assumptions from the start. In particular, whenever it was recognized that actual human epistemic agents are often biased or inattentive, it was immediately assumed that these tendencies should not be reflected in our best theories. From the point of view of ideal epistemology, this just means that their performance is decidedly sub-standard from a normative point of view: we’re entitled to expect more from them, even if we recognize that most people fall short of these standards. (This is not a problem, they argue, but rather intrinsic to any form of normative theory: after all, we don’t usually take it to be a problem for philosophical accounts of morality that people often act in violation of moral norms, whether due to ignorance or weak motivation.)

But what if some of the cognitive patterns that manifest from a certain idealized point of view as bias or inattention are really something else? To see how this might be the case, we can try to imagine what a “fully idealized” epistemology might look like. Going down this path quickly reveals, I believe, that the notion of a “fully idealized epistemology” cannot even be given a coherent description. Epistemic norms are norms that are supposed to govern inference from evidence. But inference from evidence is inherently a “risky strategy,” epistemically speaking: that is, it ineluctably carries with it a non-zero chance of ending up with false beliefs.¹⁷

By contrast, at the limits of idealization, we might imagine something like Leibniz’s omniscient God. The problem is that in virtue of being omniscient, Leibniz’s God is just not in the business of drawing inferences from evidence at all. And so, by definition, norms governing the transition from evidence to belief just don’t apply to the case. In short, relying on *evidence* to form *beliefs* is, inherently, something that only a non-ideal agent would think to do. And so, any form of normative, evidence-based epistemology, will have to recognize that it is already a non-ideal epistemology.¹⁸

¹⁷ Obviously, I’m not enamoured by the claim—taken to be axiomatic in contemporary “knowledge-first” approaches to epistemology—that one’s evidence, properly understood, is simply what one knows, and so that there is no risk of running into “false evidence.” To my mind, this is just another idealization, now masquerading as a metaphysical insight. Moreover, it doesn’t do anything to solve our problem, which now simply transposes to the question of how epistemic agents can know whether what they have is evidence (as opposed to something that merely looks like evidence).

¹⁸ For more detailed elaborations on these arguments, see Begby, 2021b; 2024.

If this is right, then we cannot really approach the central questions of epistemology without curtailing our idealizations in important ways. This is not to say that idealizations have *no* role to play in normative epistemology. Rather, it amounts to asking, which idealizations are the right ones (Ramberg, 2004)? Again, we'd struggle to find any clear recognition, let alone a sustained discussion of these points in the traditional literature. But surely it matters that what we are doing is *human* epistemology, and not, say, chimpanzee epistemology or Universal Turing Machine epistemology. The fact that we are doing human epistemology might seem obvious (it certainly does to me), but its repercussions are rarely recognized in the literature. Because if we are, then the normative picture that we are striving to articulate must in some important sense be relativized to a realistic (or if you will, "naturalistic") picture of distinctively human cognitive capacities. After all, human beings are information processing systems operating in real-time. They display distinctive limitations on short- and long-term memory, on perceptual and attentional capacities, etc. This is what I call "endogenous non-ideality." I think it imposes real constraints on our best account of epistemic normativity: it means we're no longer free to just assume that subjects have virtually unlimited capacities to absorb, process, store, and retrieve facts about their environment. For example, a subject faced with a multiple object tracking task cannot be faulted, epistemically speaking, for failing to note that a dot that started out in the upper left quadrant of the screen ended up in the upper right quadrant (even though all the evidence they needed to complete the task was in some relevant sense "right there" on the screen, plainly within their visual field). The task itself is designed precisely to demonstrate the limits of the human attentional capacities. More generally, much of what gets routinely classified as "bias" or "neglect" in the epistemological literature may turn out to be expressions of the kind of cognitive economizing we should expect to see at work in limited ideal agents.

To repeat: as a normative science, there's clearly a role for idealizations in epistemology. But the idealizations in question must be relativized to what we independently know about the upper limits of human cognitive capacities. There's still ample room to criticize people whose prejudices or intellectual laziness leads them to disregard readily available evidence. But there's no room for criticizing people for not being able to process *all* the evidence, or for failing to recognize evidential relations far beyond anyone's computational horizon.

Similar considerations apply to what I call the “exogenous” dimension of non-ideal epistemology. Social epistemology is borne from the recognition that human epistemic agency is deeply socialized from the earliest stages of cognitive development. This puts human beings at a tremendous advantage over members of other species: quite simply, the range of facts that we are capable of knowing—facts about history, about geography, about the cosmos, even intimate biographical facts about ourselves, such as where and when we were born or who are parents are—are facts that we can only know via our relations to other people. (Epistemologists call this “knowledge by testimony.”¹⁹)

But the flipside of this tremendous cognitive advantage is a distinctive kind of vulnerability. There’s an ineliminable element of trust built into our cognitive make-up from the start. We would *like* to think that our trust is rational, and in some way it is: we might not have a choice *whether* to trust, but we have some manner of discretion on the question of *whom* to trust. But on the other hand, by the time the critical capacities required to make these kinds of choices are properly developed, we are already so thoroughly socialized into a particular epistemic community that it will be virtually impossible for us to affect the kind of “critical distancing” that an ideal choice-situation requires (Begby, 2013; Cf. Nguyen, 2020).

This forces the recognition that while *in general*, our socio-epistemic embeddedness is a tremendous cognitive boon, some of us might nonetheless have the comparative misfortune of being born into communities where misinformation or prejudice predominates. For people in such situations, it will be exceedingly difficult—maybe even impossible—to find the epistemic resources to “move beyond”, by asking pointed, rationally motivated questions about why we all think like this, why certain assumptions are taken for granted, why the hypothesis-space takes this form rather than that. We can grant that there *is* a rational requirement to recognize one’s fallibility and to ask such questions. But at the same time, we recognize that critical reason is a limited resource and remains fundamentally constrained by the materials that it’s given to work with. Quite simply, we cannot but evaluate new evidence except in light of what we already presume to know. And what we must presume to know is deeply constrained by our socio-epistemic relations to others, whoever they may be, and the quality of the information that their testimony provides.²⁰

¹⁹ For a classic statement, see Hume, 1748. For an even more probing account, which might also be the first to tap into a realistic picture of human cognitive development, see Reid, 1764; 1785.

²⁰ On this see, Begby, 2021a.

One way to summarize the argument for non-ideal epistemology is this: we can certainly imagine an information-processing system with greater processing abilities than our own; but we cannot imagine a physically embodied real-time-operating information-processing system with *no* intrinsic performance limitations. Therefore, we will have to curtail our idealizations at some point: so why not just curtail them to fit a realistic picture of human cognitive capacities from the start? Similarly, we can certainly imagine an information society with a more expansive collective cognitive horizon, a flatter authority structure, and fewer incentives for using information/disinformation as a tool of deception or manipulation. But at the same time, we could hardly imagine an information-society with *no* limitations to its collective cognitive horizon and no institutionally sanctioned distinction between experts and non-experts, etc. Again, however much philosophy is naturally (and reasonably) prone toward idealization, we are going to have to recognize limitations somewhere. And these limitations will have to constrain any sensible, or if you will, pragmatist, account of epistemic rationality.

4. Whither pragmatism?

I'm not at all surprised to rediscover the influence of Bjørn's philosophy in many of these arguments. Nor am I surprised to recognize themes from thinkers like Rorty, and farther afield, Gadamer, both of which I also likely absorbed via Bjørn, either directly or indirectly.

Are these "pragmatist" themes? My default response to these kinds of questions has always been a shrug of the shoulders: does it matter? Pragmatism has a history of entangling itself in disputes about what "pragmatism" really commits one to. My general outlook has always been to think there are few things less "pragmatist" than engaging in such disputes.

But perhaps a more probing response is available. I think of epistemology as a normative discipline, centrally focused on human cognition. I *do* think (and I'm probably in a relative minority among contemporary epistemologists here) that the correct account of epistemic norms will require human cognitive agents to engage in trade-off reasoning determining what information, sought along what lines of inquiry, might be most relevant to them at any given point. The reason is simply that any pursuit of truth or knowledge in inquiry imposes opportunity-costs which draws epistemic resources away from the pursuit of truth or knowledge along other lines of inquiry. When it comes to evidence, non-ideal agents really

can't "have it all," and so, we need to be choosy. This choice-situation is always going to involve an element of risk. But clearly there is better or worse advice one can offer to someone in this situation. My advice is to choose to conduct such inquiries as might maximize the expected epistemic value of the information we would stand to gain. Differently put, we should seek the most consequential (or "relevant") information we could reasonably hope to have.²¹

I don't think this should be particularly controversial. But one consequence is immediately notable: candidate epistemic norms which counsel us simply to gather the most amount of knowledge or the most amount of truths that we can are simply barking up the wrong tree. Perhaps it's true in some abstract sense that our epistemic endeavors should be guided by "truth" or "knowledge": but we don't satisfy that norm by simply adding truths to our knowledge-base one by one, where we can find them. We have finite minds and cannot possess all the truths. Some truths are comparatively epistemically inconsequential. So we need to be resource-conscious and selective about which truths we pursue. I think this is the correct view to take whether we are considering scientific inquiry, policy-making, or mundane everyday epistemology at the individual level.

Is this a form of pragmatism? We might, perhaps, think of it as a kind of "ground-level pragmatism," one that directly seeks to influence how we conduct our inquiries. But as such, it seems rather different from the more explicitly "meta-philosophical pragmatism" that was advocated by Rorty and taken over by Bjørn. Meta-philosophical pragmatism is a position on philosophy. What I have called ground-level pragmatism is a position on inquiry more generally.

But of course, this neat distinction crumbles as soon as we recognize that philosophy is, in an important and recognizable sense, just another form of inquiry. As such, I would not be averse to thinking about meta-philosophical pragmatism simply as ground-level pragmatism applied to philosophy, understood as a form of inquiry. In fact, that would be a highly pleasing result.

But even so, it's not entirely clear to me how this all relates to the distinctively "American pragmatism" that so enamored Rorty. Instead, I find it just as natural here to seek inspiration in a range of philosophers whose contributions pre-date American pragmatism by a significant margin. Now, to be clear, I think Peirce, Dewey, and James are *also* inspirations

²¹ Begby, 2023b.

here. But if we think of them as bringing about an entirely novel chapter in the history of philosophy, we will be led down the path of asking what their distinctive methodological and doctrinal contributions were. And so, we will almost inevitably come to read them against the specific historical context they operated in—particularly, perhaps, the transition from late Hegelian idealism to the early stages of positivist empiricism and beyond. In that context, it seems reasonable to ask, for instance, what is “the pragmatist theory of truth?” And before we know it, we will be ensnared in distinctive, but to my mind, very “un-pragmatic” doctrinal questions that essentially reveal pragmatism as just another “philosophical theory.” By contrast, if we think about the contributions of American pragmatism in a slightly larger historical context, as just one chapter in a longer trajectory of a distinctive kind of broadly inquisitive, cautiously optimistic, science-inspired philosophical humanism, then things will look quite different. Accordingly, I’d rather like to think of pragmatism as an attitude that rebels against the increasing insularity of professionalized, institutionalized philosophy, and which maintains that philosophers should feel free—indeed encouraged—to have their inquiries range as widely as possible, never being overly concerned about whether either the product or the process is particularly “philosophical.”

At this point in history, we have a thoroughly institutionalized conception of philosophy. Philosophy is done by professionally trained philosophers, working in specialized Departments of Philosophy. We routinely distinguish between Philosophy and other “humanistic disciplines,” such as History, Literature, or Gender Studies. We further distinguish these humanistic disciplines from “science.” We sometimes wonder whether the “social sciences,” so-called, really are sciences or only humanistic disciplines trying to pass themselves off as sciences, thereby hoping to draw some of the prestige that this designation holds in contemporary academia and public life.

If this is what Philosophy *is*, then it stands to reason that philosophers should engage in “meta-philosophy”: they should aspire to draw a boundary around the “proper domain of philosophy” and be prepared to have something to say about the unique methodologies that we deploy to investigate questions falling within that domain. Part of what makes this reasonable is of course administrative: if Philosophy is to control its own budget, request new faculty-lines, make a distinctive contribution to the undergraduate curriculum, and so on, then it’s hardly too much to ask

that philosophers have something concrete to say about what makes “philosophy” different from any other discipline.

These developments are perfectly understandable. And conceivably, it’s been a good thing for philosophy at large. There are probably more “philosophers” now than there has ever been, publishing more “philosophy” in accredited “philosophy journals” than ever before. On the other hand, a large part of our professional pride no doubt comes from our long tradition—unlike these new upstart disciplines, philosophy can trace its roots back to the dawn of written culture. But at the same time, we tend to forget that the great philosophers of the past were scarcely “philosophers” in the current sense of the word. Nor did their contributions to philosophy in any clear sense conform to a particular methodology or constrain itself to particular subject matters or to considering only a certain kind of evidence. Obviously, this is not to say that they were necessarily something else instead—say, “scientists.” Instead, I take it that they were distinctively engaged in a kind of opportunistic, problem-oriented form of inquiry that would cast its net as widely as practically feasible, so as to work on any possibly relevant source materials and reason along any pathway that might generate a persuasive train of thought toward some conclusion. In brief: their type of inquiry pre-date what we now think about as the distinction between science and philosophy, and, I would like to think, was better off for it.

On this sort of thinking, calling someone like Aristotle, or Descartes, or Leibniz a philosopher *but also* a scientist seems like a recipe for misunderstanding their work. Their contributions to what we now recognize as philosophy and what we now recognize as science were arguably much more integrated than that: these were not insights emerging from a compartmentalized mind, now engaged in this form of inquiry, now in another. Instead, their best work embodies, in an important sense, an integrated philosophico-scientific outlook on the world. Closer to our time, none of Voltaire, Diderot, Rousseau, Hume, Wolff, Smith, or Reid (to mention but a few) were professional philosophers: some of them were employed in the University sector, but none of them in a Department of Philosophy. Moreover, while they were also making important contributions to philosophy, they were concurrently engaged in important reflection on other issues, be it medicine, politics, economics, history, or theology. I don’t think anyone should read their “philosophical” works today and imagine that these are the products of a mind particularly pre-occupied with questions like, “is this philosophy?” or “am I doing it right?”

I'm taken with the story that the first "great philosopher" in our canonical timeline who actually held an appointment as "Professor of Philosophy" would have been Kant. So it's perhaps no surprise that this appointment broadly coincides with Kant's sudden obsession with legislating the proper domain of philosophy and with outlining a distinctive methodology fit to answer those questions, namely the *kritische Methode* now familiar from the *Critique of Pure Reason*. While few of us would look to Kant's methodology as guiding their philosophical inquiries today, many of us nonetheless maintain an implicit conception of philosophy as a distinctive domain of inquiry, guided by a distinctive set of methodologies. (Though of course, when time comes around for us to fill in the dreaded "methodology" section of a grant proposal, we often discover that we have very little to say about what this methodology is.) Nevertheless, I think it's this self-conception which also leads us to think that, as professional philosophers, we must "stay in our lanes," which prompts us to police others' work by asking, "so, how is this philosophy?" etc.

Moreover, I believe this is the conception of philosophy that leads to the kind of insularity that "pragmatists" like Rorty and Bjørn were concerned with. They were correctly pointing out that as philosophers become increasingly pre-occupied with their status as "philosophers," the more insular our "philosophical problems" become, and the more these problems will tend toward intractability: they will be intractable, not—as we might like to think—because they are orders of magnitude more complex than the questions that occupy researchers in otherwise comparable disciplines. Rather, they are intractable simply because they have lost their grounding: as a result of drawing a bright circle around "the proper domain of philosophy" we have severed the connections that gave urgency and direction to our inquiries in the first place. We may have saved our budgets and our tenure-lines, but we are at risk of losing our way as philosophers. We will certainly have lost our life-giving connection to the tradition of reflection that gave our discipline its name.

What I make of pragmatism today, and which I like to think I've absorbed from Bjørn—by reading, by discussion, by example—comes to this: philosophy is at its best when it is not particularly preoccupied with its own status, but rather taps into the longer arc of its proud and unique tradition so as to engage open-mindedly and opportunistically with any form of inquiry that comes within its purview or that might conceivably generate a question of philosophical interest. And judging by the span of that tradition, this is a very wide range of questions indeed. If philoso-

phers have anything distinctive to contribute in modern academia, it's not by way of the unique methodology we deploy or the doctrinal insights we have painstakingly chiselled out over the years. To the contrary, it's due to our consistent refusal to be fettered by any restrictive methodology and our continued insistence that all questions are potentially philosophical questions.

References

- Andersen, H.K., & Mitchell, S.D. (Eds.) (2023). *The Pragmatist challenge: Pragmatist metaphysics for philosophy of science*. Oxford University Press.
- Begby, E. (2013). The epistemology of prejudice. *Thought* 2(2), 90–99.
- Begby, E. (2014). Hermeneutics and Pragmatism. In J. Malpas & H.H. Gander (Eds.), *The Routledge Companion to Hermeneutics*. New York: Routledge.
- Begby, E. (2021a). Evidential preemption. *Philosophy and Phenomenological Research* 102(3), 515–530
- Begby, E. (2021b). *Prejudice: A study in non-ideal epistemology*. Oxford University Press.
- Begby, E. (2023a). Theory of Mind deficits in language delayed deaf children? A reconsideration of evidence from false belief tasks. *Philosophy and the Mind Sciences* 4, 1–26.
- Begby, E. (2023b). Opportunity costs and resource allocation problems: Epistemology for finite minds. *Canadian Journal of Philosophy* 53(5), 400–417.
- Begby, E. (2024a). From belief polarization to echo chambers: A rationalizing account. *Episteme* 21(2), 519–539.
- Begby, E. (2024b). *Precis of Prejudice: A study in non-ideal epistemology*. *Inquiry* 67(8), 2548–2570.
- Begby, E. (Forthcoming). Signal-to-noise: Toward an alternative account of the Fake News phenomenon. In M. Popa-Wyatt (Ed.), *Mis/Disinformation*. Cambridge University Press.
- Begby, E. (In progress). *Belief Formation in Antagonistic Information Environments*. Book manuscript.
- Davidson, D. (1970). Mental events. Reprinted in *Essays on actions & events*. Oxford: Clarendon Press, 1980.
- Davidson, D. (1991). Three varieties of knowledge. In A. Phillips Griffiths (Ed.), *A. J. Ayer: Memorial Essays*. New York: Cambridge University Press.
- Davidson, D. (1999). Reply to Bjørn Ramberg. In L.E. Hahn (Ed.), *The Philosophy of Donald Davidson*. The Library of Living Philosophers, Open Court Publishing: La Salle, IL.

- Dennett, D.C. (1991). *Consciousness explained*. Boston, MA: Little, Brown & Co.
- Gadamer, H.-G. (1960). *Truth and method*. Translated by J. Weinsheimer & D.G. Marshall, London: Continuum Publishing, 1989.
- Hume, D. (1748). *An enquiry concerning human understanding*. T.L. Beauchamp (Ed.). Oxford University Press, 1999.
- Kant, I. (1781/1787). *Kritik der reinen Vernunft*. Translated by N.K. Smith. London: MacMillan, 1958.
- Mills, C. (2005). Ideal theory as ideology. *Hypatia* 20(3), 165–184.
- Nguyen, C. T. (2020). Echo chambers and epistemic bubbles. *Episteme* 17(2), 141–161.
- Ramberg, B.T. (1988). Charity and ideology: The field linguist as social critic. *Dialogue* 27(4), 637–651.
- Ramberg, B.T. (1989). *Donald Davidson's philosophy of language: An introduction*. New York: Blackwell.
- Ramberg, B.T. (1997). The source of the subjective. In L.E. Hahn (Ed.), *The Philosophy of Hans-Georg Gadamer*. The Library of Living Philosophers, Open Court Publishing Company, Chicago and La Salle, IL.
- Ramberg, B.T. (1999). The significance of charity. In L.E. Hahn (Ed.), *The Philosophy of Donald Davidson*. The Library of Living Philosophers, Open Court Publishing: La Salle, IL.
- Ramberg, B.T. (2000). Post-ontological philosophy of mind: Rorty versus Davidson. In R. B. Brandom (Ed.), *Rorty and his critics*. Malden, MA: Blackwell.
- Ramberg, B.T. (2001). What Davidson said to the skeptic or: Anti-representationalism, triangulation and the naturalization of the subjective. In P. Kotatko, P. Pagin, & G. Segal (Eds.), *Interpreting Davidson*. Stanford, CA: CSLI Publications.
- Ramberg, B.T. (2003). Illuminating language: interpretation and understanding in Gadamer and Davidson. In G.G. Prado (Ed.), *A house divided: comparing analytic and continental philosophy*. Amherst, NY: Humanity Books.
- Ramberg, B.T. (2004). Naturalizing idealizations: pragmatism and the interpretivist strategy. *Contemporary Pragmatism* 1(2), 1–66. <http://www.brill.com/>
- Ramberg, B.T. (2005). Bør dyr ha sjel? *Arr—Idéhistorisk tidsskrift* 4, 47–55.
- Reid, T. (1764). *An inquiry into the human mind on the principles of common sense*. In W. Hamilton (Ed.), *The works of Thomas Reid*, Vol. 1. Edinburgh: MacLachlan and Stewart, 1872.
- Reid, T. (1785). *Essays on the intellectual powers of man*. In W. Hamilton (Ed.), *The works of Thomas Reid*, Vol. 1. Edinburgh: MacLachlan and Stewart, 1872.
- Rorty, R. (1965). Mind–body identity, privacy, and categories. *Review of Metaphysics* 19(1), 24–54.
- Rorty, R. (1970). In defense of eliminative materialism. *Review of Metaphysics* 24(1), 112–21.

Rorty, R. (1979). *Philosophy and the mirror of nature*. Princeton, NJ: Princeton University Press.

Rorty, R. (2000). Response to Ramberg. In R. B. Brandom (Ed.), *Rorty and his critics*. Malden, MA: Blackwell.

Williamson, T. (2000). *Knowledge and its limits*. Oxford University Press.