Externalism About Knowledge: A Brief Introduction[1]

Luis R.G. Oliveira
University of Houston

1. Internalism and Externalism

In the sense that matters here, someone's *knowledge* that p is or requires a particular kind of connection between their *belief* that p and the *fact* that p (c.f., Armstrong 1973; Zagzebski 1996; Nagel 2014). Yet there are different views on the nature of this connection. Traditional internalism sees the relevant connection as a kind of reflective assurance of truth that is sufficient to put away any skeptical concerns about whether p. Knowledge is here the result of fully satisfying an uncompromising "philosophical curiosity" (Fumerton 2004, 75). Non-traditional internalism – more popular today – compromises on these anti-skeptical ambitions but remains committed to the idea that knowledge requires reflective assurance of some kind. Knowledge is here the result of getting things right by doing well-enough with what is available from the first-person perspective (e.g., one's mental states and/or seemings). Contemporary externalism, by contrast to both of these internalisms, sees the relevant connection as something broader and weaker than reflective assurance of any kind: it is something that can sometimes be instantiated by reflective assurance, but something that can also survive without it. Here knowledge and what is available from the first-person perspective – at any level of ambition – can come apart.

Traditional internalism is by far the historically more prominent view. Arguably, it is the view of knowledge presumed by Plato and Aristotle, by the Pyrrhonian skepticism of Sextus Empiricus, by Modern philosophers from Descartes to Kant, by the two father figures of the central methodological division in contemporary philosophy – Husserl and Russell – and by many of the most important epistemologists of the 20th C.[2] It is easy to see, moreover, the appeal behind their views. Most of us are ready to acknowledge that we could be wrong about almost everything that we believe about ourselves and the world around us. With respect to any proposition p, removing the possibility of error about whether p in a manner that is accessible from the first-person perspective would be an improvement to our epistemic situation. According to traditional internalism, "knowledge" is just a name for an improvement of this kind. To know is to ascend from the vulnerable

[1] For comments on previous drafts, I am grateful to Hilary Kornblith, Timothy Perrine, and Josh DiPaolo.
[2] See Schmitt (1992) for a challenge to this potted history.

uncertainty of mere belief; it is to achieve our "epistemic *summum bonum*" (BonJour 2010, 58). No wonder claims of the form "I know that p but I could be wrong" sound infelicitous to our ears.[3]

Contemporary externalism leaves all of this and more behind. To distinguish it from both traditional and non-traditional internalism, we can characterize it here – abstracting away from the internal variations to be discussed below – as the rejection of two central ideas: that knowledge is incompatible with reflective awareness of the possibility of error, and that knowledge is necessarily tied to the resources that are available from within the first-person perspective. In the next section, I outline five distinctly externalist accounts of knowledge, and two distinctly externalist methodological approaches to knowledge, all fitting this general description.

2. Varieties of Externalism

The contemporary externalist tradition in epistemology gathered momentum in the late 60s and early 70s with work from Goldman (1967), Skyrms (1967), Unger (1968), Dretske (1971), and Armstrong (1973). In response to Gettier's (1963) widely discussed argument that epistemic justification is not *sufficient* for turning a true belief into knowledge – where "epistemic justification" stood for something closer to the non-traditional internalist idea of non-conclusive assurance of truth available from the first-person perspective – all of these authors proposed theories of knowledge where epistemic justification wasn't *necessary* for knowledge either. Not soon after, starting with Goldman (1979), some externalists started to see themselves as offering rival accounts of epistemic justification as well, taking it once again as a necessary condition for knowledge, albeit a condition of a very different kind. As we will see below, both approaches are alive and well in the externalist camp today, as well as approaches that identify further externalist properties such as "warrant" (c.f., Plantinga 1993a) or "entitlement" (c.f., Burge 1993) as necessary for knowledge instead.[4]

2.1. Process Reliabilism

The first self-conscious articulation of what is today called *process reliabilism* comes from Alvin Goldman (1979, 102):

[3] For an overview of traditional internalism, see Coppenger and Bergmann (2016). For an overview of the two non-traditional versions of internalism that are more prominent today – evidentialism and conservatism – see Dougherty (2011) and Tucker (2013). For a criticism of these newer versions from the perspective of traditional internalism, see BonJour (2010).

[4] By "externalist property" I mean a property that can be instantiated in the absence of any kind of reflective assurance of its instantiation.

> If S's belief in p at t results from a reliable cognitive process, and there is no reliable or conditionally reliable process available to S which, had it been used by S in addition to the process actually used, would have resulted in S's not believing p at t, then S's belief in p at t is justified.

This is, of course, a theory of epistemic justification. Traditional process reliabilism is thus an externalist theory of *knowledge* by way of accepting as a necessary condition for knowledge an externalist property of *justification* (c.f., Goldman 1986, 53-55).

There are two central challenges to this kind of view. The first challenge argues that the central notion of a *cognitive process* is impossible to be made acceptably precise in any principled way. "Each token process that causes a particular belief," as Conee and Feldman (1998, 2) have famously put, "is of numerous different types of widely varying reliability." Which of these many types, then, is such that *its* reliability matters for the justification of a belief formed by the token process? This is the so called "generality problem."[5] The second challenge argues that, by breaking the close connection between knowledge and the resources available from the first-person perspective – that is, by making knowledge compatible with the lack of reflective assurance of any kind – process reliabilism leads to intuitively incorrect verdicts about who is justified in believing what and when (and, consequently, about who *knows* what and when). In reference to the details of a famous counterexample (c.f., BonJour 1980), this is sometimes called the "clairvoyance problem."[6] More than anything else, these two challenges have impacted the formulation, understanding, and subsequent defenses of process reliabilism.

Traditional process reliabilism is not the most widely discussed account of knowledge these days.[7] Nonetheless, as we will see below, it is fair to say that all other branches of externalist epistemology are in some sense elaborations on, or developments of, the same original idea: knowing is or requires the instantiation of an externalist property that is or reflects some kind of reliable connection between belief and truth, however we come to understand what that really means.

2.2. Tracking Views

In what is now an undeniable classic, Robert Nozick (1981, 172-177) articulated a theory of knowledge that, surprisingly, is both externalist and centrally interested in responding to skeptical concerns:

---

[5] See Alston (1995), Comesaña (2006), and Lyons (2019) for discussion. Notably, Goldman (1979, 12) himself is already sensitive to this concern.

[6] See Ghijsen (2016) and Moon (2018) for discussion.

[7] I will make similar sociological assessments of other views. These assessments are all unempirical and are not intended as assessments of their accuracy. Truth, after all, is not a popularity contest. See Goldman and Beddor (2021) for an overview of recent developments in the process reliabilism approach.

S knows that p if and only if

(1) P is true.
(2) You believe that p.
(3) If p weren't true, you wouldn't believe that p.[8]
(4) If p were true, you would believe that p.

Differently from traditional process reliabilism, here we have an account of knowledge that altogether rejects the need for justification. To know is to counterfactually "track" the truth of the relevant belief across the neighboring possible worlds. Yet since many skeptical concerns are framed in counterfactual terms, or depend on counterfactual claims, satisfying these conditions for knowledge is tantamount to deflecting these skeptical threats.[9]

There are two central challenges to this kind of view. The first is that, given Nozick's third condition, which has come to be known as "sensitivity" and has become the focus on this line of thought, it turns out that while we *do not* know that we are not in a global skeptical scenario – e.g., being deceived by an evil demon in all of our beliefs – we *do* know that more local skeptical hypotheses are false – e.g., that we don't have hands. This, however, requires denying that knowledge is closed under known entailment: even though I know that I have hands, and know that if I have hands then I'm not being deceived by an evil demon in *all* of my beliefs, I do not know that I am not being deceived by an evil demon in all of my beliefs (c.f., Nozick 1981, 200-211). We can call this the problem of epistemic closure.[10] The second challenge is that tracking views seem incapable of explaining how we can know or fail to know necessary truths. This is because knowing some true proposition in the actual world, according to tracking views, depends in part on what I believe about it in worlds where that proposition happens to be false, and yet there are no worlds where necessary truths are false. This is the problem of counterpossibles.[11]

Tracking views are also not the most widely discussed account of knowledge these days.[12] They are, however, the starting point in a trend that shifts the focus of externalist accounts from features of the *cognitive process* behind the belief to the *modal profile* of the belief itself. This is a shift, as I say, in "focus;" process reliabilism itself already makes a modal claim (see the second condition in the antecedent of Goldman's original view

[8] Much earlier, Dretske (1971) incorporated a similar counterfactual into his own account of knowledge, taking it as a condition for having what he called "conclusive reasons".
[9] Despite the simplicity and "pleasing symmetry" (c.f., Nozick's (1981, 177) of the now famous formulation mentioned just above, Nozick (1981, 179-185) is himself quick to recognize the importance of relativizing his account to the particular "methods" used in forming the relevant belief.
[10] See Warfield (2004) and Baumann (2011) for discussion.
[11] See Sosa (1999), Roland and Cogburn (2011), and Blome-Tillmann (2017) for discussion. Another prominent challenge to sensitivity views is their difficulty in accounting for inductive knowledge. See Vogel (1987) for an early discussion.
[12] See DeRose (2010) and Becker and Black (2012) for overviews and recent developments in this approach.

above), and the emphasis on "methods of belief formation" in both tracking and safety views (see below) are very soft departures from an emphasis on cognitive processes instead.

2.3. Safety Views

As just mentioned, tracking views – based around a sensitivity requirement – are precursors to safety views. Their shift in focus from cognitive processes to modal profiles, however, has an explanatory rationale: both attempt to use the modal profile of our beliefs in the articulation of the incompatibility of knowledge with a certain kind of *luck*. Together, these kinds of views have come to be known as *anti-luck epistemology* (c.f., Pritchard 2005).

Perhaps the first to suggest an alternative modal requirement to sensitivity was Steven Luper-Foy (1984, 46-7), suggesting that:

> S knows that p if and only if (i) S's belief that p is caused by a sequence of events F, and (ii) if F were to occur, then p would be true.

Variations on this idea are now known as "safety" and are now typically offered as just a necessary condition on knowledge. On Pritchard's (2005, 156) early view, for example:

> S's belief that p is safe if and only if the actual and all relevantly nearby possible worlds in which S believes that p based on the same method as in the actual world are p-worlds.

And on one of Williamson's (2000, 147) many formulations:

> If one knows, one could not easily have been wrong in a similar case.

The main appeal of this alternative modal approach, however we formulate it, is the fact that it avoids both of the central worries associated with tracking views: they can block both local *and* global skeptical threats – thus avoiding the need to deny closure – and they can account for our knowledge of necessary truths (c.f., Sosa 1999).[13]

Safety views are some of the most widely discussed versions of externalism today.[14] But Safety's peace with closure is also the source of its main difficulty – at least if we put aside the pressure from a litany of counterexamples that is common to all attempts to define knowledge. Differently from the skeptical examples discussed above, there are ordinary instances, after all, where knowledge indeed *does not* seem closed under

[13] Despite being an early proponent of the Safety View, Sosa has now distanced himself from it (c.f., Sosa 2009: 206–7).
[14] See Grundman (2018) for an overview of recent developments in safety views.

known entailment. It seems to most of us, for example, that someone can know – and *does* know, if they are like me – that they will be unable to completely pay off their mortgage loan in the next month. Indeed, we have no trouble ascribing to ourselves and others various bits of ordinary knowledge of this future-oriented kind. Just as much, it seems to most of us that we do not know things such as that I will not win the lottery this month, or that I will not receive an impressive inheritance from an unknown deceased relative this month, and so on. And yet these propositions are entailed by the propositions we seem to know. The history of Safety views, in fact, is in large part the history of how to successfully calibrate the condition in order to make sense of these (and other) patterns of knowledge ascriptions in a principled way.[15]

2.4. Virtue Epistemology

In the most important paper of his early work, Ernest Sosa (1980, 23) concluded his examination of the dialectic between foundationalism and coherentism by suggesting that, on an epistemology modeled on an ethics of virtue:

> Primary justification would apply to intellectual virtues, to stable dispositions for belief acquisition, through their greater contribution toward getting us to the truth. Secondary justification would then attach to particular beliefs in virtue of their source in intellectual virtues or other such justified dispositions.

This is a version of virtue epistemology where knowledge requires, once again, an externalist property of justification. Not all virtue epistemologies, however, have retained this feature, and indeed Sosa's own development of this approach has shifted its focus away from justification and towards a direct account of knowledge itself. What has remained constant in the entire camp, however, is Sosa's original guiding suggestion that knowledge is, in some sense, *intellectually virtuous true belief*.

Sosa's virtue epistemology associates intellectual virtue with cognitive dispositions that are conducive to true beliefs. As he puts it a bit later, "an intellectual virtue is a quality bound to help maximize one's surplus of truth over error" (Sosa 1985, 227). What Sosa has in mind here are qualities of someone's cognitive faculties such as vision, memory, introspection, and reasoning. To know, according to this approach, is to have a true belief formed by a faculty with the relevant features.[16] A bit differently, Linda Zagzebski (1996, 270) has developed a virtue epistemology around a notion of intellectual virtue that focuses instead on someone's *intellectual character*, more broadly construed:

---

[15] See Pritchard (2008) and Hawthorne and Aarnio (2009) for discussion.

[16] Notice how this kind of virtue epistemology – often called "virtue reliabilism" – harks back to talk of "cognitive processes" introduced by process reliabilism. See Greco (1999) and Sosa (2009) for discussion. A central feature of Sosa's later work, however, is the distinction between different kinds of knowledge and different levels of knowing. More importantly, Sosa's virtue epistemology is now centered on a detailed analysis of "competent performances." See Sosa (2015) and Sosa (2020).

> An act of intellectual virtue A is an act that arises from the motivational component of A, is something a person with virtue A would (probably) do in the circumstances, is successful in achieving the end of the A motivation, and is such that the agent acquires a true belief (cognitive contact with reality) through these features of the act.

To know, according to this approach, is to have contact with reality through the exercise of responsible intellectual agency, where this is partly a matter of one's motivational structure.[17]

Virtue epistemology is yet another family of widely discussed views today.[18] In fact, we have now arrived at a point of recent multiple convergence. Not only have some virtue epistemologists started to elide the internal differences in this approach (c.f., Greco 2019 and Sosa 2020), even some of the leading and early proponents of the safety view now think of their views as hybrids of the safety and virtue approaches (c.f., Pritchard 2012).[19] This is all the more interesting given the connections already noted between all of these views and process reliabilism.

2.5. Proper Function

Following the general post-Gettier externalist dissatisfaction with available theories of justification and knowledge, and a bit differently from the preceding views, Ruth Milikan (1984) was the first to offer an account of knowledge based on the teleological notion of a *proper function*. As she put it (Milikan 1984, 316):

> The proper functions of any body organ or system are those functions which helped account during evolutionary history for survival or proliferation of the species containing the organ or system.

For example, the proper function of a heart is to pump blood in certain specific ways – the "normal" ways, meaning the ways it has historically performed its function – such that a heart not doing so can rightly be said to not be functioning properly. Milikan treats knowledge in the same way. From an evolutionary perspective, at least one proper function of our belief-forming processes is the production of true beliefs.[20] The term "knowledge," then, simply picks out the complex phenomena of having a true belief that is the result of a properly functioning belief-forming mechanism (c.f., Milikan 1984, 325). There are important differences, of course, but the similarities to process reliabilism should be plain (c.f., Graham 2017).

---

[17] Developments of this approach are often called "virtue responsibilism." See Baehr (2006) and Axtell (2011) for discussion.
[18] See Turri, Alfano, and Greco (2019) for discussion of recent developments.
[19] See Greco (2020) for the suggestion that the changes in Sosa's views (from a focus on Safety to Virtues) are not substantive breaks in his thinking about knowledge.
[20] Although see Burge (2010) and Graham (2014) for discussion.

Milikan's proper functionalism offers a direct and naturalized analysis of knowledge: there is no reference to justification or any similar property in it, and the aim of, and rationale for, the theory is describing a natural phenomena in the same ways we describe other natural phenomena (e.g., by understanding its evolutionary origins).[21] There is variation within proper functionalist views on each of these aspects of Milikan's work. Most influentially, Alvin Plantinga's (1993b) proper functionalist view abandons both, offering a non-naturalist account of what he calls *warrant*: "that, whatever precisely it is, which together with truth makes the difference between knowledge and mere true belief" (Plantinga 1993a, 3). According to Plantinga (1993b, 194):

> A belief *B* has warrant for you if and only if (1) the cognitive faculties involved in the production of *B* are functioning properly… ; (2) your cognitive environment is sufficiently similar to the one for which your cognitive faculties are designed; (3) the triple of the design plan governing the production of the belief in question involves, as purpose or function, the production of true beliefs… ; and (4) the design plan is a good one: that is, there is a high statistical or objective probability that a belief produced in accordance with the relevant segment of the design plan in that sort of environment is true.

By highlighting the importance of a "design plan," and by arguing that Milikan's and others' naturalistic understanding of proper function is defective, Plantinga produces a proper functionalist account of knowledge that is essentially super-naturalist: it makes the possibility of knowledge depend on the intelligent and well-intentioned designs of a super-natural being.

Neither Milikan nor Plantinga take justification to be a requirement for knowledge. Some proper functionalists, however, disagree. According to Michael Bergmann's (2004, 44) earliest account, for example,

> In all possible circumstances, a belief B is prima facie justified iff B is produced by cognitive faculties that are (a) functioning properly (operating normally), (b) truth-aimed (their function is to produce true beliefs), and (c) are reliable in the environments for which they were "designed" (i.e., reliable in normal conditions).

On this kind of view, justification is an externalist property and once again a necessary condition for knowledge. With or without justification, at any rate, proper functionalist views are not widely discussed in the externalist camp today.[22] That said, some versions of this approach fall under the more general category of naturalized epistemology – one of the two overarching and methodological approaches to knowledge that are distinctly externalist, both to be discussed in turn.

2.6. Naturalized Epistemology

[21] For developments of this approach, see Graham (2011).
[22] See, however, Boyce and Moon (2016) for a recent defense.

Ruth Milikan's appeal to evolutionary science in a bid for understanding knowledge was not the first time a major epistemologist pushed for a close relationship between epistemological theorizing and the empirical sciences.[23] In what is widely considered the contemporary opening salvo of naturalized epistemology, Quine (1969, 75-8) suggests that instead of pursuing "creative reconstructions" of the proper relation between theory and experience (by which he means *a priori*, empirically ungrounded hypotheses), it would be more sensible for epistemologists to simply "settle for psychology." On one reading, Quine is here advocating for a kind of "replacement naturalism" (c.f., Kornblith 1994a, 3-4) that is not representative of naturalized epistemology today: namely, the complete abandonment of the traditional subject matter of epistemology.[24] On another reading, however, Quine is here advocating for what Goldman (1999, 26) calls "moderate naturalism" instead. As Kornblith (1994b, 49) puts it, in the same spirit:

> What does have priority over both metaphysics and epistemology, from the naturalistic perspective, is successful scientific theory, and not because there is some a priori reason to trust science over philosophy, but rather because there is a body of scientific theory which has proven its value in prediction, explanation, and technological application. This gives scientific work a kind of grounding which no philosophical theory has thus far enjoyed. Only by making philosophy continuous with the sciences, as Quine has suggested, may we provide it with a proper foundation.

Yet despite a shared commitment to the theoretical relevance and authority of the empirical sciences, there is disagreement on what such a commitment actually entails: disagreement, that is, about the substantive constraints imposed by science on the theoretical entities, relations, views, methods, and so on, that are available for theory building. In fact, whatever we make of Quine, naturalized epistemology is today a large and varied group.

Two prominent expressions of naturalized epistemology deserve special note. For some, naturalizing epistemology entails a commitment to what we can call a *realistic psychology constraint* on evaluative epistemology. In contrast to epistemological theories that "overintellectualize" human knowledge by grounding their prescriptive force on facts about logical or probabilistic relations, naturalized epistemologists here hold that how we *actually* form our beliefs partly determines how we *should* form them in the first place (c.f., Goldman 1976, 102).[25] For others, however, naturalizing epistemology entails something more: a departure from the goal of conceptual analysis towards the study of "a certain natural phenomenon, namely, knowledge itself" (Kornblith 1999, 161). With science being the best tool available for the study of natural phenomena in general, naturalized epistemologists here look to its results for clues to better understand the

[23] See Kitcher (1992) for a discussion of the history and rationale for naturalism.
[24] See Kim (1988) for that reading.
[25] See Dretske (1981), Harman (1986), and Kornblith (2001) for discussion.

epistemic properties of their interest. On this way of seeing things, few proponents of the approaches discussed above count as naturalized epistemologists.[26]

It is perhaps the latter that is most commonly understood as a distinctive externalist approach to knowledge. This is because we here have a call to revise both the methods and the projects that have characterized epistemology throughout its history.[27] Given what we have already learned from empirical psychology and from (negative) experimental philosophy, the line goes, it turns out that what our intuitions tell us about the proper application of our epistemic concepts is a rather poor guide to understanding the real and natural phenomena behind them.[28] For similar reasons, the line continues, epistemological projects centered around the truth-conducive value of reflection and deliberation are bound to lead us astray.[29]

2.7. Knowledge First

Some of naturalized epistemology aside, much work in externalist epistemology consists in the attempt to "analyze" the concept of knowledge into its component concepts and relations: e.g., justification, sensitivity, safety, and proper-function. Yet the intractability of debates about the correct analysis of knowledge, and especially the recurring and unrelenting threat of Gettier-styled challenges (c.f., Zagzebski 1994), have led many to believe that this analytical project is doomed. As an alternative, Timothy Williamson (2000) has suggested not only that we treat knowledge as a fundamental epistemic concept that is not amenable to analysis into more fundamental parts, but also that we reverse the traditional order of explanation: that we take knowledge as our starting point and use it to analyze notions typically used to analyze it instead. Knowledge, on this view, is an "unexplained explainer" (Williamson 2000, 10). This is the second distinctly externalist overarching and methodological approach to knowledge.

This is not to say, of course, that, on the knowledge first approach, nothing can be said that illuminates the concept or the phenomena of knowledge. Rather, it is to shift our understanding of how it is that what we say can sometimes shed light. On Williamsons' view, the knowledge first epistemologist, when focusing on the phenomena of knowledge itself (in contrast to focusing on the *concept* of knowledge, c.f. Williamson 2007), is in the business of exploring *philosophical models*. In short, this is the business of searching for truths of the kind "if a given case satisfies the model description, then it satisfies this other description too" (c.f., Williamson forthcoming, 4). Here the goal is not to define the concept of knowledge in a way that escapes

[26] See Bishop and Trout (2005), Lyons (2009), and Goldman (2015) for discussion.
[27] Again, however, see Schmitt (2014) for disagreement on this characterization of epistemology's past.
[28] See Weinberg, Gonnerman, Buckner, and Alexander (2010) and Nagel (2012) for discussion.
[29] See Kornblith (2012) and Hannon (2018) for discussion.

counterexamples – since models are self-conscious simplifications of some target phenomena – but rather to provide an approximate understanding of the phenomena of knowledge that is fruitful for our understanding of neighboring phenomena.

As an analogy, consider two approaches to understanding causation. One might try to analyze the concept of causation by reducing it to something else, for example by showing how causation is nothing but counterfactual dependence (c.f., Lewis 1973). Alternatively, one might try to show how various other phenomena – representation, perception, knowledge, intention, action, and so on – can be better understood in (unanalyzed) causal terms. That too would allow one to better understand causation itself, in turn, in an important way. Williamson sees this second approach as most likely to be the best way to illuminate what knowledge is, specially once we "give up Cartesian fantasies about the mind" (c.f. Williamson 2014, 4).

With this approach in mind, Williamson (2000, 147) himself has endorsed a version of the safety view, as mentioned above.[30] Using the same method, he has also claimed that someone's evidence is all and only what they know (c.f., Williamson 2000, 193),[31] and that knowledge is the norm of both assertion and action (c.f., Williamson 2000, 243; 2005, 231).[32] Indeed, Williamson's knowledge first approach has been put to a wide variety of applications, including the analysis of concepts previously used for the analysis of knowledge itself: *justification* and *belief*.[33] Whether those applications have helped us better understand traditional issues in epistemology, however, is still very much a matter of debate.[34]

3. Chapter Summaries

This collection brings original contributions on each of the externalist accounts and approaches to knowledge briefly outlined above. Here is a summary of the chapters.

In chapter 1, Jack Lyons argues that the best solution to the generality problem reveals important lessons for nearly any externalist theory of knowledge. While the generality problem is typically pressed against the appeal to *processes*, other externalist theories of knowledge face the same challenge when they appeal to *ways*, *methods*, *bases*, and *competences*, all of which refer to the token causal pathway to belief formation whose relevant type satisfies some necessary condition for knowing – sensitivity, adherence, safety, proper function, etc. By

[30] Williamson (2009, 307) is also sensitive to the way in which process reliabilism seems to be the backbone of all kinds of externalist accounts: "although my view is not process reliabilism in Goldman's sense, it is not quite as distant from process reliabilism as he thinks."
[31] See Dodd (2007) and Littlejohn (2017) for discussion.
[32] See Hawthorne (2004) and Hawthorne and Stanley (2008) for discussion.
[33] See Littlejohn (2013) and Simion (2019) for discussion.
[34] See Goldman (2009), McGlynn (2014), and Schechter (2017) for criticisms.

extending the "algorithm and parameters" theory of process individuation he has defended elsewhere, Lyons goes on to show how these other externalist theories of knowledge can make use of the same resources available to process reliabilism.

In chapter 2, Sanford Goldberg argues for a new answer to the clairvoyance problem – which he himself calls *the problem of idiosyncratic reliability*. After criticizing two of the leading attempts to reply to this problem – one that introduces a "no-defeaters condition" on the power of a reliable process to deliver *prima-facie* justification and one that restricts a reliable process's capacity to produce *prima-facie* justification in the first place – Goldberg goes on to defend a solution that is based on the human capacity for surprise. According to Goldberg, what the problem of idiosyncratic reliability reveals is that the presumption of reliability we must bestow on certain basic belief-forming processes depends on them satisfying what he calls an *expectedness* condition: reliance on them must be something that is to be expected by members of our species.

In chapter 3, Kelly Becker continues the focus on methods of belief formation by turning his attention to the alleged tension between the epistemological externalism of tracking views and the internal manner by which these views individuate the relevant basis of a particular belief. This challenge has been recently pressed against tracking views by knowledge-first epistemologists and disjunctivists as an unmotivated left-over from internalism. Becker here mounts a defense on behalf of tracking views, dispelling the tension by providing an externalist friendly motivation for the internal individuation of basis.

In chapter 4, Sherrilyn Roush focuses on the question of whether naked statistical evidence alone – the frequency of some property in some population – is adequate evidence for trial verdicts. Judges and jurors, she notes, are reluctant to consider such evidence as adequate, but process reliabilism struggles to explain why this could be a reasonable position. By contrasting the role of probabilities in process reliabilism and in her distinctive version of a tracking view, Roush goes on to argue that while process reliabilism lacks the resources to explain why naked statistical evidence could be inadequate, her tracking view is indeed capable of delivering that correct result.

In chapter 5, Duncan Pritchard defends what he calls *anti-risk virtue epistemology*, a moderately externalist theory of knowledge that incorporates insights from safety approaches, virtue approaches, and even internalist approaches. According to Pritchard's view, while knowing requires only an externalist property of *cognitive responsibility*, developed agents often satisfy robust internalist constrains as well in paradigmatic instances of their knowledge (e.g., perceptual knowledge). One of the advantages of this view, Pritchard tells us, is that it restores for the externalist the important connection between knowledge and rationality.

In chapter 6, John Hawthorne and Christina Dietz take a closer look at many of the finer choice points that safety theorists face – how to think of *methods*, *near-beliefs*, *epistemic counterparts*, *closeness of worlds*, the scope of the *nearly-all* condition on neighboring worlds, and more. Each of these choice points leads to different formulations of the safety condition on knowledge, leading to different strengths and weaknesses for the approach. They conclude by discussing a final choice-point: whether to think of safety approaches as providing necessary and sufficient conditions on knowledge – as Pritchard has done in his chapter – or as an exercise in what Williamson has recently called "model building" in epistemology.

In chapter 7, Ernest Sosa delivers the latest statement of his ever more comprehensive framework for what he calls a *telic virtue epistemology*. In line with his most recent work, Sosa here expands on his previous AAA structure for epistemic normativity – Accuracy, Adroitness, Aptness – with the complimentary addition of the SSS dimension of evaluation – Situation, Shape, Skill. With this updated framework in hand, Sosa proceeds to examine how his view can make sense of the normativity of "pure thinking," thinking that is not sorted or bounded by the practices of any community, and the role in such thinking of "default assumptions" – assumptions to the effect that background conditions for performance in a given domain are in place.

In chapter 8, Berit Brogaard addresses herself to the aforementioned divide within virtue epistemology – that between taking intellectual character virtues and responsible agency to be necessary for knowledge (*virtue responsibilism*) and taking the quality of one's cognitive faculties but not character virtues or responsible agency as necessary instead (*virtue reliabilism*). According to Brogaard, developments in the virtue epistemology approach to knowledge reveal that the divide is ultimately artificial. Virtue reliabilists, Broogard argues, are ultimately required to accept that intellectual character virtues involve a substantial truth-motivational component, and that knowledge requires a kind of epistemic responsibility that is far more substantive than a causal, naturalistic notion of attribution.

In chapter 9, Peter Graham begins by providing a unifying overview of proper functionalism, the approach to knowledge focused on the function and normal functioning of belief-forming processes. After distinguishing between five grades of the externalist property of *warrant* – the property enough of which turns a true belief into knowledge – Graham sets aside supernaturalist accounts of the central notion of a function and expands on two versions of the naturalist variant instead: the "selected-effects account of functions" and the "organizational theory of functions." After explaining both and endorsing the latter, Graham goes on to argue that warrant can obtain in non-normal conditions (such as demon-worlds), and argue for a solution to the Swampman objection: the suggestion that proper functionalism cannot explain how an instantaneously created mind could have warranted beliefs and knowledge.

In chapter 10, Kenneth Boyce and Andrew Moon defend an *explanationist* version of proper functionalism. After explaining proper functionalism's initial appeal, they note two major objections to proper functionalism in general: creatures with no design plan who appear to have knowledge (Swampman) and creatures with malfunctions that increase reliability. The second of these objections, in fact, leads them to explore a variety of cases of what they call *warrant-compatible malfunction*. With these cases in hand, Boyce and Moon mount a defense of a novel version of proper-functionalism – *explanationist proper functionalism* – and argue that this version can block the two major objections faced by proper-fuctionalist accounts in general.

In chapter 11, Patrick Rysiew explores the relationship between *naturalistic epistemology* (NE) and *externalism*. After providing an overview of the main positions and commitments associated with NE, Rysiew argues that while NE does not mandate externalism, the observed connection between them is not an accident. As general approaches, Rysiew argues, both NE and externalism are animated by a desire to have realistic and genuinely explanatory epistemological theories, something that more traditional ways of thinking about knowledge are regarded as failing to achieve. Rysiew then goes on to note and explore some of the significant disagreements among those within the naturalistic-externalist camp, in particular about the appropriate target of philosophical theorizing and the methods that epistemologists should employ in investigating it.

In chapter 12, Louise Antony provides a solution to an externalist version of the Sellarsian problem of "the given." She begins by explaining how this traditional *internalist* problem – how can perceptual experience both provide faithful information about the external world *and* justification for empirical belief? – arises for externalists about knowledge as well. Next, she argues that the problem can be solved within naturalistic bounds, by appealing to a category of causal relations she calls *intelligible causation*.

In chapter 13, Clayton Littlejohn provides a novel knowledge first defense of externalism about both knowledge and justification. Littlejohn's starting point is an understanding justification in terms of norms about what we should and shouldn't believe. The next step is a defense of normative externalism in general – both moral and epistemic – as the view where what we should and shouldn't *do* or *believe* is something determined by whether some external condition obtains or fails to obtain. The final step, then, is a defense of the view that this external condition, in the epistemic case, is knowledge itself: what we should or shouldn't believe (i.e., what we are justified in believing) depends upon what we are in a position to know.

In chapter 14, Timothy Williamson argues against internalism about justification by drawing out some of its unacceptable moral implications. He begins by noting that internalist accounts of epistemic justification typically imply that the bigoted beliefs of consistent bigots are epistemically justified. Given plausible connections between norms of belief and norms of action, however, internalism about the former extends to

internalism about the latter. Williamson then draws the conclusion that, by internalist standards, consistent bigots are justified in acting on their bigoted beliefs, and urges this as a *reductio*.

Conclusion

Earlier in this introductory essay I characterized the unity behind externalist approaches to knowledge as the rejection of two central ideas: that knowledge is incompatible with reflective awareness of the possibility of error, and that knowledge is necessarily tied to the resources that are available from within the first-person perspective. Despite this common and distinguishing core – in contrast to internalist approaches to knowledge – it should be plain by now that externalism is nonetheless a variegated family. There are externalists engaged in conceptual analysis, empirical investigations, and philosophical modeling. There are externalists who focus on evidence and justification and externalists who ignore one or both epistemic notions. There are reductionists and non-reductionists; skeptic-fighters and skeptic-ignorers; naturalists and supernaturalists. And while each approach has its own microcosmos of fundamental notions and conditions, it is clear that they borrow and relate to each other in a variety fruitful ways. What is also clear is that externalism is thriving in contemporary epistemology, despite representing "a very radical departure" from "the general standpoint of the western epistemological tradition" (c.f., BonJour 1980, 56). This timely collection brings together the latest developments in the work of some of those leading this radical departure today.

References

Alston, William (1995). "How to Think about Reliability," *Philosophical Topics* 23 (1): 1-29

Armstrong, David (1973). *Belief, Truth, and Knowledge*. Cambridge University Press.

Axtell, Guy. (2011). "From Internalist Evidentialism to Virtue Responsibilism," In Dougherty, T. (ed.), *Evidentialism and its Discontents*. Oxford: Oxford University Press: 71-87.

Baehr, Jason (2006). "Character, Reliability and Virtue Epistemology," *The Philosophical Quarterly*, 56(223): 193-212.

Baumann, Peter (2011). "Epistemic closure," In Duncan Pritchard & Sven Bernecker (eds.), *The Routledge Companion to Epistemology*. Routledge. pp. 597--608.

Becker, Kelly and Black, Tim (2012) *The sensitivity principle in epistemology* Cambridge: Cambridge University Press.

Bergmann, Michael (2004). "Externalist Justification without Reliability." *Philosophical Issues* 14: 35-60.

Bishop, Michael A. and J.D. Trout, (2005). "The Pathologies of Standard Analytic Epistemology", *Nous*, 39(4): 696–714.

Blome-Tillmann, Michael (2017). "Sensitivity Actually," *Philosophy and Phenomenological Research* 94 (3):606-625.

BonJour, Laurence (1980). "Externalist Theories of Empirical Knowledge", *Midwest Studies in Philosophy* 5(1): 53-74.

BonJour, Laurence (2010). "The Myth of Knowledge," *Philosophical Perspectives* 24 (1): 57-83.

Boyce, Kenneth and Moon, Andrew. (2006) "In Defense of Proper Functionalism: Cognitive Science Takes on Swampman." *Synthese* 193, 2987–3001

Burge, Tyler. (1993). "Content Preservation," *The Philosophical Review*, Vol. 102, No. 4. (Oct., 1993), pp. 457-488.

Burge, Tyler (2010). *Origins of objectivity*. New York: Oxford University Press.

Comesaña, Juan. (2006). A Well-Founded Solution to the Generality Problem. *Philosophical Studies* 129 (1):27-47.

Conee, Earl. & Feldman, Richard. (1998). The generality problem for reliabilism. *Philosophical Studies* 89 (1):1-29.

Coppenger, Brett and Bergmann, Michael (2016). *Intellectual Assurance: Essays on Traditional Epistemic Internalism.* Oxford University Press.

DeRose, Keith (2010). "Insensitivity is Back, Baby!" *Philosophical Perspectives* 24 (1): 161-187.

Dodd, Dylan. (2007) "Why Williamson Should Be a Sceptic." *Philosophical Quarterly* 57: 635–49.

Dougherty, Trent (2011). *Evidentialism and its Discontents*. Oxford University Press.

Dretske, Fred (1970) "Epistemic Operators", *The Journal of Philosophy* 67(24): 1007-1023.

Dretske, Fred (1981). *Knowledge and the Flow of Information*, Cambridge, MA: MIT Press.

Fumerton, Richard (2004). "Achieving Epistemic Ascent," in John Greco (ed.), *Sosa and His Critics*. Malden: Blackwell, 72 - 85.

Gettier, Edmund. (1963), 'Is Justified True Belief Knowledge?,' *Analysis* 23, 121-123.

Ghijsen, Harmen. (2016) "Norman and Truetemp Revisited Reliabilistically: a Proper Functionalist Account of Clairvoyance," *Episteme* 13: 89–110.

Goldman, Alvin (1967) "A Causal Theory of Knowing", *The Journal of Philosophy* 64: 357–372.

Goldman, Alvin. (1976), 'Discrimination and Perceptual Knowledge,' *Journal of Philosophy* 73, 771-791.

Goldman, Alvin. (1979) "What is Justified Belief?", in George Pappas (ed.), *Justification and Knowledge*, Dordrecht: Reidel, pp. 1-23.

Goldman, Alvin. (1986) *Epistemology and Cognition*, Cambridge, MA: Harvard University Press.

Goldman, Alvin. (1999) "*A Priori* Warrant and Naturalistic Epistemology", *Philosophical Perspectives* 13: 1-28.

Goldman, Alvin (2009). "Williamson on Knowledge and Evidence," in Patrick Greenough, Duncan Pritchard & Timothy Williamson (eds.), *Williamson on Knowledge*. Oxford University Press. pp. 73-91.

Goldman, Alvin (2015). "Review of Hilary Kornblith, A Naturalistic Epistemology: Selected Papers", *Notre Dame Philosophical Reviews*.

Goldman, Alvin and Bob Beddor, (2021). "Reliabilist Epistemology", The Stanford Encyclopedia of Philosophy (Summer 2021 Edition), Edward N. Zalta (ed.)

Graham, Peter J. (2011). "Intelligent Design and Selective History: Two Sources of Purpose and Plan." In J. Kvanvig, ed., *Oxford Studies in Philosophy of Religion*, vol. 3: 67-88.

Graham, Peter J. (2012). "Epistemic Entitlement." *Nous* 46: 449-482.

Graham, Peter J. (2014). "The Function of Perception," In Abrol Fairweather (ed.), *Virtue Scientia: Bridges between Virtue Epistemology and Philosophy of Science*. Synthese Library. pp. 13-31

Graham, Peter J. (2017). "Normal Circumstances Reliabilism." *Philosophical Topics* 45: 33-61.

Greco, John. (1999). Agent Reliabilism. *Philosophical Perspectives*, 13: 273-296.

Greco, John. (2019). Virtue, Knowledge, and Achievement. In Battaly, H. (ed.), *Routledge Handbook of Virtue Epistemology*, New York: Routledge: 273-284.

Greco, John (2020). "Safety in Sosa," *Synthese* 197: 5147–5157.

Grundman, Thomas (2018). "Saving Safety from Counterexamples," *Synthese* 197(12): 5161-5185.

Hannon, Michael (2018). "Intuitions, Reflective Judgments, and Experimental Philosophy," *Synthese* 195 (9):4147-4168.

Harman, Gilbert (1986) *Change in View*, Cambridge, MA: MIT Press.

Hawthorne, John. (2004) *Knowledge and Lotteries*. Oxford: Oxford University Press.

Hawthorne, J. and Lasonen-Aarnio, Maria. (2009), 'Knowledge and Chance,' in Greenough and Pritchard (eds), *Williamson on Knowledge*, 92-108.

Hawthorne , John and Stanley, Jason. (2008) "Knowledge and Action," *Journal of Philosophy* 105: 571–590.

Kim, Jaegwon (1988). "What is 'Naturalized Epistemology'?", *Philosophical Perspectives*, Vol. 2 (1988): 381–405.

Kitcher, Phillip (1992). "The Naturalists Return", *Philosophical Review*, 101(1): 53–114.

Kornblith, Hilary (1994a) "Introduction: What is Naturalistic Epistemology?", in Kornblith (ed.), *Naturalizing Epistemology*, 2nd ed., Cambridge, MA: MIT Press, pp. 1-14.

Kornblith, Hilary (1994b) "Naturalism: Both Metaphysical and Epistemological", *Midwest Studies in Philosophy* XIX: 39-52.

Kornblith, Hilary (1999) "In Defense of a Naturalized Epistemology", in John Greco and Ernest Sosa (eds.), *The Blackwell Guide to Epistemology*, Malden, MA: Blackwell, pp. 158-169.

Kornblith, Hilary (2001) "Epistemic Obligation and the Possibility of Internalism", in Abrol Fairweather and Linda Zagzebski (eds.), *Virtue Epistemology: Essays on Epistemic Virtue and Responsibility*, Oxford: Oxford University Press, pp. 231-248.

Kornblith, Hilary (2012), *On Reflection*. Oxford University Press.

Littlejohn, Clayton. (2013) "The Russellian Retreat," *Proceedings of the Aristotelian Society* 113: 293–320.

Littlejohn, Clayton (2017). "How and Why Knowledge is First." In A. Carter, E. Gordon & B. Jarvis (eds.), *Knowledge First*. Oxford: Oxford University Press.: 19-45.

Luper-Foy, Steven (1984). "The Epistemic Predicament: Knowledge, Nozickian Tracking, and Scepticism," *Australasian Journal of Philosophy* 62 (1): 26–49.

Lyons, Jack (2009). *Perception and Basic Beliefs: Zombies, Modules and the Problem of the External World*. Oxford University Press.

Lyons, Jack (2019). Algorithm and parameters: Solving the generality problem for reliabilism. *Philosophical Review* 128 (4):463-509.

McGlynn, Aidan (2014). *Knowledge First?*, Palgrave Macmillan.

Millikan, Ruth (1984) "Naturalist Reflections on Knowledge", *Pacific Philosophical Quarterly* 65(4): 315–334.

Moon, Andrew (2018). "How to use cognitive faculties you never knew you had." *Pacific Philosophical Quarterly* 99(1): 251–275.

Nagel, Jennifer (2012). "Intuitions and Experiments: A Defense of the Case Method in Epistemology", *Philosophy and Phenomenological Research*, 85(3): 495–527.

Nagel, Jennifer (2014). *Knowledge: A Very Short Introduction*. Oxford University Press.

Nozick, Robert (1981) *Philosophical Explanations*, Cambridge, MA: The Belnap Press of Harvard University Press.

Plantinga, Alvin (1993a) *Warrant: The Current Debate*. Oxford University Press.

Plantinga, Alvin (1993b) *Warrant and Proper Function*, Oxford University Press.

Pritchard, Duncan (2005) *Epistemic Luck*, Oxford: Oxford University Press.

Pritchard, Duncan (2008). *"Knowledge, Luck, and Lotteries." In: Hendricks, V. & Pritchard, D. (eds.)*. New Waves in Epistemology. London: Palgrave Macmillan: 28-51.

Pritchard, Duncan (2012). "Anti-Luck Virtue Epistemology," *The Journal of Philosophy* 109 (3): 247-279.

Quine, W.V.O. (1969) "Epistemology Naturalized", in *Ontological Relativity and Other Essays*. Columbia University Press: 69-90.

Roush, Sherrilyn (2005). *Tracking Truth: Knowledge, Evidence, and Science*. Oxford University Press.

Roland, Jeffrey & Cogburn, Jon (2011). "Anti-Luck Epistemologies and Necessary Truths," *Philosophia* 39 (3):547-561

Schechter, Joshua (2017). "No Need for Excuses: Against Knowledge-First Epistemology and the Knowledge Norm of Assertion," in *Knowledge-First: Approaches in Epistemology and Mind*, by J. Adam Carter, Emma Gordon, and Benjamin Jarvis (eds.) Oxford University Press: 132–159.

Schmitt, Frederick (1992). *Knowledge and Belief*. Routledge.

Schmitt, Frederick (2014). *Hume's Epistemology in the Treatise: A Veritistic Interpretation*. Oxford University Press.

Simion, Mona (2019), "Knowledge-First Functionalism," *Philosophical Issues* 29 (1):254-267.

Skyrms, Brian (1967). "The Explication of "X knows that p"" Journal of Philosophy 64 (12):373-389

Sosa, Ernest (1980). "The Raft and the Pyramid: Coherence versus Foundations in the Theory of Knowledge," *Midwest Studies in Philosophy* 5 (1): 3-26.

Sosa, Ernest (1985). "Knowledge and Intellectual Virtue," *The Monist* 68 (2):226-245.

Sosa, Ernest (1999). "How to defeat opposition to Moore," *Philosophical Perspectives* 13:137-49.

Sosa, Ernest (2009) *Reflective Knowledge: Apt Belief and Reflective Knowledge, Volume II*, Oxford: Oxford University Press.

Sosa, E. (2015). *Judgment and Agency*. Oxford: Oxford University Press.

Sosa, E. (2020). *Epistemic Explanations: A Theory of Telic Normativity, and What it Explains*. Oxford University Press.

Tucker, Chris (2013). *Seemings and Justification: New Essays on Dogmatism and Phenomenal Conservatism*. Oxford University Press.

Turri, John., Alfano, Mark. & Greco, John. (2019). Virtue Epistemology. In Zalta, E. N. (ed.) *The Stanford Encyclopedia of Philosophy*

Unger, Peter (1968). "An analysis of factual knowledge," *Journal of Philosophy* 65 (6):157-170.

Vogel, Jonathan (1987). "Tracking, Closure, and Inductive Knowledge," In Luper-Foy Steven (ed.), *The Possibility of Knowledge: Nozick and His Critics*. Rowman & Littlefield. pp. 197--215

Warfield, Ted (2004). "When Epistemic Closure Does and Does Not Fail: A Lesson From the History of Epistemology," *Analysis* 64 (1):35–41.

Weinberg, Jonathan, Gonnerman, Chad, Buckner, Cameron & Alexander, Joshua (2010). "Are Philosophers Expert Intuiters?," *Philosophical Psychology* 23 (3):331-355.

Williamson, Timothy. (2000), *Knowledge and Its Limits*, Oxford University Press.

Williamson, Timothy. (2005) "Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge," *Philosophical Quarterly* 219: 213–35.

Williamson, Timothy (2007). *The Philosophy of Philosophy*. Wiley Blackwell.

Williamson, Timothy (2009). "Reply to Alvin Goldman" in P. Greenough and D. Pritchard (eds.) *Williamson on Knowledge*. Oxford University Press: 305-312.

Williamson, Timothy (2014), "Knowledge First," in M. Steup, J. Turri, E. Sosa (eds.) *Contemporary Debates in Epistemology*. Wiley and Sons: 1-9

Williamson, Timothy (forthcoming). 'Model Building in Philosophy,' Blackford and Broderick (eds.), *Philosophy's Future: The Problem of Philosophical Progress*, Wiley.

Zagzebski, Linda (1994). "The Inescapability of Gettier Problems," *The Philosophical Quarterly* 44 (174): 65-73.

Zagzebski, Linda (1996). *Virtues of the Mind: An Inquiry Into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge University Press.