

Research

(Note: This research statement was written for an audience of non-philosophers as part of my tenure materials, fall 2005. Philosophers, please take loose formulations with a pinch of salt!)

The Big Picture: The Social and The Rational

Human beings are rational animals. We are also social animals. My overall goal is to understand these two truths, and their interrelations, with regard to our intellectual lives. My focus here is upon our activities of inquiring, evaluating evidence, forming beliefs, claiming to know things, and assessing our own and others' beliefs as reasonable or justified. I aim to attain what the Twentieth Century philosopher Ludwig Wittgenstein called a "perspicuous representation" – not an explanatory theory, but rather a description or interpretation which renders our intellectual lives fully intelligible. This requires correcting certain distortions in the current epistemological discussion and solving traditional philosophical problems which threaten to leave our intellectual lives unintelligible. The story I am telling is a large one and cuts across a number of philosophical subfields. My ultimate aim is to show the fruitfulness of the following broad claims: that a careful study of the structure of our actual epistemic practices is essential to understanding knowledge and justified belief; that traditional philosophical problems about the possibility of knowledge are best countered by appeal to our actual practices; and that the most basic structure of our epistemic practices can be rendered intelligible by relating it to the facts that we are dependent upon each other for much of our knowledge (as for so much else), that we have need to apply such notions as authority, responsibility, and entitlement in relation to our intellectual lives, and that in making claims to each other and offering reasons for them, we undertake commitments for which we are appropriately held accountable.

This statement will provide an overview of my thinking, highlighting the place of my publications and work in progress within it.

Philosophical Skepticism and The Coherence of our Epistemic Practices

Do I know that there is a table here before me? Common sense declares that I do. But philosophical reflection can quickly lead to the conclusion that common sense is incorrect: no one can know or even reasonably believe anything about the world around us. This is the problem of *philosophical skepticism*.

The arguments which generate this problem are not silly or easily refuted. They work by purporting to show that our epistemic practices – our ordinary ways of justifying and evaluating claims about the world – commit us to requirements that are incoherent or unsatisfiable. Consequently, these arguments do not merely question the adequacy of this or that philosophical theory of knowledge. They challenge the most basic aspiration of our intellectual lives: the attainment of reasoned and defensible beliefs about an objective, mind-independent world.

I aim to put skepticism to rest through careful reflection on how our epistemic practices actually work. This is a neglected approach. Philosophers have traditionally responded with elaborate theories designed to safeguard the possibility of knowledge and reasonable belief. Descartes, for instance, attempted to prove God's existence and benevolence in order to underwrite our entitlement to rely on certain beliefs (our so-called "clear and distinct perceptions"). More recently, philosophers have appealed to

sophisticated a priori theories about the nature of perceptual experience, the nature of linguistic meaning and the way in which our beliefs come to be about the world, and the conditions under which beliefs count as reasonable or constitute knowledge. Such approaches share the aspiration to stand outside our ordinary lives and provide them with an external, philosophical underpinning: a “first philosophy”. Eschewing such aspirations, philosophers such as Richard Rorty propose simply to turn their backs on the problem. But a mid-Twentieth Century tradition including G.E. Moore, J. L. Austin, the later Wittgenstein, and (in a somewhat different way) W. V. O. Quine has proposed another route: to overcome philosophical skepticism from *within* our ordinary epistemic practices. This approach has been neglected in the current discussion. One aim of my work is to put it back on the mainstream agenda.

Here is my approach. I confront the arguments for skepticism in exactly the way we would confront any other surprising claim, by asking “Why — if at all — should I believe this?” I start from where we are, with our convictions about what is the case and what counts as a good reason for believing what, our ordinary ways of evaluating objections and deciding whether someone has knowledge. I then ask, “Can I reasonably be brought to think that because our epistemic practice is incoherent or involves unsatisfiable demands, I cannot know or even reasonably believe that there is a table before me now?” I believe that the answer is, “No. You can’t get there from here.” What’s more, I believe that we cannot reasonably generate the demand for independent validation of our practices from within the framework of our ordinary epistemic lives. By exploring the extent to which these claims are correct I hope to attain insight into the relationship between traditional epistemology and our ordinary lives, an issue which should trouble philosophers more than it generally does.

Here are some examples of how I have developed this approach.

It is tempting to think that knowledge requires infallibility — evidence which “rules out” all possibilities of error. (If you recognize you could be in error, the thought goes, how can you also claim to have knowledge?) This thought gives rise to one standard argument for skepticism: we can never attain infallible evidence about the world around us, so we can’t ever have knowledge of it. But we do not insist upon completely infallible evidence in ordinary life or in science. It is often suggested that this fact does not defeat the argument. Our ordinary ways of talking about knowledge are misleading, it is said; they are shaped by the exigencies of everyday life — considerations of practicality and the demands of communication and cooperation — and so it is only when philosophical reflection prompts us to prescind from these concerns that we recognize that knowledge requires infallibility. My paper “Is Fallibility an Epistemological Shortcoming?” (2004a) takes on this view. As I argue, it misdescribes our practices of knowledge attribution: our practices are not shaped by practical considerations in the suggested way. We are consequently free to insist, in accordance with our ordinary practices, that the fact that our evidence is fallible does not prevent us from having knowledge of the world. A compelling skeptical argument must deploy requirements to which we are committed by our ordinary practices; it cannot simply abandon those practices altogether. This paper inspired a reply by Stephen Hetherington. His paper, along with my reply, is forthcoming.

An ancient argument, known as the “problem of the regress”, purports to show that it is impossible for anyone ever to be justified in believing anything. The thought is

that in order to be justified, you must be able to provide a good reason for your belief, a good reason for believing that reason, and so on and on. But then it seems that you can never justify any belief. To refuse to provide reasons at any point would be arbitrary and dogmatic. To repeat yourself is to argue in a circle. And even if you could go on infinitely offering reasons for your beliefs, it is hard to see how this could, by itself, constitute an adequate defense. My paper “A Localist Solution to the Regress of Justification” (2005c) develops a new response to this puzzle. It draws on our actual practices to show that no vicious regress results from accepting (1) that in order to be justified, you have to be able to provide a good reason for your belief, and (2) that you must be justified in believing whatever you might offer as your reason. To be justified, I propose, is to possess a certain *ability*: the ability to draw upon your background beliefs to provide what are in fact good reasons for holding the belief in question, ultimately by providing reasons which you correctly take there to be no reason to doubt. This is an ability which we can possess.

Showing the coherence and satisfiability of the demands of our epistemic practices can seem unsatisfactory in certain ways. In particular, one might also want an independent or external validation of our practices, an argument which shows how and why they enable us to arrive at knowledge or justified beliefs about the world. This demand can take several forms. For instance, traditional “Cartesian” inquiries attempt to explain our knowledge of the world without presupposing any claims about the world. (This is the project Descartes initiated in his *Meditations*.) It is arguable that once you take up this standpoint, it becomes impossible to see how we could have knowledge of the world. My paper “Epistemological Externalism and the Project of Traditional Epistemology” (Forthcoming 1) shows that the problem lies with the project. More particularly, I show that if, as I believe is the case, what’s a good reason for holding some belief about the world always depends in part upon facts in the world, then there is no compelling intellectual motivation for the traditional Cartesian project: it cannot be completed, and we can have knowledge and justified beliefs about the world even if we cannot explain, from the Cartesian’s artificial standpoint, how this is possible. This argument opens up several lines of inquiry regarding the nature of reasons: What exactly is a reason for belief? Are there any necessary truths regarding reasons for belief (aside from the laws of logic)? What (if anything) can be known independently of experience about the goodness of epistemic reasons? These questions are at the forefront of my current thinking. (I treat some of them in “Epistemic Instrumentalism and Reasons for Belief” (forthcoming 2). I will also address some of them in a graduate seminar on reasons for belief and action which I am co-teaching with my colleague Kevin Toh this fall.)

Taken together, the above papers form the core of my argument so far that “you can’t get there from here.” However, an important skeptical argument remains. It goes like this. Consider the possibility that you are just dreaming that there is a piece of paper before you (or, in keeping with our high-tech age, that you are being subjected to deceptive sensory stimulations from a sophisticated computer apparatus). It seems that you can’t know or reasonably believe that there is a piece of paper before you unless you know or have reason to believe that these possibilities are not the case. But it also seems, on reflection, that you can’t know or have any reason to believe that they aren’t. Any evidence you might appeal to could itself be the product of a dream or of the deceptive

stimulations of the computer. So it seems that there is nothing you could appeal to in support of the belief that you aren't being deceived in one of these ways. It consequently seems that you can't know or reasonably believe anything about the world around you. This argument has deep appeal. It has appeared in a variety of cultures over several millennia, and its continuing allure is shown by the immense attraction of contemporary cultural products such as the *Matrix* movies. At the same time, it seems that the conclusion is clearly incorrect. What has gone wrong?

One common response is to deny that in order to know there is a piece of paper before you, you must know or have reason to believe that you are not dreaming. This response strikes me as incredible. How can I know that there is a piece of paper before me if I don't even have reason to believe that I am not just asleep and dreaming? A satisfying, commonsensical response to the Dream Argument should grant this requirement. Moreover, it should show that this requirement is satisfiable. For one's natural response is to say, "Of course I have reason to believe that I am not asleep and dreaming; in fact, I know that I am not." In "An Empirical Refutation of Skepticism" (in progress) I argue that the claim that we can't have reason to believe we aren't asleep and dreaming depends upon hidden empirical premises – premises about the world – which we have good reason to reject. Turning this insight into a satisfying response to the argument requires showing that the appeal to empirical claims is not utterly question-begging in this context. This in turn requires taking on several large topics, including the nature of reasons for doubt, the role which requirements such as the skeptic's play in our actual epistemic practices, the requirements one must meet in order to know something, and the nature of perceptual experience and its role in the justification of our beliefs about the world. (Some of these issues are treated in my paper "Skepticism, Closure, and Sensitivity, or Why the Closure Principle is Irrelevant to Skepticism" (2004c)).

In the fall of 2004, I had a College Arts and Humanities Institute fellowship to develop my research on these topics. I focused particularly upon the nature of knowledge, developing and extending lines of thought initially broached in "Is Fallibility an Epistemological Shortcoming?" One prominent view in the current literature is that *practical circumstances* determine the requirements one must meet in order to count as knowing something. This idea is sometimes elaborated as a semantic theory about the meaning of sentences using the word "knows", sometimes as a view about the context-relativity of knowledge itself. Either way, I think the view is mistaken, and I have been developing an account which shows how merely fallible evidence can be sufficient for knowledge even though the requirements for knowledge are not responsive to the practical context. This project requires dealing with a number of difficult technical issues about the relation between knowledge and practical reasoning. Current work in progress ("In Defense of Fallibilism", "Knowledge and Decision") develops solutions to these problems. It has recently been argued by a prominent epistemologist that our practices of assertion entail that the standards one must meet in order to have knowledge vary with the practical context. In "How to Link Assertion and Knowledge Without Going Contextualist" (forthcoming 3) I show that this is not so. In "Worries for Would-be WAMmers" (2005b), I show that certain anti-contextualist ways of treating the linguistic data will not work; anti-contextualists should acknowledge that our ordinary practices of claiming and ascribing knowledge are not responsive to purely practical circumstances. What matters when we think about whether people know things is just how good their

evidence is. My work on these issues has recently been acknowledged by an invitation to contribute the article on Fallibilism for the second edition of Blackwell's *Companion to Epistemology*.

The Life of the Mind

An easy way to avert many traditional epistemological problems is to claim that one can be justified in holding a belief even if one can't come up with any reasons at all in defense of it. For a variety of reasons, this view is currently common in epistemology. In fact, many of the most prominent contemporary theorists hold that being justified *never* requires the ability to offer good reasons for one's belief. On this view our activities of publicly justifying our beliefs look like a secondary, inessential add-on, so not only do the traditional problems disappear, but it can also look perverse to take our epistemic practices to have any epistemological significance at all. I believe that this view does too much violence to our ordinary understanding of the significance of the activities of asking for, providing, and evaluating reasons. In "On Justifying and Being Justified" (2004b), I argue against the dominant view on several grounds. Most significantly, the view makes it incomprehensible that we evaluate people on the basis of the reasons for which they hold their beliefs, subjecting them to various obligations and responsibilities in this regard. The fundamental issue here concerns the very idea of a person's *holding a belief for a reason*; as I argue, the dominant view can't make good sense of this idea. I propose an alternative account which places the person's commitments, and hence her justificatory abilities, at center stage.

This account opens up several lines of further investigation, many of which forge connections between my work in epistemology and important contemporary philosophical discussions regarding freedom of the will, the relation between commonsense psychological explanations and scientific explanations, the nature of our knowledge of our own minds, and our responsibility for our actions and beliefs. The over-arching issue here concerns the relation between, on the one hand, our first-person perspective on ourselves when we deliberate about what to believe, and, on the other hand, theoretical perspectives from which our beliefs look like things which "just happen" to or in us. I plan to develop my thinking in these areas as I complete my current work on knowledge. My focus here will be upon explanations which appeal to reasons and the question of what it is to hold a belief upon the basis of a reason. (Some of this work will appear in a paper currently in progress, "Responsibility and Responsiveness to Reasons," which will focus upon the relation between the fact that we think it right to regard ourselves as responsible for what we believe and the fact that we can directly determine our beliefs — that is, bring it about that we believe certain things and not others — by weighing evidence and evaluating reasons.) The pay-off of this work will be insight into the place of notions like *justification* and *believing for a reason* in everyday psychological explanations and in our lives as social and evaluating beings.

Future Research: Social Dimensions of Knowledge

Much of our mundane, everyday knowledge depends upon our acceptance of someone else's (or some agency's or organization's) word as to how things are. (For instance, just consider your knowledge of such matters as your name and birth date, the state capitols, and current events around the world.) The same point applies to much of

our scientifically-based knowledge. How many of us have checked for ourselves whether high cholesterol increases the risk of heart attack, what the atomic weight of oxygen is, or how many chromosomes are found in human liver cells? The process by which knowledge can be transmitted is almost absurdly simple: if you know something, tell it to me, and I believe you, then I know it too. But how is this possible? How can I acquire *knowledge* — not just come to believe something — merely by believing what you tell me? Second-hand knowledge, or "knowledge by testimony" as it's often called, has been subjected to intense scrutiny by philosophers and social scientists over the past decade or so. However, it has not been sufficiently recognized that the topic cannot be treated satisfactorily within the framework of standard epistemological theorizing.

Epistemology in the Western tradition has been resolutely individualistic. Theorizing has been guided largely by the image of the lone thinker struggling to wrest the truth about the world from his sensory experiences and the insights of reason. This image has been preserved in recent philosophical discussions of second-hand knowledge. It has been widely assumed that a primary question in this area concerns the nature of one person's justification for believing what another person says: does this justification rest upon merely empirical factors (such as one's evidence about the likelihood that the speaker is lying or misinformed), or does it involve an *a priori* element (that is, is there some non-experiential reason why we are entitled to believe what other people tell us)? As a small number of epistemologists are beginning to realize, formulating the issue in this way maintains an underlying individualistic conception. It assumes that for the listener, the speaker always functions as something like an instrument or indicating device, a mere source of evidence like a thermometer or gas gauge. The listener's task is then to decide whether to regard the speaker as a reliable instrument regarding the matter at hand. If the listener does believe what the speaker says, then — on this view — his justification for doing so would always look something like this: "S said that *p* and S is most likely right about *p*, so *p* is probably true." This way of interpreting the relationship between speaker and listener still treats the listener as a lone thinker; all that is added is that the words of other people are an additional source of evidence for the lone thinker to consider. If the speaker was in fact misguided or lying, then the ultimate responsibility for the listener's incorrect belief would rest with no one but the listener himself. To put it bluntly, on this view there is never any significant epistemological difference between relying on a person's word and relying on a thermometer: people might just as well be machines.

In an important class of cases the relationship between speaker and hearer is not at all like this. The philosopher J. L. Austin once noted that when a speaker says "I know ..." to someone else, she thereby institutes a unique relationship between herself and the listener: she thereby gives the listener her guarantee of the truth of what she is saying. If it turns out that the speaker is wrong, then it may be appropriate for the listener to charge *the speaker* with responsibility for the listener's incorrect belief. Thus, in such an exchange the speaker's words do not merely function as evidence for the listener. The speaker takes on responsibility for the truth of the listener's belief by offering him a guarantee of its truth, and if the listener accepts this guarantee, his acquisition of knowledge would seem to depend upon his having entered into a relationship structured by certain interpersonal obligations and entitlements. This point applies more generally. In an important class of cases, a listener's relationship to a speaker is not like his

relationship to an instrument or source of evidence. If your thermometer gives a faulty reading, it makes no sense (anthropomorphism aside) to *blame* it for the incorrectness of your belief. But it does make sense to blame a speaker who leads you astray.

Austin's insight has been largely neglected in the epistemological literature, with the exception of several recent important papers by Dick Moran. In unpublished work in progress (growing out of a graduate seminar taught in 2001), I attempt to establish the *epistemological* relevance of the relationship between speaker and hearer. This requires showing that the speaker's guarantee, the hearer's acceptance of that guarantee, and the possibility of the hearer's appropriately holding the speaker responsible for the hearer's incorrect belief, are not merely a moral or quasi-moral overlay upon a more basic epistemological structure, but are rather a crucial part of what makes the hearer justified in believing as he or she does. They enable the hearer to depend upon, or "piggyback" upon, the speaker's justification for her belief. However, it seems that the hearer will rarely have sufficient independent evidence that the speaker is sincere or that her claim is likely to be true, so it seems that the hearer's acceptance of the speaker's guarantee is a decision to *trust* the speaker. If this is right, then much of our knowledge depends in a fundamental way upon a social web of relations of trust and responsibility, not upon evidence in our own possession. This implies a deeper connection than is often recognized between moral philosophy and epistemology.

My work in this area to date has focused on the structure of the relationship of trust and responsibility between speaker and hearer and the role it plays in the transmission of knowledge. However, I suspect that the obligations and entitlements which structure the relationship of trust and responsibility between listener and speaker provide insight into the requirements which one must meet in order to have knowledge or justified beliefs. The idea could roughly be put as follows. If I know something, then I am entitled to tell you how things are. But my invitation of your trust is open to criticism if I am wrong about how things are or if I lack a good explanation of why you should trust me. So if my invitation of your trust is to be beyond criticism, then I must be *worthy of your trust* — my belief must be true and I must be able to back my invitation with good reasons. This requirement yields important results concerning the nature of knowledge, the relation between knowledge and justification, and the extent to which being justified requires the ability to articulate an adequate defense of one's beliefs — all topics of hot debate in the current epistemological literature. In particular, since knowledge (at least the kind of knowledge which is transmitted through conversation between mature human beings) involves the authority to tell other people how things are, it requires justified belief, which in turn requires the ability to articulate reasons in defense of one's beliefs. I therefore suspect that reflection upon the transmission of knowledge through speech will provide a rich framework for understanding our epistemic lives more generally. Whether this suspicion is correct will be a central question in my future work.

Publications while at Indiana University:

(Forthcoming 1). "Epistemological Externalism and the Project of Traditional Epistemology," *Philosophy and Phenomenological Research* (main text 9,491 words (31 pages), including notes 14,481 words (52 pages)).

- (Forthcoming 2). “Epistemic Instrumentalism and Reasons for Belief: a reply to Tom Kelly,” *Philosophy and Phenomenological Research* (main text 3,589 words (12 pages)).
- (Forthcoming 3). “How to Link Assertion and Knowledge Without Going Contextualist: A reply to DeRose’s ‘Assertion, Knowledge, and Context’,” *Philosophical Studies* (main text 5,000 words (19 pages)).
- (2005a). “On Williamson’s Arguments that Knowledge is a Mental State,” *Ratio (new series)*, XVIII:2, 165 - 75.
- (2005b). “Some Worries for Would-be WAMMers,” *Grazer Philosophische Studien*, 69, special issue on contextualism, 101 - 25.
- (2005c). “A Localist Solution to the Regress of Epistemic Justification,” *Australasian Journal of Philosophy*, 83:3, 395 - 421.
- (2004a). “Is Fallibility an Epistemological Shortcoming?” *The Philosophical Quarterly*, 54: 215, 232 – 251.
- (2004b). “On Justifying and Being Justified,” *Philosophical Issues* (a supplement to *Nous*), 14, 219 – 253.
- (2004c). “Skepticism, Closure, and Sensitivity (or Why the Closure Principle is Irrelevant to Skepticism),” *The Croatian Journal of Philosophy*, special issue on Robert Nozick, IV: 12, 335 - 350.