

KNOWLEDGE BY ACQUAINTANCE & IMPARTIAL VIRTUE

Emad H. Atiq

Abstract

Russell (1911/12) argued that perceptual experience grounds a species of non-propositional knowledge, “knowledge by acquaintance,” and in recent years, this account of knowledge has been gaining traction. I defend on its basis a connection between moral and epistemic failure. I argue, first, that insufficient concern for the suffering of others can be explained in terms of an agent’s lack of acquaintance knowledge of another’s suffering, and second, that empathy improves our epistemic situation. Empathic distress approximates acquaintance with another’s suffering, and empathic agents who are motivated to help rather than disengage exhibit an important epistemic virtue: a variety of intellectual courage. A key upshot is that an independently motivated account of the structure and significance of perceptual experience is shown to provide theoretical scaffolding for understanding a famously elusive idea in ethics—namely, that the failure to help others stems from a kind of ignorance of their situation.

Keywords: Acquaintance, knowledge, moral motivation, moral objectivity, suffering, intellectual courage, interoception, the self.

Introduction

A familiar problem in ethics begins with the depressing observation that there are agents, both actual and possible, who cannot seem to find motivation to care about the suffering of others. Such agents often appear rational and knowledgeable. And yet it is hard to shake the conviction that in addition to being atypical or amoral, they are likely missing something or else mired in some kind of error or confusion. The conviction that inadequate concern for others should be explainable in terms of distortions in an agent's perspective may not be universally shared or impervious to doubt, but it finds expression in a broad range of philosophical and literary works. The challenge is to clarify and, if possible, vindicate it, as part of a broader effort to make vivid the appeal of the moral point of view.¹ Street (2016: p165) puts the challenge thusly:

The Holy Grail in secular metaethics would be a vindication of this inchoate sense—a philosophical account that clearly explained, with no mysterious metaphysical or epistemological remainder, and in a way that rang deeply true with our pre-philosophical sense of things, exactly what... an agent's mistake would consist in when he or she contemplated the suffering of others and saw it as counting for nothing.

Traditional responses to this challenge share a common feature in that they've tended to rely on assumptions that are internal to ethics. The Platonist appeals to *sui generis* moral facts about

¹ Many morally conscientious agents experience their concern for others as an expression of their objectivity. Indeed, being objective seems to be part of the appeal of the moral point of view. For example, Boyd (1988, p341) describes “the very strong intuition which many philosophers share that the person for whom moral judgments are motivationally indifferent would not only be psychologically atypical but would have some sort of cognitive deficit.” The challenge is to clarify and vindicate this intuition. Blackburn (1984, p222) calls it moral philosophy's “holy grail... the knock-down argument that people who are nasty and unpleasant and motivated by the wrong things... are reasoning badly, or out of touch with the facts.”

the Good that elude the callous and indifferent, and that, if known, would by brute metaphysical necessity produce motivation towards other-regarding action.² The Kantian argues based on assumptions about the nature of agency that, as far as I can tell, do no theoretical work outside of ethics, that immoral action necessarily mires an agent in a kind of practical contradiction, since acting against the requirements of morality turns out to be inconsistent with what one allegedly must believe about one action's to be acting at all.³ The Hobbesian appeals to a surprising alignment between moral demands and an agent's self-interest, so that agents indifferent to the needs of others end up frustrating their own desire to live a happy and fulfilling life.⁴

These traditional approaches may turn out to be defensible. But to those most in need of convincing, the grounds for convicting the callous and indifferent of error will likely seem invented just for that purpose. For the Platonist, Kantian, and Hobbesian don't just rely on controversial claims concerning inherently prescriptive moral facts, or essentially moralistic commitments implicit in action, or thoroughgoing moral integrity as the key to happiness. Being independently persuaded of morality's importance seems to be a prerequisite for finding such claims compelling.⁵ That wouldn't be a problem if our sole concern was to bolster the moral resolve of those who aren't

² Both Glaucon and Thrasymachus raise the question of why one ought to behave justly when it isn't in one's self-interest to do so. The *Republic* offers several answers and at least one casts the failure to act justly as an epistemic failing: those who've seen the "form of the Good" will see the point in just action (517b-c).

³ See discussion in Korsgaard (2009, p90-91) & Singh (2022, p159-160).

⁴ See discussion in Rutherford (2003, 370-31).

⁵ It might be going too far to deem the entire theoretical background of the traditional views "a just so story, a myth, that many moral philosophers tell themselves" (Nielsen 1984, p.90), but it is hard not to sympathize with the doubts of the unpersuaded. See e.g., Blackburn (1984, p222). To be clear, Nielsen and Blackburn aren't moral skeptics. They are merely skeptical of arguments in moral philosophy that aim at "holy grails."

generally skeptical of ethical or metaethical considerations, but there are real skeptics among us who think that morality is a fraud on the intelligent. An argument that aspires to convince not the choir but the unconverted that indifference towards others stems from a failure to see things clearly is more likely to succeed if its premises, in virtue of doing theoretical work outside of ethics, can be motivated on wholly non-moral grounds.⁶

This paper explores the prospects of explaining the indifferent agent's failure based on a Russellian idea that is enjoying a revival in contemporary epistemology—namely, that perception, including seeing, feeling, hearing, and tasting, grounds a species of non-propositional knowledge: “knowledge by acquaintance.”⁷ I argue that *if* this Russellian claim can be successfully defended and we adopt the general theoretical framework in which it is embedded, then we can explain an agent's indifference to the suffering of others in terms of a lack of knowledge. Indifference to suffering, on this view, turns out to be a distinctly epistemic defect, a form of ignorance that we can mitigate by trying to empathize with others. While some recent work has begun exploring the connection between acquaintance and moral virtue (Marshall 2019, Atiq 2021, Atiq & Duncan 2024), I offer a novel characterization and defense of this connection in terms, principally, of the

⁶ Convincing the moral skeptic on grounds they can accept seems to me an important desideratum unique to the moral context, and not a general constraint on solutions to skeptical challenges. The skeptic about the external world or other minds may be a philosophical invention, but there are real skeptics among us who think morality is a con and who take false epistemic pride in their moral indifference. If we can ground moral demands in epistemic virtue, we will have discovered a means of reaching such skeptics at least insofar as they care about the epistemic good. Another way of understanding the intellectual and practical stakes is in terms of the importance of vindicating the self-conception, and epistemic pride, of ordinary moral agents.

⁷ See e.g., Conee (1994), Tye (2009), Stump (2010), Johnston (2011), Gertler (2012), Coleman (2019), Duncan (2020), Atiq (2021), Giustina (2022), and Kriegel (2024).

idea that certain forms of knowledge and epistemic virtue turn out to be *essentially* motivating and the relevant motives turn out to be other-directed.

The discussion is structured as follows. §1 clarifies the key commitments of an acquaintance-based epistemology. The goal, here, is not to establish that the background epistemology is true. Rather, it is to clarify the theoretical context for evaluating the arguments of subsequent sections. §2 demonstrates that the epistemology can support a robust explanation of why indifference to the suffering of others amounts to an epistemic defect. I argue based principally on the phenomenology of suffering and of experience more generally that, necessarily, if a subject had acquaintance knowledge of the sensible content of another's suffering, then that subject would be motivated to alleviate the other's suffering. Additionally, given plausible assumptions about ordinary agents, the only psychologically feasible choice compatible with preserving one's acquaintance knowledge of another's experience involves ameliorating the other's suffering. §3 argues that although no subject can be directly acquainted with another's suffering, we can improve our epistemic situation by trying to empathize. Empathic distress, by virtue of being an unpleasant experience, simulates acquaintance with another's suffering, and empathic agents who are motivated to help others exhibit an important epistemic virtue in relation to the knowledge they lack—a variety of intellectual courage. Being willing to ameliorate the suffering of others rather than one's own empathic distress is analogous, I argue, to being immune to wishful thinking and related epistemic vices. §4 takes stock, considers several outstanding objections, and further clarifies the proposal.

§1 Knowledge by acquaintance of suffering

In recent years, there has been a resurgence of interest in the view that in perceptual experience subjects bear an epistemically significant cognitive relation directly to the objects of perceptual awareness that is different from thinking truths about such objects. Some call this relation “acquaintance” and take acquaintance to ground a species of non-propositional knowledge. Although Russell (1911) is often credited with introducing the concept of knowledge by acquaintance to analytic philosophy, the general idea—that there is a way of knowing things that does not necessarily involve being in a truth-evaluable mental state—has recurred throughout the history of philosophy.⁸ This way of knowing involves directly perceiving or sensing an object. For example, when I see a red apple or experience a dull headache, I gain acquaintance knowledge of the immediate objects of my sensory experience, and the knowledge I gain is not reducible to believing truths about the objects. Neither is it reducible to practical know-how or any dispositions to act, infer, or have other thoughts. Rather, the mental state’s status as a form of knowledge is grounded in the relation of direct awareness that it essentially involves.

At its core, the view consists of two related theses: a metaphysical thesis that perception essentially involves a form of awareness of objects that is especially direct and not mediated by awareness of facts or anything else; and an epistemic thesis that this direct form of awareness grounds a species of knowledge of objects that is non-propositional, so that the only way to know an object in the relevant sense is to perceive it directly.⁹ The resulting account contrasts sharply

⁸ See discussion in Raleigh (2019) and Duncan (2021).

⁹ Coleman (2019, p59) writes:

acquaintance is a cognitive achievement: it involves making substantive *contact* with something...

In acquaintance, awareness achieves metaphysical and epistemic connection with another

with what is perhaps the more conventional view that the world can only be grasped through representations, either linguistic (sentences) or non-linguistic (propositions and propositional attitudes), which mediate between perceivers and reality.¹⁰ Hence, Gertler (2012, p101) writes “[p]erhaps the most controversial commitment of the acquaintance approach is the claim that reality can intersect with the epistemic: that is, that one’s epistemic grasp of a bit of reality—a fact, event, or property—can be partly constituted by that reality itself.”

A serious defense of the epistemology lies well beyond the scope of this paper. However, it might prove helpful to briefly review the kinds of arguments on offer, primarily to show that the motivating considerations bear no obvious connection to ethics. Some theorists argue from linguistic evidence (Benton 2017): the verb “to know,” in addition to taking “that” clauses (“I know that the sun is shining”), also takes noun-phrases (“I know the pain of losing someone”), which suggests that we ordinarily conceive of agents as being able to stand in the knowing relation directly and non-propositionally to things. Others have argued that perceptual experience exhibits

existent... Had it not enlarged the subject’s world in this way, acquaintance’s claim to be a kind of knowledge would be jeopardized.

In a similar vein, Gertler (2012, p95) observes:

[w]hen I am acquainted with an object, my awareness of that object is epistemically direct: it is noninferential and does not epistemically depend on an awareness of anything else. My awareness is also metaphysically direct: there is no object, fact, event, or process that mediates my access to the object.

¹⁰ See, e.g., Williamson (2000). Duncan (2020, p3560) observes:

Whether or not most contemporary philosophers explicitly believe that all knowledge (save know-how) is propositional, this doctrine is consistently reflected in philosophical practice—propositional knowledge is all that’s talked about, for one thing—and it is presupposed in many debates—including debates that other kinds of knowledge would be highly relevant to such as debates about the structure of knowledge, the nature of justification, the epistemic significance of experience, certainty, intuition, and self-knowledge.

intrinsic or final epistemic value, and is not just instrumentally valuable as a means of knowing truths, and that the best account of perception's final epistemic value appeals to knowledge by acquaintance (Johnston 2011, Atiq 2021).¹¹ Russell's (1911, p108) own argument appeals to intuitive considerations. Knowledge is associated with a capacity for being in touch with the world, and perceptual experience puts us in touch with worldly phenomena: colors, smells, sounds, and sensations are made available for observation and reflection.

Several of the available arguments have the structure of an inference to the best explanation. Acquaintance knowledge best explains what Mary in Jackson's (1982) famous thought experiment learns when she sees red for the first time (Connee 1994, Tye 2009: 131-33), and why the knowledge she gains is elusive.¹² The acquaintance knowledge hypothesis has also been invoked to explain why perception seems to satisfy, independently of propositional content, knowledge's unique functional role. In general, what is known serves as evidence and inferential

¹¹ Johnston (2011) argues that something worth epistemically wanting seems lost when we imagine ourselves in the circumstances of a person who develops perfectly reliable blind-sight. The person reliably forms true beliefs about her immediate surroundings, the physical layout of a room and the location of objects within it, without visually experiencing her surroundings. We can stipulate that the agent has propositional knowledge of the room's layout perhaps because she is aware of the reliability of her blind-sight. Yet, intuitively, the agent lacks an epistemic good, one constitutively tied to being able to visually perceive objects in one's environment. An explanation of our intuitions about the case is that the agent lacks the good of knowledge_A. Elsewhere, Johnston (2001, p206) writes that "[s]ensing is a form of knowing, even if it is not yet knowing that." For an argument that we must invoke knowledge by acquaintance to explain the unique epistemic appeal of "Edenic" counterfactuals involving unmediated and comprehensive perceptual awareness of all worldly objects and their qualities, see Atiq (2021, pp14040-42). On "Edenic" counterfactuals generally, see Chalmers (2006, p66-7).

¹² As Tye (2009: 98) puts it "Mary may not know that shade of red a few moments later after turning away." For Tye, knowing_A red is "restricted to the moment in which one is conscious of it" (Crane 2012). On the "elusiveness" of knowledge by acquaintance, see Atiq (2021: 14043-47).

basis for belief. And the contents of perception (the non-propositional objects of acquaintance) serve as evidence and inferential basis for our beliefs about the external world (Duncan 2020, 2021). The knowledge-like functional role of perception is especially evident in the case of evaluative reasoning, where the objects of perceptual acquaintance (for example, a sensation of pain) serve as non-propositional evidence for evaluative belief (that pain is bad) (Atiq 2021).

We can afford to be brief in summarizing this literature since the present goal is not to defend the acquaintance approach. In fact, the above summary is far from exhaustive.¹³ The point in highlighting some of the available arguments is that the considerations theorists have offered in defense of the acquaintance knowledge hypothesis are ethically neutral. There is little reason to worry that the key concepts of acquaintance and acquaintance knowledge have been invented to solve a distinctly moral problem. Moreover, an argument in ethics that relies on the epistemology cannot be accused of begging the question against any figure in the moral dialectic, including the skeptic who denies that caring about others has anything to do with objectivity. For the existence of acquaintance knowledge is ethically and metaethically uncontested. This is worth flagging given its importance to my overall argument. I'll argue later that the fact that my starting points are independently defensible is a significant reason for preferring the account I'll offer of the amoral agent's error over competing accounts in the literature. For now, the point is just that there is enough ethically neutral support for an acquaintance-based epistemology to justify provisionally accepting it for the limited purpose of exploring its ethical implications.

¹³ Stump (2010, ch1) defends knowledge by acquaintance on several grounds, including based on its theoretical usefulness in the philosophy of religion. On acquaintance knowledge being a necessary basis for propositional knowledge, see Kriegel (2024). A comprehensive review of the burgeoning literature on acquaintance would take us well beyond the scope of this paper.

The account requires a little more set up, however, since we haven't yet clarified what it is that we are acquainted with in visual, auditory, or somatosensory perception. It turns out that there is little consensus on this question among acquaintance theorists. According to "direct realists," in non-hallucinatory perception subjects are acquainted with mind-independent physical particulars qualified in various ways by sensible qualities.¹⁴ When Mary sees a red apple, she is aware of the facing front of a physical object instantiating qualities like redness. The red facing front of the apple is the directly known_A object of her visual perception immediately present in her experience, and not some mentalistic or representational item that mediates between the apple and her awareness of it (Johnston 2011). By contrast, the "sense datum theorist" maintains that perception acquaints us with non-physical mental particulars, volumes of qualitative stuff located in a subject's sensory field (Price 1950, Broad 1959, Robinson 1994). While I do not intend to settle such intramural debates, I will be making an important assumption about the objects of acquaintance involved in *suffering*. I assume that suffering, defined as the conscious experience of pain and felt aversions to pain, involves acquaintance with phenomenal particulars unique to each suffering subject: specifically, a token pain and a token aversion. Some clarifying remarks about this assumption are in order.

¹⁴ According to some direct realists, hallucination acquaints us with *sensory profiles*: "When the visual system misfires, as in hallucination, it presents uninstantiated complexes of sensible qualities and relations, at least complexes not instantiated there in the scene before the eyes" (Johnston 2011, p135). A key feature of this view is that the qualitative and relational complexes are not mental items but fall under their own objective kind (p146). Other direct realists have offered other accounts of hallucination, as involving, for example, acquaintance with sense data.

First, the definition of suffering as involving conscious pain and aversion is meant to be stipulative though not arbitrary.¹⁵ The target phenomenon is morally significant: agents ought to ameliorate conscious suffering of the relevant sort.

Second, the idea that token pains and token aversions are objects of acquaintance can be defended based on the background metaphysics—in particular, the assumption that states with phenomenal content are best analyzed in relational terms, as involving acquaintance with things. Pain exhibits such manifest phenomenal features as intensity, location, and duration. We needn't decide whether it is a quality of bodily states (tissue damage and the like) or some type of somatic/affective field; or whether an experience of pain involves representational content. The key claim is that there is a phenomenal particular or quality of which one is directly aware in pain perception. Likewise, felt aversions to pain have a phenomenology of their own, a recognizable sensory dimension. Aversions as well as attractions vary in intensity, for example. And although they may seem more ethereal than, say, the colors, “experiential” desires and aversions can be “phenomenologically salient,” presenting as objects of perceptual acquaintance.¹⁶

¹⁵ For an alternative view of suffering “without subjectivity,” see Carruthers (2004).

¹⁶ For further discussion on this point, see Ayded & Fulkerson (2019):

Clearly, some sort of a desire-like state is involved somewhere in the phenomenology of sensory affect. When we have a pleasant sensation such as the taste of the ripe sweet strawberry, there is some sort of a tug that we feel toward the taste, we feel some kind of pull that makes us ‘want to continue having the sensation’....If there is any sense in which there is a kind of desire involved in affective experiences, it is a phenomenologically salient experiential desire.

See also Johnston (2001, p188):

there is a perfectly good non-philosophical sense of ‘desire’ in which desire is not only one of the springs of action, but a state which makes certain kinds of actions readily intelligible. In this sense of ‘desire’, which we might distinguish by the somewhat pleonastic name ‘affective desire’, we desire other things and other people, we are struck by their appeal, we are taken with them. This is

Third, the idea that two suffering subjects are directly aware of numerically distinct pains and aversions is similarly defensible. The direct realist takes the object of pain perception to be a bodily disturbance qualified by the sensible qualities associated with pain. The painfully qualified bodily disturbance is directly present in conscious awareness. And since distinct subjects in pain are aware of distinct bodily disturbances, the parts of reality of which they are directly aware, and so know_A, count as distinct. Distinctness is, likewise, implied by sense-datum theory. The sense datum theorist takes the direct object of pain perception to be a non-physical mental particular. On standard versions of sense datum theory, the non-physical qualia are essentially located in each subject's unique phenomenal space or somatic field so that two subjects experiencing pain perceive distinct qualia.

My assumptions about pain and suffering, while controversial, complement the epistemology. I do not consider, here, views that I find harder to square with the acquaintance approach. For example, some theorists contend that the mental state of being in pain is exhausted by its representational or propositional content: there is nothing more to being in pain than being in a state that represents tissue damage or some other kind of disturbance in one's body (Tye 2006). Representationalists may be able to accommodate knowledge by acquaintance, but I doubt their view adequately explain the phenomenology of experience. In particular, 'pure' representationalism about pain seems to me to mischaracterize the most salient aspect of a painful experience—the affective phenomenology—which is best understood, I think, in relational and objectual terms. I am also not sure whether representationalists can adequately explain what *grounds* the unique epistemic value of acquaintance consistent with all the cases in which that

part of how things are manifest to us.... On the face of it, appeal is as much a manifest quality as shape, size, color and motion.

value is intuitively realized (including, e.g., in hallucinations). On the relational approach, the value is grounded in a metaphysically real relation of “substantive [cognitive] contact with something” (Coleman 2019, p59), a relation that secures an “epistemic grasp of a bit of reality... partly constituted by that reality itself” (Gertler 2012, p101). It’s unclear whether the representationalist can offer a comparable explanation. But I won’t press the points here. I take the above claims about knowledge by acquaintance and suffering largely for granted, without claiming to have systematically ruled out alternative views, including views that may be in line with the paper’s overall argument.

§2 Knowledge_A that’s Essentially Motivating

The discussion so far has been limited to stage-setting. With relevant aspects of the background epistemology and metaphysics clarified, we can turn to the argument for an ethical upshot. Consider the basic moral duty to mitigate others’ suffering. Morality may demand more than just the provision of aid to those in desperate need. For example, it is not clear that the obligation not to lie is best explained in terms of the importance of minimizing suffering. Moreover, not all ways of being motivated to help others may be equally virtuous from the moral point of view. Nevertheless, indifference to the suffering of others represents a paradigmatic kind of moral failure, and the main claim I wish to defend in this section is that this familiar species of moral failure can be explained in terms of an epistemic defect in the indifferent agent: a form of ignorance associated with not knowing_A the sensible content of other people’s suffering. Throughout, I use the phrase “sensible content of another’s suffering” (and sometimes, for convenience, “another’s suffering”) to refer to the token pain and token aversion that the suffering agent is, *ex hypothesi*, acquainted with.

As an initial observation, it seems unlikely that any subject has ever had first-hand acquaintance knowledge_A of the sensible content of another's suffering. We may have sensed experientially identical pains, and having known_A our own pains, we might be propositionally knowledgeable about other people's pain, knowing truths about what pain feels like, the behavior it tends to produce, and perhaps even evaluative truths about pain. But none of this knowledge would be equivalent to or a substitute for knowing_A another's *token* pain by feeling it.¹⁷ Recall that the epistemic value of acquaintance is supposed to be grounded in the direct cognitive contact that is achieved with some aspect of reality. And we lack such contact with the token pains experienced by others, given our assumptions about the nature of pain perception. A surgeon operating on a patient might visually sense the bodily disturbance responsible for the patient's pain. But visual experience cannot deliver acquaintance knowledge of that aspect of the disturbance that is painfully qualified, the aspect that is immediately present in the patient's own awareness. Knowing_A another's pain in the relevant sense involves directly experiencing its painful character.

Set aside for the time being concerns about why anyone would *care* to have knowledge of the relevant sort, though of course the question needs to be addressed eventually. One might reasonably wonder why it would be valuable for A to know_A B's pain having known many experientially indistinguishable pains, or for Mary to know_A many more instances of redness after having known_A a single instance. How valuable could the additional knowledge_A be? I'll address this concern more fully later, but to foreshadow the rest of the argument: all knowledge *qua* knowledge exhibits final epistemic value, but this does not mean its overall value has to be

¹⁷ See discussion in Atiq & Duncan (2024).

particularly high in individual cases.¹⁸ What does critical work in our discussion is the fact that the knowledge_A in question would influence our choices in ways that are experientially intelligible, as I'll argue shortly. Its practical import helps dispel the indifferent agent's misconception that their choices are informed by a fully objective take on the world. Second, the epistemic value of acquaintance knowledge resides in an agent's perceptual grasp of aspects of reality. Two introspectively indistinguishable acquaintances can put one in contact with distinct parts of reality and, as a result, represent distinct and non-intersubstitutable cognitive achievements. Third, the knowledge_A at issue will turn out to be instrumentally significant for the way it tests an agent's intellectual courage and epistemic well-functioning. In short, by the end of the discussion it should become apparent why knowing_A another's pain is important, even if it might seem impossible, uninteresting, or trivial at first.

Setting aside concerns about the value of knowing_A, what would it be like for A to share, in the direct acquaintance sense, the phenomenal particulars involved in B's suffering (that is, the token pain and token aversion to the felt pain)? Although we cannot answer this question based on having had the experience ourselves, we can make reasonable inferences based on the phenomenology of suffering and experience more generally. There is an argument to be made that if A were to directly experience B's pain and B's aversion, then A would, necessarily, be averse to B's pain (note that all instances of "B's pain" and "B's aversion" in what follows should be understood *de re*, as referring to the token pain-aversion originally felt by B, unless I indicate otherwise). The argument involves motivating the claim that A's experience of being directly aware of B's aversion towards a token pain would *constitutively* entail A's being averse to the

¹⁸ Like Kvanvig (2003, p. 232), I take the "value assumption" that knowledge is worth valuing for its own sake to be an essential feature of our concept of knowledge.

token pain. A's awareness would be sufficient for being averse because A should be able to experience the aversion *as their own* in virtue of the awareness, and not as, say, *essentially B's aversion* as if B's aversion exhibits some unique phenomenological flavor or 'tag' that prevents it from experienced by anyone else *as their own*.¹⁹ Why think this? Because the self does not impinge on phenomenal content in such a way as to make our pains and aversions essentially self-involving—a claim we can defend in several different ways.

One way appeals to a famous and somewhat radical Humean observation that the self makes *no* discernable contribution to the phenomenology of experience.²⁰ When I experience an aversion to a dull headache, it is the felt force of the aversion I perceive, not *myself feeling averse*. Neither is there any phenomenal feel that reveals the felt aversion to be *essentially mine*, just as there is no analogous feel in the case of seeing a red apple. To be sure, we self-ascribe pains and aversions in ways that we do not colored objects, but that is compatible with the observation. An experience of pain may be accompanied by the *belief* that the pain is one's own, or that no one else is currently experiencing the pain, or that the pain is located in one's body. However, such beliefs needn't be based on the phenomenology of pain being "self-involving." On this point, one might be tempted to follow Hume:

¹⁹ The condition of experiencing an aversion *as one's own* shouldn't be misconstrued as involving anything intellectual or cognitive. We can be alienated, judgmentally, from our desires and aversions, which doesn't prevent them from being *ours* in the thin sense at issue. The relevant sense of "ownership"—of being the subject of the aversion—is phenomenological, peripheral (in that it's not usually attended to), and secured by awareness alone (as opposed to, say, sensing oneself *in* the object of awareness). I expand on these points below. I'll argue that if A is directly aware of B's aversion, then A experiences the aversion as their own in a way that's sufficient for A to be averse.

²⁰ *Treatise of Human Nature*, Book I, Part IV, Section VI.

But self or person is not any one impression, but that to which our several impressions and ideas are suppos'd to have a reference... For my part, when I enter most intimately into what I call myself, I always stumble on some particular perception or other, of heat or cold, light or shade, *love or hatred, pain or pleasure*. I never can catch myself at any time without a perception, and never can observe any thing but the perception (*Treatise*, Book I, Part IV, Section VI).

Hume's observations find some confirmation in my own phenomenology, and I suspect others' as well (see e.g., Shoemaker 1994).

But perhaps the Humean position is *too* radical.²¹ Plausibly, we experience aches, pains, and other sensations in a way that's not completely neutral with respect to the fact that such experiences are had by *a* subject. There is a less controversial way of motivating our main point about the self and phenomenal content. We can grant that subjects might be aware of themselves as subjects of experience while denying that the self somehow 'tags,' or is essentially intermingled with, the ordinary contents of experience (that is, aches, pains, aversions, and so on). One way we might be aware of being the subject of our experiences is through *awareness* itself, or the experiencing relation. But this limited degree of self-awareness doesn't entail awareness of a self that impinges on objects of awareness and their phenomenal character. It certainly doesn't entail the presence of phenomenological markers in experience that track distinctive features of the particular subject undergoing the experience. One reason for thinking that the self makes no such "substantial" contribution to phenomenal content is that we can conceive of two subjects sharing a brain, experiencing the same token pains and aversions, such that their experiences are

²¹ For some criticisms of the Humean view, see Duncan (2018).

phenomenologically indistinguishable. In the grip of a dull headache, it is not hard to imagine another agent simultaneously experiencing the headache, perhaps in complete ignorance of the fact that the headache she experiences is shared by another. Each subject might experience the headache *as their own*, but not in a way that prevents the other from standing in the same experiencing relation to the headache. There may be physical barriers to such a scenario, but its conceivability seems hard to deny.²² And the fact that we can conceive of multiple, co-located selves, sharing one and the same phenomenal experience, suggests that the self does not impinge on the phenomenology in the sense at issue. If we were aware of some contribution we make as unique subjects to the phenomenal character of our headaches, the scenario would be hard to conceive.

One might worry that there is risk of conflating, here, multiple selves having qualitatively identical headaches, which is clearly conceivable, with multiple selves experiencing the same token-identical headache, which is less obviously so. Still, having clearly distinguished the two scenarios, the latter remains conceivable it seems to me. And I take the intuition to be broadly accessible. Certain meditative practices, for instance, are predicated on the possibility of recognizing, through sustained introspection, the aware subject as a distinct existent from the various aches, pains, and aversions that populate one's field of awareness. That's one perspective from which the scenario should be readily conceivable. But since intuitions might differ, and

²² For example, Zagzebski (2013) invites us to imagine an "omnisubjective" agent who simultaneously occupies the experiential perspective of all subjects at once.

because the scenario's conceivability plays an important role in the argument, it would be good to consider some reasons why it might *seem* inconceivable.²³

Difficulties conceiving the scenario may be based on a theory-laden interpretation of an important phenomenological difference between our awareness of bodily sensations, or interoceptive awareness, and our awareness of external objects and their qualities, or exteroceptive awareness. This general difference in phenomenal character is sometimes described in terms of sensing the self or some aspect of the self: the interoceptively perceived bodily ache and the perceiving subject are somehow essentially unified as objects of direct awareness. We *sensorily feel* a union, to employ an old Cartesian metaphor.²⁴ No such subject-object elision occurs in ordinary cases of exteroceptive awareness (say, witnessing the red and gold of autumn leaves). I suspect that the Cartesian interpretation of interoceptive phenomenology comes naturally to many and makes it harder to imagine multiple and essentially distinct subjects experiencing the same token headache. After all, if I regard *this* token headache as inseparable from me, it's hard to imagine it as the object of another's awareness.²⁵ But the assumed inseparability is not *sensed* and

²³ I thank an anonymous reviewer for pushing me to say more on this point. The argument should win wider acceptance if readers who struggle to conceive of the scenario can be convinced to, perhaps, see their interoceptive phenomenology slightly differently than they might be used to, and in a way that, I think, unlocks the conception.

²⁴ See Simmons' (2017: 18) interpretation of Descartes' famous sailor-in-the-ship analogy:

Through these sensations of pain, hunger, thirst, etc., nature also teaches me that I am present in my body not merely as a sailor is present in his ship, but rather that I am very closely joined and as it were intermingled with it, so that I compose a single thing with it (AT VII 81).

²⁵ What else besides subject-object union might make the scenario seem inconceivable? I suppose one could take the awareness relation itself to exclude more than one subject being aware of a token headache, as though it were like the relation of being one's lawful spouse. But that can't be right since we don't

the phenomenology doesn't entail it. We can appreciate this by considering a simpler account of the phenomenal contrast between interoceptive and exteroceptive awareness that seems equally consistent with its general character. On a view I find plausible, interoceptive awareness involves experiencing an object of awareness as subject-dependent: the felt headache doesn't seem like it could possibly exist without *some* subject experiencing it. We don't experience colors as subject-dependent in quite the same way. Chalmers (2006, 76-113) puts this point nicely:

The phenomenology of color does not seem to be the phenomenology of properties that require a perceiver in order to be instantiated. (The phenomenology of pain is arguably different in this respect...). It seems coherent to suppose that there is a world in which perfect colors are instantiated, but in which there are no perceivers at all. Can one conceive of one's ankle being in perfect pain without *anyone* experiencing the pain? It is not clear that we can.²⁶

But experiencing headaches as subject-dependent is not the same as experiencing headaches as essentially unified with, or inseparable from, a specific subject. The phenomenology of subject-dependence shouldn't prevent us from conceiving of multiple subjects experiencing the same token headache.

think of awareness as a "rivalrous good" generally (consider, e.g., visual awareness of colors). It certainly isn't entailed by the phenomenology of interoception for reasons I discuss.

²⁶ Chalmers' "perfect" colors and pains are supposed to be qualities or properties that are just as colors and pains phenomenologically *appear* to be in ordinary perceptual experience. For instance, colors appear to be primitive qualities pervading the surfaces of external objects that are directly present in awareness (49). Pain appears to be both a primitive quality or property of a body *and* something that requires a perceiver: "We might think of [perfect pain] as an intrinsic property that, if instantiated, necessarily 'broadcasts' further constraints on the world" (114).

Put more positively, the phenomenology of suffering, and of interoceptive experience more generally, *taken at face value* favors the hypothesis that an acquaintance with pain/aversion is not “substantially” self-involving. The hypothesis simply and straightforwardly explains the phenomenology.²⁷ And having motivated it, we have grounds for thinking that if A were acquainted with B’s pain and aversion, it would feel just like being in pain and being averse to the pain that is felt. A would be averse to B’s pain given that A would be experiencing a token aversion that takes B’s pain as its object. The fact that the pain and aversion are also experienced by B shouldn’t affect A’s experience in a way that might make us reluctant to attribute the aversion to A. There is no reason to think A would experience the aversion as *B’s aversion* (or the pain as *B’s pain*) since we do not experience our own pains/aversions as uniquely or essentially ours. If anything, A should experience the aversion *as their own* in the thin experiential sense at issue simply in virtue of their awareness. To be sure, A might have beliefs about the pain-aversion being simultaneously experienced by B based in part on the phenomenology. For example, the pain might be felt in a physical location outside of A’s body in a manner akin to “supernumerary” phantom limb pain (Halligan et al. 1993). But the knowledge that the pain is located external to one’s body (or in someone else’s body) is compatible with the pain being experienced *as one’s*

²⁷ The phenomenology might be *consistent* with the possibility that the self figures more substantially in the content of experience, as I’ve already acknowledged, but this would not entail that the fact that experience is the way it is fails to favor the self’s non-involvement. The hypothesis that the self does not feature more substantially in the phenomenology, paired with our observations about subject-dependence, offers a *simpler* and less metaphysically demanding explanation of what the experience of pain and suffering is like. For a related point about the evidential import of experience, see Skow (2011, §3): “This is an instance of a general epistemological truth: a body of evidence can support some hypothesis without entailing that hypothesis.”

own. Indeed, the patient suffering from supernumerary phantom limb very much finds herself in pain and averse to the pain she feels.²⁸

Given our earlier assumptions about acquaintance, the claim I'm defending can be construed epistemically: if A were to have knowledge (by acquaintance) of B's pain and aversion, then A would, necessarily, be averse to B's pain. Moreover, so long as A intends to preserve their knowledge_A of the sensible content of B's experience, A must ameliorate B's suffering at least insofar as (1) A's aversion to the felt pain dominates A's other desires, and (2) A is instrumentally rational (that is, desires the necessary means to the satisfaction of their final desires). At a minimum, we can draw the perfectly general conclusion that an agent with knowledge of another's

²⁸ My argument has been based principally on introspective and phenomenological evidence. Could one appeal to other sources of evidence to motivate the main point? An anonymous reviewer helpfully draws attention to the case of Krista and Tatiana Horgan, conjoined twins whose brains are connected by a band of neural tissue joining each twin's thalamus (Pyke 2017). The twins seem to have extensive introspective and experiential access to the other's experience. Cochrane (2021: 1022-23) describes several striking incidents:

The twins also demonstrate a common receptivity to pain. When one twin's body is harmed, both twins cry...At one point the twins experience a headache (Pyke 2017: 12.00–12.48). Though it is hard to tell from the clip, it appears to hit Tatiana first because Krista reports "my sister has a headache" but immediately afterwards the clip shows both twins expressing distress, including a precisely synchronized howl of pain.

In another incident, the twins fall while sledging and Krista reports that she felt her sister's pain but that it didn't hurt her. One interpretation is that Krista was fully aware of Tatiana's pain without being affected by it aversely. Alternatively, Krista might have experienced a milder pain, or the thalamic bridge was not robust enough to convey all affective information, which would explain the difference in their aversions (1023). The latter seems more likely given that their "mother ... notes that both twins [eventually] cried and Krista agree[d] that she was crying because her sister hurt her buttock." Ultimately, I'm not sure what to make of this fascinating case. It is possible that the twins had direct awareness of each other's pains and aversions to some extent.

suffering would, necessarily, feel *some* motivation to ameliorate the other's suffering. It may not be a dominant desire in all cases, but there is no possible subject such that if that subject were acquainted with another's suffering, they would fail to be motivated (to some degree) to alleviate the other's pain. We should acknowledge, of course, that motivation-talk can be ambiguous. It can refer to a disposition to act or, as I noted earlier, phenomenally salient aversions and affective desires with which we are acquainted (see e.g., Johnston 2001, p188). What I'm claiming is impossible, given our assumptions, is for a subject to know by acquaintance another's suffering without experiencing a *felt* aversion to the suffering or a *felt* desire to stop it (with the same felt force as the sufferer's aversion). However, since the disposition to act isn't an acquaintance state, it remains possible to know another's suffering without taking any action. The disposition to act may require additional non-epistemic assumptions about the agent, such as means-ends rationality, or, as I'll discuss later, some additional epistemic assumptions. Nevertheless, we've made manifest progress by explaining in purely epistemic terms an agent's lack of felt care and concern for the suffering of others.

The argument's success does not depend on the knowledge in question being physically possible. It may well be impossible to become directly acquainted with someone else's token phenomenal qualities, in which case it would be impossible to know certain aspects of the world. However, the impossibility of knowing_A another's suffering would remain explanatorily relevant, for a scenario's impossibility does not prevent it from being explanatorily useful. To borrow one of Bernstein's (2016) examples, Sam's failure to win the Fields Medal might be explained by the fact that it is impossible to prove the proposition Sam was trying to prove (cf. Baron & Colyvan 2020). Likewise, the impossibility of knowing_A the suffering of others sheds explanatory light on the indifferent agent's failure to care.

Our epistemic explanation of indifference to others' needs shouldn't be viewed as a substitute for other, more familiar explanations—for example, one that appeals to an agent's inability to experience ordinary empathy. The alternative explanations needn't compete. That Susan proved a long-standing mathematical conjecture provides one explanation for why Sam didn't win the Fields Medal without undermining the explanatory significance of the impossibility of Sam proving his favorite mathematical claim. Likewise, ordinary explanations for an agent's failure to care about the suffering of others do not undermine the explanatory significance of the agent's lack of knowledge.

Other concerns might be raised, however, about the fact that the epistemic condition that is sufficient for experiencing motivation to help others turns out to be impossible for anyone to meet. For one, it seems to imply that the way in which the callous and indifferent have gone wrong, epistemically, is not something anyone can be blamed for, since it was impossible to cure the underlying epistemic defect and blame arguably assumes an ought-implies-can principle. Fortunately, there are other, more accessible ways to act morally that remain grounded in epistemic virtue, as I argue below. The goal of this section was limited to establishing, on grounds that even the amoral and utterly indifferent can accept, that a species of ignorance explains an agent's failure to care about the suffering of others. We have made progress insofar as we have succeeded in showing that our perspective on the world which informs our choices is anything but objective or complete, and our inability to be fully objective—to have a complete perceptual grasp (and so knowledge_A) of what's real—explains an agent's partiality.

§3 Empathy as an epistemic virtue

Let us recap the argument so far. If we adopt (a) the philosophically defensible view that perception grounds a species of non-propositional knowledge of things perceived, along with (b) a complementary metaphysics of suffering, we discover a basis for thinking that callousness and indifference can be explained as a form of ignorance. The explanation appeals to an essential feature of having knowledge_A of other people's suffering: it constitutively involves feeling averse to the suffering of others, given phenomenologically supported assumptions about the content of suffering.

A potentially counterintuitive upshot of the view, however, is that the morally conscientious and the morally indifferent appear similarly situated since both lack knowledge_A of other people's suffering. One might have expected—indeed, hoped—that the epistemic deficiency that explains inadequate concern for others would not simultaneously cast doubt on the epistemic standing of the morally virtuous. However, it might have been a mistake to assume, at the outset of inquiry, that the epistemic defect that explains radical moral indifference might not also explain entirely ordinary ways of falling short of moral standards. For if morality is as demanding as the utilitarian maintains, requiring us always to act in ways that maximize aggregate happiness and minimize suffering, then it is a feature of the view, not a bug, that ordinary agents lack the kind of knowledge that is sufficient for ideal moral compliance given the scale and magnitude of suffering in the world.²⁹ That said, we should be able to explain why the morally conscientious do *better*

²⁹ Traditionally, vindicating ordinary decency has not been seen as a requirement for a successful demonstration of the objectivity of moral demands. Kant, for example, famously maintains that the ideal of the morally good person free of rational incoherence is instantiated by no one. For Kant, rational incoherence is pervasive (*Religion* 6:22-23). True appreciation of and conformity to moral demands takes a revolutionary act of subordinating the maxims of self-love to the imperatives of morality, an act that is

epistemically than the radically amoral, and explaining the asymmetry involves taking into account the different ways in which we can epistemically evaluate agents.

It is tempting to suppose that in addition to the familiar epistemic goods like truth, knowledge, and understanding, there are related epistemic virtues (Zagzebski 2001; Roberts & Wood 2007). These include traits conducive to the discovery of truth, the avoidance of error, and the pursuit of other intellectually valuable ends. It is widely acknowledged, for example, that epistemically virtuous agents seek the truth, apportion their beliefs to the evidence, and do not exhibit various pathologies of reason that interfere with knowledge acquisition. Far less clear is what epistemic virtue demands in the case of the unknowable, but a familiar and plausible line is that epistemically virtuous agents aspire to approximate knowledge. For example, even if the fundamental truths about physical reality turn out to be unknowable, we should nevertheless aim to believe scientific theories that approximate the truth (Maher 1993, Oddie 2008). The notion of approximate truth is often motivated based on the need to explain the character of scientific progress (Boyd 1988). And the question this raises is whether there might be an analogous phenomenon in the case of acquaintance knowledge.

It seems to me that we can provide a promising account of what approximately knowing_A another's pain amounts to and why the approximation might be epistemically desirable. Several moral philosophers have suggested that empathy and compassion are epistemically desirable states insofar as they involve accurately representing the mental states of others (Boyd 1988: p341;

incompatible with even momentary lapses in doing what is right. As a result, the "holy will" remains out of reach for ordinary mortals, but nevertheless serves as a kind of regulative ideal: "It is a human being's duty to strive for this perfection, but not to reach it . . . and his compliance with this duty can, accordingly consist only in continual progress" (*Metaphysics of Morals*, 6:446). As I argue below, knowledge_A of others' suffering can serve as a regulative ideal for epistemically virtuous and empathic agents.

Marshall 2019: ch4).³⁰ The epistemic achievement involved in empathy is best characterized, in my view, not in terms of representational accuracy or propositional knowledge, but based on the fact that empathy simulates the state of knowing by direct acquaintance another's pain in a way that affords significant opportunities for cultivating epistemic virtue. When A experiences empathic distress after coming to believe that B is in pain, A becomes acquainted with, and so

³⁰ Like Marshall, I take compassionate attitudes to be epistemically evaluable, but our accounts diverge on the explanation. Marshall argues that there is an epistemic good—“being in touch” with another's pain—that only the compassionate realize. Marshall characterizes “being in touch” in terms of representational accuracy: “Compassionate reactions... significantly resemble pains like the wombat's, and a representation's resembling pain is necessary and sufficient for it to reveal pain, and so for letting the subject know *what the property is like*” (p68). Marshall's account is attractive but confronts several important challenges. Knowing what pain is like, being propositional knowledge, does not seem to depend on the subject having a concurrent painful experience, even if it depends on having previously experienced pain (Lerner 2019). Accordingly, a view like Marshall's needs to explain why having discovered what pain is like, the compassionate agent shouldn't just distract herself or walk away without helping. Moreover, it remains unclear on Marshall's view why an agent who feels compassionate distress but “looks away” without helping fails to accurately represent another's experience. After all, the agent experiences an aversion that is directed at the distress she feels, and, so, her experience accurately represents the ‘self-directed’ character of other people's aversions. Marshall (2019, p48 fn. 47) acknowledges the worry, though I'm not sure he does enough to relieve us of it.

The view I defend differs in key respects and avoids the above challenges. First, on my view, a key epistemic condition that is sufficient for being motivated to help others—acquaintance with their pain—isn't realized in compassion. On the contrary, it may well be *impossible* to be relevantly acquainted. Nevertheless, we can invoke this epistemic condition to explain what the morally indifferent are “missing,” and so dispel the skeptic's misimpression that their choices are informed by a fully objective take on the world (see section II). Second, on the present proposal, empathic distress is epistemically useful not because it allows us to correctly represent the pains of others or even to have acquaintance knowledge of another's pain. Empathy's value lies in the opportunity it affords epistemically virtuous agents to act with intellectual courage in relation to the unknowable_A experiences of others. For reasons I explain below, the empathic agent cannot just mitigate her own distress while the painful object of another's acquaintance endures because it would fall short of the demands of intellectual courage.

knows_A, her own pain. The knowledge_A gained through empathy is unpleasant and difficult to bear just as genuine knowledge_A of B's pain would be. And this similarity between the two states affords A an opportunity to demonstrate the virtue of intellectual courage, or so I believe we can argue.

First, the nature of the relevant virtue invites clarification. On Roberts and Wood's (2007) account, intellectual courage consists in the ability to respond appropriately to threats to one's intellectual goals and epistemic wellbeing. Such threats include the temptation to believe what is convenient. Baehr (2011) associates the virtue with a willingness to inquire after the truth even when it might be difficult to bear. To my mind, while intellectual courage can take various forms, what makes it a distinctly epistemic virtue is that it is often motivated by love of knowledge, or a final desire to know *for its own sake*.³¹ When this final desire is more robust than, say, a final desire for comfort, it manifests as an unwillingness to sacrifice knowledge at the altar of convenience. Hopefully, these claims strike the reader as intuitive. I take intellectual courage in the sense defined—as a disposition, sustained by a final desire for knowledge, to act in ways that are knowledge-preserving even when the knowledge gained is uncomfortable—to be very much part of our ordinary repertoire of epistemic concepts.

Armed with an account of the virtue, we can explain why empathic distress affords A an opportunity to demonstrate intellectual courage. Consider, first, how A's commitment to knowledge-preservation would be tested if A, counterfactually, had direct experiential access to B's inner mental life. The acquaintance knowledge would be hard to bear since A would be

³¹ As I discuss below, the argument of this paper is compatible with other accounts of epistemic virtue, including an account on which being motivated by knowledge's instrumental value is sufficient for intellectual courage.

acquainted with B's suffering. By alleviating B's pain, A makes her experiential insight into B's inner mental life more bearable without interfering with her capacity to know. In other words, the choice is compatible with intellectual courage. By contrast, asking a neuroscientist to disable A's 'mind-viewing' capacities would not be. By analogy, an epistemically virtuous agent immune to wishful thinking might take actions in the world that aim at making the facts more bearable but will not succumb to the temptation to ignore or deny hard truths. In the actual world, A lacks such mind-viewing capacities, but her empathic distress serves as an imperfect substitute. In the grip of empathy, it can *seem* as though one is directly acquainted with an unpleasant aspect of another's mental life. As Adam Smith writes concerning the empathic imagination, "[b]y the imagination,... we enter as it were into [another's] body, and become in some measure the same person with him, and thence form some idea of his sensations, and even feel something which, though weaker in degree, is not altogether unlike them."³² Imaginative empathy can therefore activate, in epistemically virtuous agents, the same intellectually courageous dispositions they would exhibit if they knew by direct acquaintance the suffering of others. When A helps alleviate B's pain instead

³² *The Theory of Moral Sentiments* (I.i.1.2). Along the same lines, Smith observes:

But though sympathy is very properly said to arise from an imaginary change of situations with the person principally concerned, yet *this imaginary change is not supposed to happen to me in my own person and character, but in that of the person with whom I sympathize*. When I condole with you for the loss of your only son, in order to enter into your grief I do not consider what I, a person of such a character and profession, should suffer, if I had a son, and if that son was unfortunately to die: but I consider what I should suffer if I was really you, and I not only change circumstances with you, but I change persons and characters. (VII.iii.1.4)

For a helpful discussion of Smith's view, see Bailey (2016). Smith's observations strike me as most accurate in the case by empathic distress experienced in response to the suffering of close family and friends, especially children.

of tending to her own empathic distress, A manifests her intellectual courage: a desire and ability to preserve knowledge_A of B's experience notwithstanding the practical costs.

There are, of course, less intellectually resilient ways of responding to another's suffering. Suppose A ignores B's pain and tends to her own empathic distress while B's pain persists (say, by distracting herself or by taking an anti-empathy pill). I submit that the choice would demonstrate the same kind of epistemic weakness as a willingness to believe convenient falsehoods or a resistance to believing readily available truths. For as long as B's pain endures, opting out of empathic distress betrays a subject's willingness to prioritize practical comfort over (apparent) knowledge_A of another's pain. The only psychologically viable means of meeting the epistemic challenge posed by other people's suffering as an empathic agent—one who experiences distress at another's suffering that is imaginatively felt *as the other's distress*—involves intervening in the world to relieve other people's suffering. It is possible of course that A is under no perceptual illusions about being acquainted with another's pain, and so opts out of empathic distress because it is merely ersatz and not genuine knowledge_A of B's pain. The way things *seem* to A in empathy, on this picture, corresponds to the fact that A is acquainted only with A's own inner mental life (her own felt distress). As I've already indicated, this objection seems false to the phenomenology of empathy. But even if it's true that some people experience empathy without the externalizing seeming, there might still be reasons for them to try to imaginatively access the phenomenology I claim is essential to full-blooded empathy, if for no other reason but to test and exercise their intellectual courage. On the view I am defending, a kind of literal-mindedness when empathizing with others hinders the cultivation of epistemic virtue. Opting out of empathy deprives agents of the opportunity to exercise intellectual courage in relation to other people's suffering.

The argument turns on the fundamental claim that choosing not to be acquainted, even if only by empathic approximation, with another's suffering represents a kind of intentional ignorance or epistemic vice. But the objection might be raised that we are constantly losing and gaining acquaintance knowledge of a broad range of things as we navigate the world, and it would be odd to think that we demonstrate an epistemic vice whenever we look away from one object to perceive another.³³ Fortunately, our argument does not have this implication since there is an important difference between ordinary shifts in perceptual attention and an unwillingness to perceive another's suffering. We have limited capacities for knowledge_A acquisition and, so, it makes sense to perceptually prioritize things that directly engage our interests. But that is different from shifting attention away from other people's suffering because the associated acquaintance is uncomfortable or hard to bear. Giving up on knowledge because what is known is hard to bear involves being *averse*, and not just indifferent, to what is known. And the epistemic relevance of this distinction is evident in our ordinary evaluation of agents. Normally, we do not criticize agents who fail to retain truths they find uninteresting, whereas we do think that a resistance to believing truths one finds uncomfortable, or a desire to believe regardless of the truth, amounts to an epistemic vice: the vice of wishful thinking or willful ignorance (see e.g., Heil 1984, p64). In the latter case, the fault lies not in our limited capacities for knowledge-acquisition and -retention, but in our inability to bear certain truths.³⁴ And while it's easy enough to find pleasant opportunities

³³ I thank an anonymous reviewer for pressing the point.

³⁴ This distinction should help us see why the view does not imply that it's epistemically courageous to devote one's cognitive efforts to counting "blades of grass in ... park squares and well-trimmed lawns" (Rawls 1971: 432). Widespread suffering, unlike trivial facts concerning the contents of one's lawn, represents an aspect of what's real that we'd prefer not to be aware of, and not just because we're interested in other things. Many of us would turn away from the suffering of others *even if* it did not require much effort or forgoing other knowledge to acquire knowledge of others' suffering. In general,

for cultivating concentration, retention and other epistemic virtues, opportunities for the development of intellectual courage require empathic distress and imagination when only a small corner of reality is known through direct acquaintance.³⁵ And intellectual courage is worth cultivating not just because, like other virtues, it is of intrinsic epistemic worth, but because without this general disposition, our epistemic well-functioning is seriously and regularly threatened in a world as non-ideal as ours. Sadly, human and non-human suffering is not a local or unusual phenomenon. It is terribly widespread. Thus, without the ability to confront it, we will often find ourselves turning away and so willfully choosing a distorted view of things: a crude and painless vision of what the world is like.

A different objection to our account of the epistemic importance of empathy might be based on the normal motivations of empathic agents. Presumably, empathic agents are ordinarily moved to hold on to their awareness of another's suffering, despite their own discomfort, out of care and respect for the other, not because they desire to know another's pain for its own sake. And this kind of empathic resilience, which we might call "moral courage," might not seem like *intellectual* courage, at least given the account of epistemic virtue I've adopted in this paper which requires being moved directly by the final value of knowledge. However, the reasons that motivate sustained empathy in ordinary cases are likely complex. The value an agent places on knowledge

knowledge-acquisition and -retention require effort and skill, but not of the same kind that's involved in facing up to knowledge we find hard to bear. Thanks to a reviewer for a productive exchange on this point.

³⁵ As Johnston (2007: p235) memorably puts it, "we are... highly selective Samplers of Presence," where "Presence" refers to the immediate perceptual presence before numerous subjects of real "objects, qualities or whatever" that "disclose some aspect of their nature" (233). What I'm pointing out, here, is that our selective sampling relieves many of us of the burdens of knowledge and, simultaneously, undercuts our capacity for epistemic virtue.

for its own sake, or on having an undistorted view of the world, can reinforce her capacity to be morally courageous, and in ways that needn't be motivationally transparent, especially given that a final desire to know and understand another's experience is often inherently a way of caring about them. At the very least, an agent capable of sustained empathy for moral reasons does not so clearly exhibit the epistemic vice manifested by those who turn away from the suffering of others. For the morally courageous agent treats sustained knowledge of another's pain as *instrumentally* valuable, as a means of expressing care and solidarity, in resisting the temptation to turn away, a disposition that may qualify as a form of intellectual courage on other accounts of epistemic virtue that are compatible with the argument of this paper.³⁶

To sum up, empathic experience and imagination enable possibilities for epistemically virtuous action. On the present proposal, we should not think of empathy as though it were a strategy intentionally adopted by knowledge-maximizers. Empathic distress combined with a willingness to help others is an expression of epistemic virtue to the extent that it reflects an agent's commitment to prioritize knowledge over practical comfort. Hence, while the empathic agent and the radically amoral both labor under an ignorance that might explain their general failures to be precisely as helpful towards others as morality demands (and morality may be quite demanding), only the empathic agent's response to what is unknown_A demonstrates greater epistemic virtue—the fortitude to know_A hard reality by approximation.

§4 Questions and clarifications

I have addressed questions and concerns about the proposal as they have come up. However, questions remain that serve as opportunities for further clarifying the proposal. The

³⁶ I thank a reviewer for an illuminating exchange on this point.

following discussion is organized around doubts one might raise about: (a) the background epistemology on which the argument depends; (b) the overall value of acquaintance knowledge and whether its sufficient to motivate agents to care about others; and (c) whether being motivated to relieve other people's suffering for epistemic reasons could ever be sufficient for moral virtue.

(a) The background epistemology

The proposal's main vulnerability is that it rests on a controversial set of claims about the epistemic significance of perception, claims that might reasonably be questioned. However, the fact that its starting points occasion controversy does not render the proposal speculative, unproductive, or alien to philosophy. Philosophical arguments often rely on controversial starting points. And while direct arguments in favor of the background epistemology, which I have offered elsewhere, would take us well beyond the scope of this paper, the present proposal amounts to an indirect defense based on the theoretical work the epistemology can do for us. We have discovered overlooked reasons for taking the acquaintance approach to knowledge seriously because it can help us make progress in ethics. At the very least, the argument I have offered should be of interest to those who find the epistemic claims independently plausible.

Another reason why the assumptions at the proposal's core should not detract from the progress we have made is that they cannot be accused of being either ad hoc or question-begging. I know of no ethical or metaethical reasons for denying that acquaintance delivers non-propositional knowledge, or for rejecting the underlying metaphysics. Meanwhile, traditional efforts at explaining why indifference towards others involves some kind of error or confusion have tended to rely on assumptions that are not just controversial but "internal" to ethics in the ways highlighted earlier. The Platonist invokes inherently prescriptive moral facts. The Kantian makes sweeping assumptions about agency (e.g., that we can only choose actions under the guise

of the good or by regarding them as *free and principled* actions) in an effort to convict the selfish and amoral of a practical contradiction (that is, of acting contrary to their beliefs). An argument that aspires to convince moral skeptics that they're missing something is more likely to be effective if it's based on premises that do theoretical work outside of ethics.³⁷ And the present argument fits the bill, even if several of its assumptions invite further examination and defense.³⁸

Finally, while some of the assumptions I've made are generally contested, they have the virtue of being easily grasped. As Street (2019) points out, a successful vindication of our inchoate sense of the connection between moral and epistemic failure should ring true to "our pre-philosophical sense of things." I submit that the argument's (relative) simplicity, especially at the level of its basic assumptions, constitutes evidence of its success at capturing ordinary ethical intuition.

(b) The value of knowledge_A

A different challenge involves doubts that are internal to the epistemic framework. One might reasonably question the value of knowing_A another's pain. Even if felt acquaintance with other people's suffering counts as knowledge, its value might seem trivial in particular instances, especially once it is admitted that knowing the truth that others are suffering, or truths about what

³⁷ See the earlier discussion of the skeptic's challenge and why it is worth taking seriously in ethics.

³⁸ An objection might be raised that while the argument does not rely on any controversial ethical or metaethical assumptions, it does rely on meta-normative premises concerning the nature of knowledge. If the skeptical challenge in ethics is best seen as a meta-normative one, with the ethical being one instance of the normative, then the argument is question begging. We can certainly imagine a radical skeptic who treats all epistemic categories as defective, but the skeptical challenge that has been our target is reasonably construed as taking the epistemic domain for granted. Epistemic normativity is not nearly as controversial as moral normativity, and the possibility of being a skeptic about truth, knowledge, or epistemic reason does not detract from the philosophical progress we have made.

suffering feels like, does not require being relevantly acquainted. And if knowing_A another's suffering is of limited value, it becomes unclear why it should motivate anyone towards other-regarding action.³⁹

There are two kinds of replies we can give to this complex challenge. The first contests the assumption that the knowledge at issue is of trivial value. One measure of the value of knowledge is the difference it makes to one's choices, and by that measure the value of knowing_A another's suffering could hardly be trivial given how radically it would alter our behavior towards others. In fact, we've discovered not just that this knowledge would shape our choices, but that its substantial impact would be experientially intelligible. We've also observed that knowing another's suffering is a means of expressing care and solidarity and, so, of great moral and interpersonal value. There are still other ways of contesting the triviality charge. Recall that acquaintance is supposed to be an epistemic good because it affords a direct cognitive grasp of aspects of reality. As Coleman (2019, p59) puts it, "in acquaintance, awareness achieves metaphysical and epistemic connection with another existent... Had it not enlarged the subject's world in this way, acquaintance's claim to be a kind of knowledge would be jeopardized." Given this, we shouldn't expect the distinctly epistemic value of being acquainted with someone's else experience to be introspectively

³⁹ Blackburn (2006: p157) raises a related but more general concern about attempts to ground moral virtue in epistemic norms or practical rationality:

The popularity of rationalism, and the general feeling that there 'must be something to' [rationalist argument] are very deep-rooted. Partly, they represent a noble dream. They answer a wish that the knaves of the world can be not only confined and confounded, but refuted – refuted as well by standards that they have to acknowledge.... We sentimentalists do not like our good behaviour to be hostage to such a search. We don't altogether approve of Holy Grails. We do not see the need for them. We are not quite on all fours with those who do. And we do not quite see why, even if by some secret alchemy a philosopher managed to glimpse one, it should ameliorate his behaviour, let alone that of other people.

transparent, for the degree to which an acquaintance enlarges a subject's world needn't be obvious from introspection. More importantly, the *final* epistemic value of knowing_A another's suffering does not have to be especially high for the knowledge_A to have considerable *instrumental* epistemic significance as a means of cultivating intellectual courage.⁴⁰ As I have argued, empathic agents who are also epistemically virtuous will be disposed to resist the desire for comfort insofar as it interferes with their (approximate) knowledge_A of, or quasi-acquaintance with, other people's suffering. In these and other ways, the argument has clear potential for engaging (and reinforcing) an agent's motives towards other-directed action.

A different reply emphasizes that the argument's success does not turn on its ability to motivate agents to be moral. It would be a mistake to regard the problem of convincing those who are indifferent to others' needs that their lack of objectivity explains their indifference as identical to the problem of finding an argument that motivates them to be better. The two problems are distinct, and our discussion has been principally concerned with the former.⁴¹ The dialectic in which we have intervened is sparked by the indifferent agent's demand to know from a suitably neutral standpoint why indifference to the suffering of others is epistemically deficient. That demand we've addressed: those who are indifferent to the suffering of others lack knowledge and intellectual courage in ways that explain their indifference. And by addressing the skeptic's

⁴⁰ See the earlier discussion of why intellectual courage, unlike other epistemic virtues such as curiosity or productive concentration, requires empathy and why it matters in a non-ideal world.

⁴¹ That rational argument alone might persuade agents to be good and decent is considerably more controversial than the idea that agents might be shown that callousness and indifference are grounded in mistake or ignorance. It is no part of the Platonist's view, for example, that the morally indifferent will necessarily be moved to moral action once they admit their ignorance of inherently prescriptive moral facts. Knowledge of their own ignorance needn't move the morally indifferent because, according to the Platonist, it is the moral facts that (if known) essentially motivate, not the fact that one does not know the moral facts.

demand, we have undercut any false epistemic pride the skeptic might take in their lack of motivation to help others, which should, in turn, render the skeptic more open to other, moral arguments that convince them to be better. However, solving the skeptic's motivational problem was not the main goal. Furthermore, we can make sense of the stakes from the point of view not of the skeptic but of ordinary, morally motivated agents. As Street and others have suggested, many morally conscientious agents take their own concern for others to be an expression of their objectivity. The present account attempts to vindicate that self-conception, and success in the enterprise does not turn on whether the account motivates *amoral* agents to be better.

(c) The elements of moral virtue

Finally, a question that we haven't yet fully addressed concerns the character of moral motivation. Genuine moral motivation, one might argue, involves caring about the pains of others despite not experiencing those pains as one's own. The hypothetical agent with knowledge_A of others' suffering may be motivated to help but only because helping brings an end to her own suffering, a mode of concern that seems selfish rather than other-regarding.⁴² In the same vein, the account portrays the empathic as having a reason to help others because helping is the only means

⁴² One might worry, also, that the knowledgeable_A agent has some vicious desires. An agent capable of experiencing other people's felt desires might experience some desires to inflict suffering, for there are agents who are not just indifferent towards others, but actively immoral in virtue of having final desires to hurt others. However, consider the utilitarian's radically impartial manner of concern for the happiness of others. The utilitarian places some weight on the desires of even immoral agents, but the happiness of such agents is outweighed by the pressing needs of others. If the scope of our hypothetical agent's concern resembles that of the utilitarian's, it cannot be assumed without argument that her moral integrity is compromised by her limited sympathy for vicious agents. Furthermore, vicious agents have final desires that are likely predicated on ignorance of others' suffering, and so there may well be epistemic reasons to avoid empathizing with their distorted point of view.

of satisfying the complex goal of preserving one's ability to (approximately) know_A while mitigating one's own distress, which does not seem like a *moral* reason to help.

There are several things to say in reply. First, as clarified at the outset, our aim was not to show that knowledge and epistemic virtue are sufficient for ideal moral virtue. Rather, our aim was to show that knowledge and epistemic virtue are sufficient for motivating compliance with an important moral duty. We have made progress even if we have not quite demonstrated that the knowledgeable agent is worthy of moral praise in light of her reasons for acting. A more general point seems worth emphasizing in this context. Many moral requirements are not just about minimizing other people's pain. If we can wrong others in ways they cannot feel—for example, by lying to them—knowledge and the associated virtues will not necessarily motivate us to avoid those wrongs (though, even here, I suspect that there is a connection with the epistemic good). I have argued in this paper that there is at least one moral duty—namely, the duty to relieve suffering—such that having knowledge and, in such knowledge's absence, intellectual courage is sufficient for being motivated to comply with the duty. But this does not entail that the epistemically virtuous agent ends up with the full range of moral virtues.

Second, our motives in action are often complex. The standard of selfishness being wielded in criticism threatens to convict ordinary agents of selfishness who desire to help in part because the suffering of others causes them a personal distress that they intend to alleviate by helping. It also undercuts views on which the fact that helping others is conducive to personal happiness is a perfectly good reason to help. Indeed, the precise character of other regard is morally disputed, but the general claim that one ought to care about the suffering of others (for whatever reason) is not. This general and less controversial claim invites vindication on epistemic grounds quite apart from whether it satisfies a contested and demanding standard for selfless action.

Third, there are viable views on the nature of selfishness that don't entail that the concern displayed by the hypothetical agent with knowledge_A of others' suffering must be objectionably self-regarding. I've argued elsewhere that selfish choices make essential reference to the self: an outcome is chosen by an agent under a description that's essentially identity-dependent (Atiq 2016). For example, Susan wants to make a major scientific breakthrough and hopes that none of her peers will. It is an essential feature of her desired outcome that *she* be the one responsible for the breakthrough and not someone else. By contrast, a brute aversion to the pain one is acquainted with is not so obviously identity-dependent or self-privileging. Although it might be true that I care to stop the pains I experience because they are my pains, the claim should not be misinterpreted. It draws attention to the explanatory relevance of my *awareness* of the pain, and not necessarily the fact that it is *my* pain (even if it is and I experience it as such). A pain's *being one's own* can remain in the periphery of one's practical experience, and not something one attends to or cognizes as a normative or motivating reason for action. Put differently, the pain-alleviating action needn't be chosen under the guise of stopping *a pain of one's own*. Instead, it can be chosen as such actions often are: as a means of stopping *this* pain. For reasons discussed in section two, I find the idea that our aversions to pain track some subject-involving feature that's given in the phenomenology suspect. Even if awareness always involves some degree of self-awareness, the latter is phenomenologically peripheral. Hence, if the earlier discussion is on the right track, then A's becoming acquainted with B's felt aversions to B's pain needn't make the aversion A thereby feels objectionably self-focused. A can be moved to act by the identity-independent aversive phenomenology alone, and without any self-elevating judgment that one's own wellbeing matters more than the wellbeing of others. In short, A's motivational psychology and normative judgments may qualify as paradigmatically impartial on the right view of impartiality.

I have argued that recent work in epistemology and the philosophy of perception warrants greater attention from moral philosophers searching for a robust explanation of why indifference to the suffering of others amounts to an epistemic defect. Even if reasoning or the facts won't logically compel us to show concern for others, a species of knowledge arguably would, and, in the absence of such knowledge, a variety of intellectual courage will. This conclusion invites a fuller defense than I have offered in this paper. But by precisely locating it within an independently motivated framework for understanding the structure and significance of perceptual experience, I have hopefully gone some way towards demonstrating its plausibility and clarified the terms for a more conclusive assessment. A long-standing ethical problem concerning the relationship between moral and epistemic failure seems much more tractable when viewed through the lens of an acquaintance-based epistemology. The framework deserves serious consideration from anyone in the market for a morally uncontested basis for thinking that indifference to the needs of others stems from a kind of ignorance of their experience.

Acknowledgements: I thank Olivia Bailey, Matt Duncan, David Enoch, Rachana Kamtekar, Julia Markovits, Andrei Marmor, Colin Marshall, and Rachel Schutz for their helpful comments on the paper and for illuminating conversations on its main themes. I am grateful, also, for the invaluable feedback I received from two anonymous reviewers for *Philosophical Studies*. The paper in its final form reflects their thoughtful suggestions.

References

- Armstrong, D.M. (1962) *Bodily Sensations*. London: Routledge.
- Atiq, E. (2016) “How to be Impartial as a Subjectivist.” *Philosophical Studies* 173:757–779.
- (2021) “Acquaintance, Knowledge, and Value.” *Synthese* 199:14035-14062.
- Atiq, E. & Duncan, M. (2024) “I Feel Your Pain: Acquaintance & the Limits of Empathy.”
- In Kriegel, U., (ed.), *Oxford Studies in Philosophy of Mind* 4. Oxford: Oxford University Press.
- Aydede, M. & Fulkerson, M. (2019) “Reasons and Theories of Sensory Affect.” In *The Philosophy of Pain: Unpleasantness, Emotion, and Deviance*, edited by Bain, D., Brady, M., & Corns, J. New York, USA: Routledge. pp. 27–59.
- Baehr, J.S. (2011) *The Inquiring Mind: On Intellectual Virtues and Virtue Epistemology*. Oxford: Oxford University Press.
- Bailey, O. (2016). “Empathy, Care, and Understanding in Adam Smith's The Theory of Moral Sentiments.” *The Adam Smith Review* 9.
- Baron, S. & Colyvan, M. (2020). “Explanation Impossible.” *Philosophical Studies*
- Benton, A.M. (2017) “Epistemology Personalized.” *The Philosophical Quarterly*. 67 (269):813-834.
- Bernstein, S. (2016) “Omission Impossible.” *Philosophical Studies* 173 (10): 2575-2589.
- Blackburn, S. (1984) *Spreading the Word*. Oxford: Oxford University Press.
- (2006) “Must We Weep for Sentimentalism?” In J. Dreier (ed.), *Recent Debates in Moral Theory*. Blackwell Publishers.

- Boyd, R. (1988) "How to be a Moral Realist." In G. Sayre-McCord (ed.), *Essays on Moral Realism*. Cornell University Press, pp. 181-228.
- Brewer, B. (2011) *Perception and its Objects*. Oxford University Press.
- Broad, C. D. (1959). *Scientific Thought*. Paterson, NJ: Littlefield Adams
- Carruthers, P. (2004). "Suffering without subjectivity." *Philosophical Studies* 121, 99–125.
- Chalmers, D. (2002) "Content and Epistemology of Phenomenal Belief." In Q. Smith & A. Jokić (eds.), *Consciousness: New Philosophical Essays*, Oxford, OUP.
- (2006). "Perception and the fall from Eden." In Tamar Gendler & John Hawthorne (eds.), *Perceptual experience*. New York: Oxford University Press. pp. 49—125.
- (2013) "The Contents of Consciousness: Reply to Hellie, Peacocke, and Siegel," *Analysis*, 73(2), 345–368.
- Cochrane, T. (2020). "A case of shared consciousness." *Synthese* 199 (1-2):1019-1037.
- Conee, E. (1994) "Phenomenal Knowledge." *Australasian Journal of Philosophy*. 72: 136—50.
- Crane, T. (2012) "Review: Tye on Acquaintance and the Problem of Consciousness." *Philosophy & Phenomenological Research*. 84: 190—198.
- Duncan, M. (2020) "Knowledge of Things." *Synthese*. 197 (8):3559-3592.
- (2021) "Acquaintance." *Philosophy Compass* 16:1-19.
- (2018). "Subjectivity as Self-Acquaintance." *Journal of Consciousness Studies* 25 (3-4):88-111.
- Farkas, K. (2019) "Objectual Knowledge." In (Knowles, J. & Raleigh, T., eds.) *Acquaintance: New Essays*. Oxford: Oxford University Press.

- Gertler, B. (2012) “Renewed acquaintance.” In D. Smithies & D. Stoljar (Eds.), *Introspection and Consciousness* (pp. 89–123). New York, NY: Oxford University Press.
- Giustina, A. (2022) “Introspective Knowledge by Acquaintance.” *Synthese* 200:128.
- Halligan, P., Marshall, J., & Wade, D. (1993). “Three arms: A case study of supernumerary phantom limb after right hemisphere stroke.” *Journal of Neurology, Neurosurgery, and Psychiatry*, 56, 149–166
- Heil, J. (1984). “Doxastic Incontinence.” *Mind* 93 (369):56-70.
- Hume, David (1711–1776), 1738 [1896], *A Treatise of Human Nature*. Reprinted, L.A. Selby-Bigge (ed.), Oxford: Clarendon Press.
- Jackson, F. (1982) “Epiphenomenal Qualia.” *Philosophical Quarterly* 32: 127–136.
- Johnston, M. (2004) “The Obscure Object of Hallucination.” *Philosophical Studies* 120(1-3): 113–183.
- (2001) “The Authority of Affect.” *Philosophy & Phenomenological Research* 63(1): 181–214
- (2007) “Objective Mind and the Objectivity of Our Minds.” *Philosophy and Phenomenological Research*, 75: 233-268
- (2011) “On a Neglected Epistemic Virtue.” *Philosophical Issues* 21 (1):165-218.
- Kvanvig, J. (2003). *The value of knowledge and the pursuit of understanding*. Cambridge University Press.
- Korsgaard, C. (2009) *Self-Constitution: Agency, Identity, and Integrity*. Oxford: Oxford University Press.
- Kriegel, U. (2024). “Knowledge-by-Acquaintance First.” *Philosophy and Phenomenological Research* 109 (2):458-477.

- Lerner, A. (2019) "Review of 'Compassionate Moral Realism,' by Colin Marshall." *Notre Dame Philosophical Reviews*.
- Lewis, D. (1990) "What Experience Teaches Us." In *Mind and Cognition: A Reader*, ed. by W. Lycan, (Oxford: Blackwells).
- Mackie, J.L. (1977) *Ethics: Inventing Right and Wrong*. Penguin Books.
- Maher, P. (1993) *Betting on Theories*. Cambridge: Cambridge University Press.
- Marshall, C. (2019) *Compassionate Moral Realism*. Oxford University Press.
- Nielsen, K. (1984) "Why Should I be Moral? Revisited." *American Philosophical Quarterly* 21(1):81-91.
- Oddie, G. (2008) "Truthlikeness and Value." In Pihlstrom S. (ed.), *Approaching Truth: Essays in Honour of Ilkka Niiniluoto*. College Publications. pp. 225-40.
- Pyke, J. (2017). Inseparable: 10 years joined at the head. *Curious Features in association with Canadian Broadcasting Company*. Documentary.
- Rawls, J. (1971) *A Theory of Justice*. Cambridge, MA: Harvard University Press.
- Raleigh, T. (2019) "The Recent Renaissance of Acquaintance." In (Knowles, J. & Raleigh, T., eds.) *Acquaintance: New Essays*. Oxford: Oxford University Press.
- Roberts, R.C., & Wood, W.J. (2007) *Intellectual Virtues: An Essay in Regulative Epistemology*. Oxford: Oxford University Press.
- Robinson, H. (1994) *Perception*. London: Routledge.
- Russell, B. (1911) "Knowledge by Acquaintance and Knowledge by Description." *Proceedings of the Aristotelian Society*. 11: 108–128.
- (1912) *The Problems of Philosophy*. London, UK: Thornton Butterworth Limited.

- Rutherford, D. (2003) “In Pursuit of Happiness: Hobbes’s New Science of Ethics.”
Philosophical Topics, 31: 369–93.
- Shoemaker, S. (1994) “Introspection and the self.” In Cassam, Q. (ed.), *Self-Knowledge*, pp.
118–139, Oxford: Oxford University Press
- Simmons, A. (2017). “Mind-Body Union and the Limits of Cartesian Metaphysics.”
Philosophers’ Imprint 17.
- Singh, K. (2022). “What’s in an Aim?” *Oxford Studies in Metaethics* 17:138-165.
- Skow, B. (2011) “Experience and the passage of time.” *Philosophical Perspectives* 25:
Metaphysics, pp. 359–387.
- Street, S. (2016) “Constructivism in Ethics and the Problem of Attachment and Loss”
Aristotelian Society Supplementary 90(1):161-189.
- Stump, E. (2010) *Wandering in Darkness: Narrative and the Problem of Suffering*. Oxford
University Press.
- Tye, M. (2000) “Knowing What it Is Like: The Ability Hypothesis and the Knowledge
Argument.” In Preyer, G. (ed.), *Reality and Humean Supervenience: Essays on the
Philosophy of David Lewis*.
- (2009) *Consciousness Revisited: Materialism without Phenomenal Concepts*. MIT
Press.
- Williamson, T. (2000) *Knowledge and its Limits*. Oxford: Oxford University Press.
- Zagzebski, L. (1996) *Virtues of the Mind: An Inquiry into the Nature of Virtue and the
Ethical Foundations of Knowledge*. Cambridge: Cambridge University Press.
- (2013) *Omnisubjectivity: A Defense of a Divine Attribute*. Marquette University
Press.

