

## CHAPTER 1

### CONCEPTUAL AND DOCTRINAL PROJECTS IN NATURALIZED EPISTEMOLOGY

There is a widespread belief among intellectuals that the domain of philosophy shrinks as the domain of the special sciences expands, and that, someday, science might swallow up philosophy entirely. Some philosophers—philosophical naturalists—believe that this day may have already arrived. Naturalists hold that philosophy does share or should share the basic concepts and methodologies of natural science.

To determine whether the naturalists are right, one useful approach is to examine proposals for naturalistic (or naturalized) *epistemology*, the recent attempt to transform theory of knowledge into a branch of natural science. In Western philosophy, epistemology has long been considered one of the most distinctively philosophic subjects. If even it can be naturalized, the days of philosophy as an autonomous discipline could be numbered.

Traditional epistemologists object to naturalized epistemology on the grounds that it eliminates the distinctively philosophical content of epistemology. In this dissertation, I argue that traditional epistemologists are correct to reject naturalism, but that new arguments are needed to show why they are correct. I establish my thesis first by critiquing two prominent versions of naturalism—which I call “optimistic” and “pessimistic”—and then by offering a proposal for how a renewed non-naturalistic epistemology must move forward.

Before I can outline how I plan on critiquing these two varieties of naturalism, I need to provide some important background exposition. In this introductory chapter, I will describe just what naturalized epistemology is supposed to be, in particular what it means for epistemology to “share the same concepts and methodologies of natural science.” It turns out that apart from sharing this very generic credo, advocates of naturalized epistemology have deep differences over what it means for epistemology to be continuous with science. I will show how these different recognized approaches fit into my categorization of “optimistic” and “pessimistic” naturalized epistemology.

Having surveyed different conceptions of the naturalist's project, I will then describe one of the most prominent objections to it: the charge that naturalism unnecessarily eliminates the normativity of epistemology. I will briefly sketch the responses naturalists typically offer. With this as a background, I will describe how my own distinctive critique of naturalized epistemology compares to this traditional objection, and outline the course this objection will take through the rest of my dissertation.

### **The variety of conceptual and doctrinal projects**

In his influential essay "Epistemology Naturalized" (1969a), W.V. Quine draws a distinction between "conceptual" and "doctrinal" projects in the traditional epistemology to which his naturalism is presented as an alternative. I find it useful to invoke this distinction to explain distinct but related projects within naturalized epistemology itself. Even though Quine critiques the manner in which traditional epistemology attempts to base its doctrinal project on its conceptual one, I find that many versions of naturalism follow the same pattern. (Whether or not Quine's naturalism does the same is somewhat more obscure.)

Quine begins by discussing the conceptual project in mathematics, which he compares to a similar project in epistemology. This project is concerned with "clarifying concepts by defining them, some in terms of others" (69). The doctrinal project is concerned with "establishing laws by proving them, some on the basis of others." (69–79). Quine then notes that the two projects are closely connected:

For, if you define all the concepts by use of some favored subset of them, you thereby show how to translate all theorems into these favored terms. The clearer these terms are, the likelier it is that the truths couched in them will be obviously true, or derivable from obvious truths. (70)

In epistemology, the doctrinal project attempts to explain how we might justify "our knowledge of the truths of nature in sensory terms" (71), whereas the conceptual project aids by defining the terms of that knowledge. Famously, Quine argues that the traditional epistemological project of translating

concepts of physical objects into the language of sense data had to fail, because of his indeterminacy of translation thesis. This failure, combined with the failure of traditional foundationalist proposals, spelled the death of traditional doctrinal projects in epistemology—not only the classical empiricist attempt to justify scientific knowledge by reference to the senses, but even the modern empiricist attempt to legitimize scientific discourse by “demarcation.”

No naturalized epistemologist is interested in traditional epistemology’s reductivist conceptual project or foundationalist doctrinal project. However the conceptual-doctrinal distinction is still at play for many naturalists, although at a higher, *meta*-epistemic level. While naturalized epistemologists no longer concern themselves with translating the content of empirical knowledge for the sake of justifying it, many are still concerned with analyzing or in some way defining the *concept* of “knowledge” itself, in order to answer the *doctrinal* question of whether and to what extent we have any knowledge in the sense provided by that definition.

In what follows, I first classify naturalized epistemologists according to their “optimistic” and “pessimistic” answers to the doctrinal question. The optimism and pessimism here is in relation to the traditional goals of epistemology, which I myself share. “Optimists” affirm that we can show human beliefs to be justified, by applying some naturalized conceptual project. “Pessimists” deny this, but would not consider themselves to be pessimists, because they urge that epistemology adopt new goals.

Optimistic naturalized epistemologists are united in the conviction that the empirical methodology of natural science can somehow show our beliefs to be justified, but there is a variety of views about what this methodology amounts to. Not surprisingly, every major semantic theory of the twentieth century—analytic, two-factor, natural kinds—has been applied to the project of understanding the reference of the concept of knowledge. I will, therefore, classify subvarieties of optimistic naturalism according to the semantic theories they rely upon.

Having presented these optimistic projects, I will turn to the pessimists. The first of these is Michael Williams (1996), who offers a “deflationary” approach to the concept of knowledge, which

focuses on the use of the *term* “knowledge,” rather than its reference in the world. The most prominent pessimist, however, is Quine himself. Though Quine would, in some moods, speak of human knowledge, the concept of “knowledge” does not figure prominently as a technical concept in his naturalized epistemology. Quine’s behaviorism generally rendered the epistemologist’s reference to subjects’ internal cognitive states to be of largely passing concern. As we will see in later chapters, Quine’s deep commitment to the principles of naturalism not only caused him to distance himself from the very idea of a conceptual project, but from many of the philosophical mechanisms used by epistemologists (naturalistic or otherwise) to engage in this project.

### **Optimistic naturalized epistemology**

#### *Analytic naturalized epistemology*

The first putatively naturalist epistemology worth discussing engages in a meta-epistemic conceptual project with deep ties to traditional epistemology. This approach seeks to offer genuine conceptual analyses of epistemic concepts such as “knowledge” and “justification,” but hopes to analyze these concepts into more basic concepts that are naturalistically respectable. This approach is exemplified in the epistemology of Alvin Goldman. Goldman’s early views sought to analyze the normative language of “justification,” for example, into the purely descriptive terms such as “believes that’, ‘is true’, ‘causes’, ‘it is necessary that’, ‘implies’, ‘is deducible from’, ‘is probable’.” These latter terms are “(purely) doxastic, metaphysical, modal, semantic, or syntactic expressions” and therefore neither epistemic nor normative (Goldman 1979, 2).

Examples of this analytic approach to the naturalistic conceptual project originally gained prominence as responses to the Gettier problem. One challenge of that problem was to identify a condition for knowledge that would explain why justified true beliefs that were merely “accidentally” true did not count as knowledge. A natural solution was to individuate knowledge by the *causal* origin of the belief. David Armstrong’s account, for example, treats knowledge as a kind of reliable indicator,

like a thermometer, in which “reliability” is understood as a lawlike connection between beliefs and facts (1973). Robert Nozick’s (1981) theory speaks of knowledge as “tracking” the truth, and analyzes “reliability” in counterfactual terms: a true belief counts as knowledge just in case the following holds: if it was true, it would be believed, but not otherwise. Goldman’s own (1986) version of reliabilism holds that a belief is justified just in case it results from a reliable belief-forming process, one that yields a greater percentage of truths than falsehoods, and counts as knowledge if it is both true and discriminates the truth from “relevant alternative” possibilities.

In the next chapter we shall examine whether the mere non-epistemic or non-normative status of doxastic, metaphysical, modal or semantic concepts is sufficient to guarantee their status as naturalistic. For the time being, however, the more interesting question is whether the approach of conceptual analysis itself is consistent with naturalism. Recognizing that the armchair approach of analysis has long been rejected by naturalists, Goldman urges that any adequate epistemology seems to “involve, or presuppose, analyses (or ‘accounts’) of key epistemic terms like ‘knowledge’ and ‘justification’” (Goldman 1986, 36). He goes on to protest against Quine’s infamous (1953b) attacks on analyticity, by insisting that there must be “some substance to the commonsense notions” of meaning and synonymy, that even philosophers who reject analyticity often perform something *like* conceptual analysis when they reason philosophically, and that presenting necessary and sufficient conditions is an indispensable approach to philosophical reasoning, even if it has a long record of failure (1986, 38–9).

In chapter 3, we will examine attempts to address Goldman’s first concern, and make naturalistic sense of analyticity. Suffice it to say that it is no small task. As to Goldman’s second concern, we will shortly discuss whether there is something sufficiently *like* conceptual analysis to do the philosopher’s task. This is particularly urgent, because Goldman’s third point about the indispensability of analysis in the face of its failure looks particularly implausible twenty years later, after the analytical debate over the Gettier problem has long fizzled out, and if any consensus has been

reached, it is only that a new approach to epistemology is needed. Naturalists, now under the guise of “experimental philosophy,” stress the diversity and cultural dependence of philosophical intuitions (Nichols, Stich and Weinberg 2003). Indeed it is arguable that the analytic “naturalists” whose roots are found in the Gettier problem are only accidentally related to naturalists like Quine, whose motivations were very different, as we shall find in chapter 2 and chapter 5.

If some version of analytic naturalism can be salvaged as a conceptual project, however, its doctrinal implication becomes apparent. Combining a successful analysis of “knowledge” (in terms of reliability, etc.) with results from cognitive psychology enables us to determine whether and to what extent human knowledge exists. Goldman thinks that his analysis at least permits us to accept that knowledge is logically possible, even if the analysis does not entail that such knowledge exists and doesn’t permit a knock-down answer to skepticism (1986, 55–6). To know if we know would require that we know our beliefs to result from a reliable process, and it is logically possible to know *this* (56–7). Only our best psychology, not any analysis, can inform us as to whether that possibility is actual. It is at this point that objections of circularity usually enter, but Goldman has the option of noting that arguments for skepticism only arise because of conceptions of knowledge uninformed by reliabilism, conceptions that require ruling out Cartesian alternatives that are not relevant. Of course not all naturalists are as confident as Goldman about the power of cognitive psychology to deliver good news (Stich 1990). And if the success of this doctrinal project depends on the success of conceptual analysis, doubts about the latter could turn the former into a “degenerating research program.”

#### *Two-factor semantical naturalized epistemology*

We need, therefore, to seek an approach to the conceptual project that is *like* traditional conceptual analysis, but not committed to the same substantive presuppositions about meaning and synonymy. Even if a philosopher is not tied to philosophic intuitions about the meaning of concepts like “knowledge” and “justification,” it may profit him to *begin* with those intuitions as an entrée to a

more sophisticated scientific theory, the results of which may or may not end up bearing much resemblance to the original intuitions. What counts for this kind of conceptual project is not so much allegiance to prior intuitions, but the predictive and explanatory power of the theorist's ultimate conceptualization. The current literature features proposals for theories of concepts supplanting the classical theory of concepts drawn on by conceptual analysis, and these proposals are relied upon, implicitly or explicitly, for alternative formulations of naturalized epistemology. I will mention two such theories of concepts, and some paradigm applications in epistemology.

The classical theory of concepts drawn on by conceptual analysis held that concepts expressed conjunctions of necessary and sufficient conditions, which could be discovered by the introspective reflection of the theorist. This theory was called into question by the Twin Earth thought experiments, which seemed to indicate that meaning of concepts could not be "in the head," because the reference of a term like "water" seems to vary in relation to the environment in which it is originally deployed (whether it is an environment containing H<sub>2</sub>O or XYZ). A recent view of concepts seeks to capture the insight of these thought experiments, while also preserving an element of the classical view. These "two factor" or "causal-descriptive" theories urge that one factor of meaning is determined by *a priori* factors, while a second is determined by external aspects of the natural or social environment. In the view of Frank Jackson (1998), for example, we begin with a description of water as a clear, liquid stuff found in rivers and streams around here. We *need* to grasp at least this much, if ever we are to eventually *discover* the reference of "water" in the external world (either H<sub>2</sub>O or XYZ). Importantly, we may end up revising our concept of "water," but we need to appeal to our intuitions about it before we can ever make that discovery. Other two-factor theories are even more unabashedly naturalistic than Jackson's, and urge that the descriptive component of reference is not *a priori*, but a manifestation of background theoretical knowledge acquired through ordinary empirical means (Boyd 1991; Stanford and Kitcher 2000)

How might the two-factor view of concepts be implemented in naturalized epistemology? One theorist who seems to be implicitly committed to the view is Philip Kitcher. In his essay “The Naturalists Return” (1992), Kitcher considers Goldman’s reliabilism to be a holdover of analytic epistemology, and claims that while “reliabilism gives a promising *start* to formulating a meliorative naturalistic enterprise,” it is “not the panacea for the problem of analyzing justification” (69). He believes that when analytical naturalists define ideal standards of justification in advance of inquiry, they invite skepticism and fail to shoulder the proper task of epistemology. Goldman’s reliabilism, treated as an analysis of knowledge, invites counterexamples of true beliefs caused by reliable processes in a bizarre manner, for instance. It is always possible to refine definitions to better capture our intuitions about knowledge, but this does little to improve our understanding of worthwhile cognitive goals or improve our ability to reach them. What Kitcher means by the “meliorative project” is precisely the kind we might guess to be recommended by a two-factor approach to reference<sup>1</sup>:

Traditional epistemology has an important meliorative dimension. Bacon and Descartes were moved to epistemological theorizing by their sense of the need to fathom the ways in which human minds can attain their epistemic ends. If analysis of current concepts of rationality and justification, or delineation of accepted inferential practices, is valuable, it is because a clearer view of what we now accept might enable us to do better. Conceptual clarification has a role to play in advance of inquiry, even when we understand that our current concepts might give way to improved ones. (64)

Kitcher speaks here of the meliorative project of traditional epistemology, but it is clear from the rest of the essay that he sees naturalism as sharing the tasks of traditional epistemology, if not the means.

How, on Kitcher’s view, do we come to understand these worthwhile cognitive goals and assess our prospects of achieving them? He would implement the doctrinal project of naturalized epistemology by looking to the history of science, and more fundamentally, to our evolutionary heritage. In the course of examining our actual cognitive practices, and the basic equipment we inherited to undertake them, we may discover that achieving our cognitive goals is not always consistent with our *a priori* epistemic standards. We may find that we need to replace rather than

---

<sup>1</sup> Stanford and Kitcher (2000) develop an explicit two-factor theory of reference.



analyze the dichotomies of “rational/irrational” or “justified/unjustified,” out of the need to give a richer portrait of factors contributing to the limited human animal’s achievement of its cognitive goals.

Kitcher is aware, of course, that not all naturalists would find epistemological solace in an examination of the history of science or in the human evolutionary heritage. The bulk of his doctrinal studies concentrate on answering their worries. These skeptics might doubt, for example, that the cognitive equipment of our ancestors needed to be geared towards the acquisition of significant truths in order for the race to evolve successfully. But even if our ancestors developed some remedy to possible evolutionary shortcomings, the more serious naturalist challenge to the possibility of outlining the means and ends of human cognitive progress is that posed by Quinean and Kuhnian underdetermination arguments. These suggest that science has not developed by a series of logically-sanctioned steps aimed at an ultimate cognitive goal, but instead by a series of paradigm shifts that could have been otherwise, because of pragmatic decisions about auxiliary hypotheses, etc. Kitcher believes that the only response to this challenge is to examine the historical record even more carefully, to show that instances of underdetermination are not as pervasive as critics suggest. (In the final chapter of this dissertation, we will return to the topic of the underdetermination, which underpins some of the most basic naturalistic assumptions—a point Kitcher does not seem to fully appreciate.) Kitcher also believes he can examine the history of science to answer persistent objections from Larry Laudan (1984) and to show that the putative diversity of historical scientists’ goals can be reduced to “a single, compelling, conception of cognitive value,” which Kitcher calls “significant truth” (1992, 102). Kitcher delivers a ground-level examination of these very questions in his exhaustive treatment, *The Advancement of Science* (1993).<sup>2</sup>

Another naturalized epistemologist, Hilary Kornblith (2002), subscribes to the same conceptual project as Kitcher, but goes further still. Appealing explicitly to Boyd’s two-factor semantics, Kornblith argues that the epistemologist’s reference to knowledge can be understood as

---

<sup>2</sup> For more on the debate between Laudan and Kitcher, see Rosenberg (1996).

reference to a *natural kind*, understood on Boyd's model of natural kinds as causal homeostatic mechanisms. A homeostatic mechanism is a natural cluster of highly correlated properties or elements, the combination of which promotes a self-reinforcing stability, such that predicates describing the cluster are readily projectible. The combination of hydrogen and oxygen atoms in the water molecule is a good example.

Kornblith thinks that knowledge is a natural kind like water because cases of knowledge "have a good deal of theoretical unity to them" rather than being a "gerrymandered kind" constructed by human convention (10). The "theoretical unity" of knowledge is first understood by reference to the theoretical unity of *belief*. Kornblith looks to animal ethology's extensive use of intentional idioms to describe, explain and predict a variety of animal behavior. Even ants returning to the colony seem to "represent" their direction and distance traveled. More sophisticated animals exhibit genuine beliefs and desires, when the information represented comes to form a stable product available for multiple uses, depending on the animal's desire. Kornblith understands knowledge as a species of belief, and adds that it features an extra dimension of explanatory/predictive value, also recognized by animal ethologists. Whereas the actions of individual animals could always be easily explained by reference to mere beliefs, explaining how it is their species possesses the cognitive capacities that permit successful interaction with the environment requires the appeal to *reliable* belief-forming processes, i.e. knowledge. In short, nature has selected these cognitive capacities for their survival value, which in turn ensures the perpetuation of the capacities themselves (57–9).

Kornblith's natural kinds-oriented conceptual project has important doctrinal implications. To show that organisms really do possess the relevant reliable capacities, he must answer critics like Brandom (1998) who allege that judgments about reliability vary in relation to the scope of the organism's environment, and theorists may circumscribe environments arbitrarily, according to their interests. Kornblith (2002, 65–9) responds that the concept of an environment is itself a technical concept of ecology, one that is just as naturalistically respectable as many used by biologists.

Knowledge, then, is specifically an *ecological* natural kind. Kornblith must also oppose popular positions in epistemology according to which animals cannot possess knowledge or beliefs, because both concern the essentially social practices of giving and asking for reasons (69–102), and because knowledge requires a kind of self-conscious reflection of which animals are incapable (89–136).

As we have progressed from analytic naturalism to Kornblith’s two-factor natural kinds naturalism, we have become less focused on the *concept* of knowledge and more focused on the *metaphysics* of knowledge itself. He goes the furthest here, seriously downplaying the need to appeal to philosophic intuition. Responding to Goldman’s (1993) contention that naturalized epistemology should at least describe our epistemic “folkways” (our inherited intuitions about knowledge) before engaging in object-level study, Kornblith notes that in chemistry, we do not bother cataloguing folk chemistry; instead we “can simply skip straight to the project of understanding the real chemical kinds as they exist in nature.” He concludes that “we should take seriously the possibility that a similar strategy might be equally fruitful in epistemology” (2002, 19). Arguably the next version of naturalism would seem to push Kornblith’s suggestion to the extreme, avoiding discussion of concepts entirely and going straight to the metaphysics of knowledge.

#### *Epistemic supervenience naturalized epistemology*

In an influential critique of Quine’s naturalized epistemology—and cognizant of Quine’s antipathy towards conceptual analysis—Jaegwon Kim (1988) proposes a method of formulating epistemological criteria that avoids controversial reliance on philosophical accounts of meaning. Utilizing a concept he has developed in detail largely in connection with topics in the philosophy of mind, Kim argues that it must be that epistemic properties *supervene* on natural ones:

[I]f a belief is justified, that must be so *because* it has certain factual, nonepistemic properties, such as perhaps that it is “indubitable”, that it is seen to be entailed by another belief that is independently justified, that it is appropriately caused by perceptual experience, or whatever. That it is a justified belief cannot be a brute fundamental fact unrelated to the kind of belief it is. There must be a *reason* for it, and

this reason must be grounded in the factual descriptive properties of that particular belief. (399)

A number of other philosophers, including Van Cleve (1985) and Sosa (1980), have endorsed the notion of epistemic supervenience, without necessarily seeing it as a naturalization proposal. Although Kim is widely recognized as a critic of Quine's naturalism, his critique acknowledges the viability of naturalistic projects rivaling Quine's, such as Kitcher's and Goldman's (Kim 1988, 394–9). His own supervenience proposal, in fact, can be transformed into a kind of naturalism, provided that the properties that epistemic properties supervene upon are themselves natural properties, and also provided that the nature of the supervenience relation itself is naturalistically respectable.

Speaking loosely, supervenience is the determination of a higher level property by a lower level property. To say that higher-level property A supervenes on lower-level property B is to say that any two objects which do not differ in lower-level B properties must not differ in their higher-level A properties. Or: there cannot be a difference in A properties without a difference in B properties. Supervenient A properties must have *some* subvenient B properties of some type or other, but if anything has these subvenient B properties, the supervenient A properties must obtain. The nature of that “must” is of some importance. The strong notion of supervenience needed to support a *determination* relation between B and A properties requires some kind of necessity. At one point in his discussion of supervenience of the mental, Kim's favored option is to find a kind of *nomological* necessity (1985). If there is a lawlike relationship between B and A properties, that would secure the necessary strong supervenience. We will discuss this view of necessity in chapter 2.

With the concept of supervenience in hand, Kim seems to have formulated a metaphysical stand-in for the conceptual project in epistemology, and can proceed to look for answers in the doctrinal project. He can search the relevant science to see if a lawlike relationship does exist between any properties and epistemic properties. In an essay on the supervenience of the *mental* on the physical, he considers the possibility of psycho-physical laws in the context of the problem of the multiple realizability of the mental. He proposes that the physical instantiation of these psycho-

physical laws may consist of lengthy disjuncts of distinct properties. Whether science could ever uncover or deal with laws of this type is not clear. To the extent that epistemic properties are themselves dependent on doxastic ones, the same problem may apply to epistemic supervenience.

The nomological supervenience concept is at best a placeholder for scientific discoveries waiting to be made. To the extent that it requires the discovery of nomological relationships, it may draw strength from discovery of the very kind of homeostatic mechanisms that Kornblith believes animal ethology to have catalogued. Indeed if supervenience requires a notion of nomological necessity, there may be little difference between Kim's and Kornblith's views in the end. Later (2005), Kim appears to rely on a conceptual form of necessity. Either way, supervenience has affinity to conceptual projects we have already considered.

Common to Goldman, Kitcher, Kornblith and Kim is the conviction that knowledge *really is something*. Consequently they look to the natural sciences to “uncover” knowledge of what that something really is. But this is not the only possible naturalistic approach to answering skepticism. It is possible to affirm the truth of statements concerning knowledge without being ontologically committed to the substantive *existence* of knowledge-stuff. This possibility is one that has been explored by the next category of naturalized epistemology, one that has not always been recognized as such: deflationary naturalism. In discussing this next category, however, we enter into the realm of what I call “pessimistic” naturalized epistemology.

### **Pessimistic naturalized epistemology**

Dividing philosophical views according to the labels of “optimistic” and “pessimistic” is, of course, loaded with value judgments. An optimistic expects success; a pessimist, failure. The present category of pessimistic naturalized epistemologies counts as pessimistic only insofar as they expect failure to achieve traditional epistemological goals. But these epistemologies are not absolutely pessimistic: they believe that their proposals offer alternative goals that can be readily achieved. I can

only state here that I myself happen to side with (most of) the goals of traditional epistemology, and for this reason I am exercising the privilege of categorizing epistemologies relative to that position. At the end of the dissertation, I hope to have established that traditional epistemological goals—including some of the traditional means to these goals—should not be abandoned for the reasons naturalists are wont to abandon them. So hopefully the present categorization will prove to be useful.

*Deflationary naturalized epistemology*

Deflationary views in philosophy are generally concerned with explaining how one might affirm a type of philosophic truth without being committed to the existence of substantive properties related to predicates expressed in those truths. The classic example is the deflationary view of truth, which holds that the meaning of the truth predicate is exhausted by the disquotational formula: “Snow is white” is true if and only if *Snow is white*. This conception avoids the commitment to a substantive truth relation, and consequently avoids thorny metaphysical questions about the nature of correspondence or of facts to which truths must correspond. For some time now, deflationary views of “knowledge” have also been available, particularly from the contextualist wing of epistemology.<sup>3</sup> Until recently, however, it has not been obvious how deflationism might also count as a form of *naturalism*.<sup>4</sup> If knowledge is not a substantive existent, what about it would scientists have to study?

One clue is offered by Huw Price (2004). Speaking of philosophic issues apart from epistemology, Price notes that we can make a distinction between *object-naturalist* and *subject-naturalist* approaches to central concepts in these fields. The object-naturalist is concerned with discovering the substantive properties to which philosophic concepts refer, and as such employs the methods of natural science to discover them. The naturalized epistemologies we have considered so far surely count as object-naturalist. But the subject-naturalist is not so much concerned with substantive

---

<sup>3</sup> See Pritchard (2004) for a survey of prominent deflationists about knowledge, including Sartwell, Foley, and Williams.

<sup>4</sup> Henderson (1994) actually considers contextualism as a kind of naturalism, but not for the reasons I outline.

properties as he is with subjects' *use* of philosophic *terms*. The subject-naturalist in epistemology, then, would be primarily concerned with human use of the term "knowledge." This, as it happens, is the celebrated project of the contextualists.

An excellent case in point is Michael Williams (1996). Williams himself characterizes his position as deflationary (111–3), drawing explicit inspiration from Quine's deflationism about truth. For reasons we will discuss later in chapter 5, Williams later critiques Quine's views on naturalized epistemology (254–65). Nevertheless, his own deflationary view may count as a form of subject-naturalism, if Price's conception here is useful. He is surely no object-naturalist: contrary to Kornblith, he denies that knowledge is anything like a "natural kind," or any *thing* at all, denying the position he calls "epistemological realism." Williams's motivation for adopting this position emerges out of his critique of traditional epistemology. He has argued that skepticism is a consequence of foundationalism, in particular the view that our beliefs have foundations in the senses, and a consequence of the "totality condition," the idea that all of our knowledge can be assessed at once. When the skeptic considers these possibilities, he loses confidence in the possibility of sensory foundations, and in doing so loses confidence in the totality of knowledge. Williams urges that we abandon foundationalism and the totality condition in order to avoid the problem of skepticism. But this solution to skepticism is very dissatisfying: without foundationalism, we crave some other assessment of the source of our knowledge. Williams, therefore, takes it upon himself to explain what is wrong with the craving in the first place. His main response is that knowledge is not an object of theory in need of any explanation.

Williams thinks that it may be true that we know many things in the proper contexts, but that this is not in virtue of anything in common among the cases called "knowledge." To show that there is something significant in common among such cases, one would need to demonstrate that cases of knowledge have a kind of "theoretical integrity" (103). But our beliefs—to say nothing of our knowledge—are not *topically* integrated. We do not store them in the form of a single, all-

encompassing axiomatized system. All that remains is the possibility that they are *epistemologically* integrated, i.e., subject to the same constraints, tracing from the same sources. Of course Williams believes that foundationalism only leads to skepticism, so the claim that knowledge exhibits this kind of integrity is in danger of giving “knowledge” the status of a theoretical term (like “phlogiston”) that fails to refer if the theory behind it is false. There are also terms such as “table” or “heat” whose reference is thought to be fixed pre-theoretically or theory-independently. But Williams can find no reason to think “knowledge” functions in the same way (109–10). In chapter 4, after considering evidence about our formation of the concept of “know” that casts doubt on deflationism about “belief,” we will return to the question of the pre-theoretical integrity of “knowledge.”

For this reason, Williams thinks all we can hope for from epistemology is a deflationary account of knowledge:

A deflationary account of “know” may show how the word is embedded in a teachable and useful linguistic practice, without supposing that “being known to be true” denotes a property that groups propositions into a theoretically significant kind. We can have an account of the use and utility of “know” without supposing that there is such a thing as human knowledge. (113)

This is as close as Williams comes to stating a conceptual project for his epistemology. Unlike previous object-naturalists, he is not concerned with the question of what knowledge *really is*. He is primarily interested in the concept itself, and even then, mainly the *word*. The naturalistic investigator can then make use of this proposal for the conceptual project, by examining our actual linguistic practices to see if they stand up to Williams’s contention that our attributions of knowledge lack any obvious theoretical integrity. If the investigator determines that this is true, this amounts to Williams’s version of the doctrinal project in epistemology: by debunking knowledge as a natural kind, he will have dissolved our craving for epistemological explanations, and in doing so he will have shown why we can reject skepticism without needing to assess the totality of our knowledge.

We need, then, to briefly describe the kinds of investigations that would be relevant to supporting Williams’s contentions about the linguistics of “knowledge.” These, I think, would be little



more than the familiar examples entertained by contextualists, concerning the shifting standards of justification from context to context. To defeat the foundationalist view of theoretical integrity, Williams believes that he need only show that there is never any single type of proposition which, in virtue of its contents, “will have an epistemic status it can call its own” (113). Here is a sample of the kind of ordinary survey that would support this:

In both science and ordinary life, constraints on justification are many and various. Not merely that, they shift with context in ways that are probably impossible to reduce to rule. In part, they will have to do with the specific content of whatever claim is at issue. But they will also be decisively influence by the subject of inquiry to which the claim in question belongs (history, physics, ornithology, etc.).... Not entertaining radical doubts about the age of the Earth or the reliability of documentary evidence is a precondition of doing history *at all*. There are many things that, as historians, we might be dubious about, but not these. (117)

To these “disciplinary” constraints, Williams also adds “dialectical” constraints and “situational” constraints, which derive from idiosyncrasies of conversational and evidential contexts. The role of context is even more important to Williams than simply providing evidence against foundationalist theory. It not only helps to show why we shouldn’t worry about skepticism, but shows the consequences of what happens if we do. The disciplinary constraints he mentions not only keep us on task as historians and physicists, but *stop us from doing epistemology* (122). Paradoxically, it turns out that this “methodological necessity” is epistemically good for us, the epistemologist’s questions about the totality of our knowledge actually cause us to *lose* our knowledge insofar as we share his doubts, insofar as they cause us to suspend the “interact relations with our environment that...are crucial to much ordinary knowledge” (358).

At one point, I almost decided to classify Williams’s deflationism as a kind of “optimistic” naturalized epistemology. Williams does attempt to show how deflationism helps respond to skepticism by affirming our knowledge of many things. For all of this, however, his theory may still be deeply dissatisfying to the traditional epistemologist. He would blame this dissatisfaction on philosophers’ lingering foundationalism, which he takes to be hopeless. (In the final chapter of this dissertation, we will revisit the question of whether better formulations of foundationalism might solve

rather than cause the problem of skepticism.) In this respect he has much in common with the category we are about to examine, Quine's naturalized epistemology. As we shall see, however, the chief difference between Williams and Quine is that Quine is on the whole less concerned with dealing with the skeptic, and does not even make substantial theoretical *use* of the concept of "knowledge" in his proposal.

### *Quinean naturalized epistemology*

Quine's position is the last we'll survey in this chapter, but it was also the first significant proposal for naturalized epistemology of the 20<sup>th</sup> century. Many of the previous views took inspiration from Quine, and take themselves to be following his research program. But even among those who revere his example, there seems to be a general consensus that Quine went too far, that his position represented an unreasonable abandonment of core elements of genuine epistemology. We can see why this is generally accepted by considering a widely-quoted representative passage from "Epistemology Naturalized" (1969a, 82-3) in which he describes the proper subject matter of naturalized epistemology:

Epistemology, or something like it, simply falls into place as a chapter of psychology and hence of natural science. It studies a natural phenomenon, viz., a physical human subject. This human subject is accorded a certain experimentally controlled input—certain patterns of irradiation in assorted frequencies, for instance—and in the fullness of time the subject delivers as output a description of the three-dimensional external world and its history. The relation between the meager input and the torrential output is a relation that we are prompted to study for somewhat the same reasons that always prompted epistemology; namely, in order to see how evidence relates to theory, and in what ways one's theory of nature transcends any available evidence.

Quine's emphasis here on studying the actual psychological history of a subject's cognitive processes led many critics to regard him as abandoning the normative element of epistemology, the attempt to assess the *justification* of our beliefs. In this respect he seems very much like Williams. This has led scholars of the field to classify Quine's views as "eliminative" naturalism (Maffie 1990), or "replacement" naturalism (Almeder 1990; Feldman 2006). Many critics (e.g., Kim 1988) have seen

this alleged abandonment of concern with normativity to be the chief flaw in Quine's position.

Whether and in what sense Quine really does abandon normativity is, of course, a subject of some debate, of course, and in the next section and in the next chapter we will examine the question in some detail.

For the time being, however, I want merely to focus on the fundamental *methodological* uniqueness of Quine's proposal in relation to the other naturalisms so far surveyed. I have categorized naturalisms so far according to their particular conceptual and doctrinal projects, a distinction which comes from Quine himself. Quine, of course, would lend no quarter to *analytic* naturalism, as he is famously (or infamously) skeptical about notions of meaning and synonymy underlying the idea of conceptual analysis (Quine 1953b). Quine would not be any happier with the two-factor theory of concepts that underpins the second version of naturalized epistemology, on the grounds that it relies on a picture of reference inconsistent with his infamous inscrutability of reference thesis (Quine 1969c). Even the doctrine of natural kinds—drawing as it does on cherished examples from the natural sciences—feels the brunt of Quine's withering critique (Quine 1969b). Science based on natural kinds is “rotten to the core,” says Quine—albeit a kind of rot necessary for progress to a better, more naturalistically respectable science based on mathematics (1969b, 133). Even the epistemologically minimalist doctrine of supervenience would irk Quine, owing to its essential reliance on the concept of “necessity.” “Necessity” is an intensional concept, like “belief,” that Quine judges to be incompatible with his naturalistic extensionalism (1953b). We will explore many of these Quinean objections to optimistic naturalism in chapters 2 and 3.

Without recourse to standard philosophical methodologies, by what means does Quine hope to naturalize epistemology? Here it is worth repeating that although the distinction between conceptual and doctrinal projects in epistemology is Quine's, it is not one that he applies to his own work. He does not begin by analyzing the concept “knowledge” and then use science to determine to what extent the concept is applicable. In fact Quine has little interest in the concept of “knowledge” to begin with.

He is more concerned to examine the relationship between “evidence” and “theory,” where the latter is understood primarily *linguistically*, rather than cognitively. Surprisingly, he also has very little concern with assessing the *reliability* of this linguistic output. “Reliability” is a concept intimately connected to truth, but for Quine, epistemology is chiefly concerned with explaining the *pragmatic* value of our linguistic outputs, i.e. their facility in allowing the prediction of experience and subsequent control of our environment. At the same time, some of Quine’s writings on epistemology appear to offer a naturalistic answer to skepticism (Quine 1974; 1975b). It is not immediately obvious how to reconcile this approach with Quine’s pragmatism, but we will attempt to do it in chapter 5.

In the end, I will argue that Quine’s naturalism is the most naturalistic of the naturalisms, and that if we want to evaluate naturalized epistemology in its most fundamental terms, we must evaluate Quine’s most fundamental theses (this will be the subject of chapter 6). But before presenting my polemical strategy, I will survey one of the usual objections raised by prominent critics of naturalism, to show why I think my strategy raises concerns that are more fundamental.

### **A representative objection to naturalism: the normativity objection**

Objections to naturalized epistemology vary with the variety of naturalism. Every version considered so far depends on a distinctive philosophic methodology—and on the philosophic presuppositions of that methodology—and any objection to this methodology would sensibly count as an objection to the associated epistemology. It is curious, however, how little most critics of naturalized epistemology focus on these basic philosophic considerations. Many instead focus on objectionable features and alleged internal inconsistencies of any approach that makes natural science the essential methodology of a discipline traditionally thought to be purely philosophical. In this section, I will survey one such objection that contains elements of many of the other typical objections, and so serves as a good representative. This is the objection that naturalized epistemology fails to do justice to a central aspect of serious epistemology: its status as a *normative* discipline. I will

argue that traditionalist objections to naturalism raised on these grounds are not ultimately convincing. If the naturalists' basic methodological and philosophical presuppositions are left unaddressed, the normativity challenge can be answered through any number of artful dodges.

One of the most prominent critics to register the normativity objection to naturalism is Jaegwon Kim (1988). Kim specifically targets Quine's version of naturalized epistemology, but in a manner that could generalize to many of the other naturalisms we have examined (even though in a way his own position is naturalistic). Kim alleges that a naturalized epistemology such as Quine's aims only to describe or even explain how our beliefs are formed, using the resources of descriptive cognitive psychology. Kim assumes that Quine is only interested in a descriptive or explanatory effort, and *not* in the normative project of explaining how our beliefs are *justified*. There is something right about Kim's observation here: there is some conception of justification that Quine is not interested in discussing. As we shall see later in this section, however, this does not imply that Quine rejects the possibility of *any* normative role for epistemology. Whatever the proper interpretation of Quine, however, it is clear that this is a natural objection to the view that epistemology could somehow be based on science: since science is thought to be mainly descriptive and/or explanatory, how could science concern itself with a matter traditionally associated with philosophy, the question of right and wrong in the way of believing?

One critic of naturalism who has developed the normativity objection in more detail is Harvey Siegel. In one essay (1989), Siegel critiques the naturalized epistemology of Ronald Giere (1985; 1988). Giere's "evolutionary" naturalism attempts to "explain how creatures with our natural endowments manage to learn so much about the detailed structure of the world" (Giere 1985, 339–40). Siegel alleges that Giere's epistemology abandons concern with the *rationality* of science, and further alleges that the concept has "no place in a naturalized philosophy of science" (Siegel 1989, 366). Siegel (1989) objects that the scientific study of science, while unobjectionable in its own right, cannot answer traditional philosophical questions about science, questions such as "Is there a scientific

method which warrants or justifies claims arrived at by its use?, Are there principles of theory choice which render such choice rational?, What are the epistemic relationships between theory and evidence, and between those two and truth?” (368). But Siegel contends that Giere has not shown these to be non-viable questions. So if naturalized epistemology is really incapable of examining questions of rationality, and these questions are viable, it is failing to deal with important, viable *philosophical* questions, and also failing to do the job of any philosophy of science worth its salt.

Of course Siegel’s objection only works if we take it for granted that there is no room in the scientific study of science for questions of rationality. Giere’s position on this question is qualified. He says that rationality “is not a concept that can appear in a naturalistic theory of science—*unless reduced to naturalistic terms*” (Giere 1985, 332, emphasis mine). As if to anticipate the suggestion that Giere or someone else could actually provide a naturalistic reduction of rationality, Siegel offers another objection also typically associated with traditionalist critics of naturalism. Science could never answer questions about rationality, he argues, because

To answer [these questions] scientifically would be to beg the question—e.g., any answer to the question of the relationship between evidence and a justified theory, if arrived at scientifically, would depend upon exactly the same relationship between it and the evidence for it as it recommends for the relationship between any justified theory and the evidence for it. Because these general questions about the epistemology of science cannot be answered naturalistically without begging the question, they cannot be so pursued. (Siegel 1989, 369)

So there are two questions in need of answering. First, is it true that science has no resources to deliver an account of rationality? Second, is it true that any putatively scientific account of rationality would beg the question?

Regarding the first question about the ability of science to deliver an account of rationality, Giere evidently does intend to cash out his qualification that rationality might still be reduced in naturalistic terms. In his response to Siegel, Giere contends that there is a notion of rationality readily available to the naturalist: *instrumental* rationality:

To be instrumentally rational is simply to employ means believed to be conducive to achieving desired goals. . . .Actions that do not fit the model are labeled “irrational”. .

. . [T]here is also a more *objective* sense of instrumental rationality which consists in employing means that are not only believed to be, but are *in fact* conducive to achieving desired goals.

This latter, objective, sense of instrumental rationality provides the naturalistic theorist of science with ample means for making normative claims about science. These claims, however, are not autonomous but are grounded within science itself. It requires empirical research to determine whether a particular strategy is in fact likely to be effective in producing valuable scientific results. (Giere 1989, 380)

As an example of this kind of rationality, Giere mentions scientists' adoption of the continental drift hypothesis in the 1960s. Although the hypothesis had been widely regarded as implausible until that time, theorists reasoned that if true, the hypothesis would imply that strips of the ocean floor should exhibit distinctive patterns of magnetism. Upon finding that these patterns existed, scientists choosing to adopt the continental drift hypothesis would be making an instrumentally rational decision: treating successful novel predictions of hypotheses as evidence for the same hypotheses has, historically, been an effective means of adopting models for real processes.

Before we address the second question, about whether Giere's proposal would beg the question, it is useful to examine Siegel's likely response to the naturalist's answer to the first question. In a later essay, Siegel (1990) addresses Laudan's similar (1984, 24, 34) proposal for how normativity can be naturalized via instrumental rationality, by considering an example of Laudan's about how scientists might discover the preferability of double-blind to single-blind experiments in medical research. Siegel contends that evidence of the history of the different types of experiment would establish the preferability of double-blind experiments *only* if investigators value learning the genuine effectiveness of drug treatments, as opposed to, say, learning the *approximate* effectiveness of drugs in the quickest and cheapest manner possible. If investigators valued approximate effectiveness more, they might well side with the value of single-blind experiments. Siegel's point is that preferring double-blind methodology is the only acceptable preference, and that this preference is justified "not instrumentally, but *epistemically*: double blind experimentation provides better *evidence* for a drug's efficacy than single-blind experimentation, because it controls for an additional source of possible error" (1994, 301). In suggesting this difference between epistemic and instrumental rationality, Siegel

is relying on some of the same ideas behind his charge that science begs the question in attempting to provide a scientific account of rationality: that charge presupposes the possibility of a non-scientific, presumably *a priori* standard for assessing the justificatory value of evidence—a presupposition that we will examine shortly. Another way of looking at Siegel’s objection is that even if we do conceive of the justification of double-blind experimentation instrumentally, we can do so only if we have offer *a priori* justification of the goal of truth, rather than the goal of approximate truth found in a cost-effective way. His assumption is that naturalists can naturalize instrumental rationality *only*, not the intrinsic or categorical rationality of the ends of inquiry. Should we accept this assumption?

Siegel realizes that naturalists such as Laudan (1987, 29) have offered proposals for naturalizing not only epistemic means, but epistemic ends as well. Laudan’s proposal claims, for example, that certain ends can be disqualified by science if they are “utopian” and unachievable by human beings, or if their truth conditions cannot be specified, or if the means to their achievements cannot be specified. But Siegel insists that the unachievability of a goal is only a reason to disqualify it if we presuppose an instrumentalist conception of rationality (which would beg the question) (Siegel 1990, 307). Why he thinks this is unclear: it seems that Laudan need only rely on the conventional wisdom that “ought” implies “can” (see also Kitcher 1992, 83–7). To say that X is an instrumental value is to say that that X is valuable only if it achieves value Y. This is not the same as saying that X is valuable only if it *can be achieved*, full stop. A better objection to Laudan’s position, I think, would be to say that just because science can help *disqualify* ends in this manner does not imply that it can help *establish* them in the first place. The challenge for the naturalist is to show how this might work, and perhaps Siegel is presupposing that the ends of inquiry can be justified only by appeal to *a priori* considerations. We shall examine that presupposition presently.

Laudan (1990a, 316) observes that a central assumption of Siegel’s objection is that “there is something called epistemic rationality or epistemic justification which stands *outside* the analysis of ends/means connections.” As we have seen, this presupposition is relevant not only to Siegel’s last



objection to the possibility of naturalizing ultimate ends, but also to his more general contention that any such attempt would beg the question. In his article, Laudan never purports to offer a naturalization of ultimate ends. He contends that all rationality is instrumental in relation to desired ends, *whatever those ends happen to be*. This is certainly a position naturalists could accept provisionally in the absence of establishing ultimate ends of their own: they could simply conclude with Hume that reason judges not of ultimate ends, that it is merely the slave of the passions. The more important point that Laudan makes, however, is that even though Siegel's objection relies on the existence, or at least the possible existence, of some *a priori* standard of rationality, he never tells us what it is or how we can discover it. Laudan mentions the example of the thesis that successful, novel predictions tend to confirm their associated hypotheses. Philosophers have sought an *a priori* standard according to which this counts as evidence, but have failed. All that philosophers managed to determine was that "we could attempt to ascertain whether theories which had proved themselves successful over the long run were theories that could have successfully satisfied the demands of this rule" (321). Indeed, naturalists such as Quine (1953b) would insist that there is no way in principle of defining objective standards of evidential confirmation, because theory is underdetermined by evidence, and theory choice is driven, ultimately, by pragmatic factors. Laudan himself would object to Quine's underdetermination thesis (in our final chapters, we will see why), but the present point is that even attempts to define standards of evidence by reference to the history of science (like Laudan's) face an uphill battle against the underdetermination thesis—to say nothing of attempts to define these standards *a priori*. As far as I can tell, Siegel's subsequent (1996) response to Laudan does nothing to address this problem. Nowhere does he attempt to tell us what the *a priori* standard of evidence is or how we are to discover it. As far as we are concerned now, therefore, Laudan's view that all rationality is instrumental may be true: Siegel has offered no reason to think that reason is anything other than a slave to the scientist's passions.

Of course we should not discount entirely the possibility that other naturalists might still succeed at naturalizing epistemic *ends* themselves. Laudan does not consider this possibility, but other naturalists have. We have already discussed how, in response to Laudan's (1984) contention that the history of science is scattered with a plethora of different ends of inquiry, Kitcher (1993) argues that, with appropriate care, these can be reduced to a single end, the pursuit of "significant truth." Kitcher argues that current scientific theories survive because of having undergone a process of "natural" selection, in which theories with the greatest predictive and explanatory power are the ones that "survive." If Darwinian natural selection can explain the teleological function of biological processes in terms of adaptive success (Wright 1976), it seems reasonable that Kitcher could offer a similar quasi-Darwinian explanation of the teleological function of scientific theories, and underwrite the "significance" of his significant truth (Rosenberg 1996, 18). Of course whether or not this explanation is consistent with *realism* about the products of science is a matter of some controversy; it depends in large part on the verisimilitude of the stock of hypotheses scientists begin with before they eliminate all but the "fittest." Whether science can account for the original verisimilitude of our ancestors' theories as a product of cognitive evolution is a matter of some controversy, by Kitcher's own admission (1992, 92–3; Rosenberg 1996, 23).

Whatever the nature of the ends of science, a naturalist's formulation of it need not always cohere with our pre-theoretic conception of rationality. As we discussed in some detail while examining proposals for "analytic" naturalism, a naturalist may be content to begin with intuitions about rationality, without remaining faithful to them at the end of investigation. The folk intuition of rationality may not have the same predictive and explanatory power as the conception fully informed by the history of science and cognitive psychology.

In any case, as Alexander Rosenberg (1996, 25) observes, naturalists can characterize the end of science in terms even more generic than "significant truth": without resolving the realism/anti-realism debate, they can point to the ultimate end of science as nothing more than "prediction and

control” (of experiences, if not of real entities). Why then are prediction and control to be taken as ultimate ends? Rosenberg considers the possibility that these might also be justified through their adaptive value, via Darwinian natural selection. But he thinks that this would go too far and lead to a vicious circularity. His alternative runs as follows:

The only way naturalism can avoid [conceding the possibility of incommensurable goals of inquiry] is to show how those who reject naturalism in fact willy-nilly embrace prediction and control as the ultimate epistemic value, despite their claims to the contrary. . . . It must claim that its competitors’ rejection of prediction and control as the ultimate aims of enquiry is belied by their own actions, choices, decisions, and provisions. (26–7)

But this self-refutation strategy seems quite dubious. Perhaps non-naturalists in everyday life can *appear* to embrace predictive and explanatory value without its being the standard of their inquiry. And it seems likely that non-naturalists could always formulate some alternative characterization of their predictive/explanatory behavior, as most theorists charged with self-refutation are usually able to do. Naturalists respond to non-naturalists like Siegel in the same way. If Siegel claims that naturalists have to rely on non-naturalistic conceptions of evidence, naturalists can simply provide a naturalistic explication of their practice.

The more important question to ask is: is Rosenberg correct that explaining the value of prediction and control via Darwinian theory would be viciously circular? Certainly it would be circular in some sense: the question is whether the circularity would be *vicious*. For as much as Rosenberg appreciates the willingness of naturalists like Giere, Laudan and Kitcher to embrace science and disavow the *a priori* approach to epistemology, he seems to ignore Quine’s reason for seeking a naturalization of epistemology in the first place: the fact that attempts at ground-up reconstructions of science have always seemed to fail, and that foundationalism as a doctrinal project has long been a dead-end.<sup>5</sup> Yet like Siegel and other non-naturalists who charge naturalists with begging the question, Rosenberg seems to presuppose that something like a rational reconstruction of knowledge ought to be possible, if not through the direct appeal to foundations, then at least through a coherentist-style

---

<sup>5</sup> And, Rosenberg (1999) has aligned himself explicitly with Quine’s project.

transcendental argument from self-refutation. If we heed Quine's original intentions, however, we should demur at all requests to offer rational reconstruction. We should forget about creative reconstructions and simply settle for psychology. We should show how, assuming the best science of the day, we explain the origins of that same science. So if we can understand the value of prediction and control using Darwinian theory (a point that Quine himself endorses), then so be it. That is a substantial achievement, even if it is not in keeping with the goals of traditional epistemology, and merely an achievement of Quine's "pessimistic" naturalized epistemology. We should be pragmatic naturalists of this manner with Quine—that is, if Quine and all of the other naturalists are correct that foundationalism and all other forms of rational reconstruction are hopeless.

If I am correct in the above, naturalists have the resources to answer the normativity objection in a variety of ways. Non-naturalists who pursue this objection usually fail to appreciate the different methodologies furnished by the many varieties of naturalism I have surveyed in my earlier section. In essence, non-naturalists fail to realize that naturalists can offer a naturalization of rationality or normativity in the same manner that they seek to offer to naturalize knowledge itself. This reduction need not comport with our pre-theoretic intuitions about evidence or justification. It need only draw on the history of science or cognitive science in a way that serves a useful scientific purpose. Depending on which naturalist we side with, the naturalist's conception of normativity can bear approximate or only minimal resemblance to our original intuitions of the ends and standards of inquiry. It can do so and avoid charges of circularity, because these charges hold up only on the presupposition that *a priori* standards can be formulated, and the entire project of naturalism has been motivated, of course, by the conviction that they cannot be. If we want to refute naturalized epistemology, we will need to challenge that conviction. It is just such a challenge that I will pose in this dissertation.

### **Outline of the dissertation**

In order to establish the thesis that we must challenge the fundamental philosophic convictions of the naturalists in order to undermine naturalized epistemology, I seek to identify those convictions

by playing naturalisms against each other to see which of them will yield to the others under pressure. In the end I will argue that “optimistic” naturalisms must give way to “pessimistic” ones of the Quinean variety. I will establish this point in the first half of my dissertation.

In our next chapter, chapter 2, I will examine Jaegwon Kim’s (1988) influential critique of Quine’s naturalized epistemology. Kim is best known for having advanced the normativity objection to Quine’s thesis, but his critique is actually more extensive than this. Kim also proposes that Quinean naturalism is not the only alternative epistemologists should consider, given the assumption that deductivist foundationalism is moribund. He proposes additional alternatives that we can now recognize as versions of optimistic naturalism, including reliabilism and his own supervenience view. In my arguments above, I alluded to the fact that these optimistic naturalisms may conflict with naturalism’s methodological proscriptions against analyticity, modality, and other intensional notions as represented by Quine’s basic theses. In my examination of Kim, I will argue in explicit detail that this conflict is real, and that pending the shouldering of a massive burden of proof, Kim and other optimistic naturalists must yield to Quine’s basic position.

In chapter 3, I apply further pressure to optimistic naturalism, and argue that one of its core concepts, the concept of “belief” (of reliable “belief” formation fame) cannot itself be naturalized by the same rigorous standards that naturalists want to apply to “knowledge” and “justification,” at least not in a way that yields a concept of “belief” usable by the naturalized epistemologist. I will demonstrate this by enumerating a variety of belief-naturalization proposals utilizing methodologies similar to the optimistic proposals enumerated above. I will argue that they all face the same methodological tensions as the naturalisms analyzed in the Kim chapter. In the end, whatever naturalization of “belief” the basic naturalistic constraints permit will longer describe the kinds of beliefs needed by epistemologists to account for the possibility of sophisticated scientific beliefs. In addition to cutting off this aspect of the optimistic naturalists’ naturalization project, I believe that focusing on the concept of “belief” will also help us understand some of the more basic

methodological shortcomings of naturalization proposals of any kind. In particular, I will return to the issue of naturalizing normativity, and show how the appeal to Darwinian theory can only yield limited satisfaction whenever applied to naturalizing philosophical concepts.

The difficulty in naturalizing a substantive concept of “belief” need not ruin the project of naturalizing epistemology, of course. In chapter 4, I will look at the attempt by pessimistic naturalists to *deflate* the concept of belief, by offering a subject-naturalization of the deployment of the *term* “belief” using a contemporary theory in philosophy of psychology, the simulation theory of Alvin Goldman and Robert Gordon. I will argue that the most naturalistic version of the theory, Gordon’s, fails to account for all of the evidence, particularly in the domain of developmental psychology. In offering this critique, I will begin to move away from applying pressure to optimistic naturalism in favor of doing the same to pessimistic naturalism. Also of interest in this critique will be the fact that evidence from developmental psychology suggests that children understand *knowledge* before they understand *belief*. Understanding this will help us to understand how the concept of “knowledge” can have theoretical significance prior to the development of adult epistemological theory, a point that helps undermine Michael Williams’s arguments against epistemological realism and case for epistemological deflationism.

I will begin the second half of my dissertation in chapter 5, by explaining in more detail the content of Quine’s pessimistic proposal, first by showing that it is in fact pessimistic. There, I show how Quine’s naturalized epistemology deals with traditional skeptical worries, not by trying to refute them in the ordinary sense, but by showing that traditional justificatory goals of epistemology must be abandoned and replaced by pragmatic ones. The resulting naturalized epistemology concerns itself with identifying the various steps (justified or otherwise) by which human subjects develop their current, instrumentally successful scientific theory. Furthermore, I argue that Quine’s pragmatic approach to epistemology faces difficulties of its own, particularly because of challenges posed by

more radical pragmatists who see no pragmatic basis for privileging natural science over other forms of human discourse.

Of course, the putative motivation for pursuing a pragmatic rather than a traditional route stems from the alleged failure of foundationalism and the inevitability of Quine's indeterminacy of translation, assuming the viability of the underdetermination thesis. In chapter 6, I argue that looking at the wider context of scientific practice—and at some specific scientific results—undermines the underdetermination thesis and draws attention to the possibility of a new foundationalism. I argue that the underdetermination thesis does not itself bear naturalistic scrutiny, in particular because of its reliance on a crude and unrealistic hypothetico-deductivist conception of confirmation. Having undermined underdetermination naturalistically, I show how to generalize this anti-skeptical strategy (itself inspired by statements from Quine about skepticism): whenever skeptics themselves assume points of science for the sake of *reductio ad absurdum*, anti-skeptics themselves have the right to make appeal to science to show how the *reductio* does not go through. I demonstrate that the classical Humean problem of induction can be partially resolved by appeal to a material theory of induction that recognizes diverse methods of confirmation practiced by scientists in different domains of fact. Finally, by appealing to psychological evidence regarding perception and concept-formation, I show how the regress of inductive justification can be terminated in perceptual foundations. At the same time, I argue that as skeptical problems become more dependent on questions about epistemological foundations, the problems become more and more philosophical and less purely scientific. This suggests that Quine is ultimately incorrect that skeptical doubts are prompted entirely by scientific problems. This means we cannot generalize his anti-skeptical strategy to solve all skeptical problems, but it also means that the naturalistic proposal to make philosophy continuous with natural science is not consistent with the fact that naturalism itself arises as a pragmatic solution to problems generated by non-naturalistic philosophic presuppositions.

The outcome of this anti-skeptical strategy robs naturalized epistemology of its *raison d'être*. At the same time, the strategy also draws attention to flaws in traditional epistemology, including its inability to articulate a workable empirical foundationalist theory of justification. These flaws must be corrected if philosophers are to preserve the autonomy of their discipline.