

COMMON GROUND AND DISCURSIVE JUSTIFICATION:
Approaching the Traditional Epistemological
Questions from an Untraditional Angle

BY
RYAN SIMONELLI

A Thesis

Submitted to the Division of Humanities
New College of Florida
in partial fulfillment of the requirements for the degree
Bachelor of Arts
Under the sponsorship of Aron Edidin

Sarasota, Florida
January, 2014

Acknowledgments

This thesis would not have been possible without Aron Edidin and all the time and effort he put in over the past few years. In addition, I would like to thank the rest of my thesis committee, Douglas Langston and April Flakne. I also owe my gratitude to Rebecca Kukla for giving very helpful feedback on earlier drafts as well as Edward Kanterian and Bob Pasnau for their feedback. Finally, I'd like to thank my family and friends for their continued support through this whole process

Table of Contents

Acknowledgments	ii
Table of Contents	iii
Abstract	iv
Chapter One: The Quest for Common Ground	1
Chapter Two: The Discursive Community	13
Chapter Three: How to Tell the Skeptic to Get Lost	30
Chapter Four: The Seesaw Model of Discursive Justification	55
Chapter Five: From Discursive Justification to Justified Belief and Knowledge	74
Works Cited	94

COMMON GROUND AND DISCURSIVE JUSTIFICATION:
Approaching the Traditional Epistemological Questions from an Untraditional Angle

Ryan Simonelli

New College of Florida, 2014

ABSTRACT

One of the classic issues in the field of epistemology is the question of how we are justified in holding the beliefs that we do in fact hold. Traditionally, epistemologists, drawing from the tradition of Descartes and Locke, have thought that this question is to be answered purely by internal reflection. In this thesis, I argue for a conception of justification that takes its explanatory starting point on the social practices surrounding the *act of discursive justification*, turning the traditional individualist conception of justification on its head. This often been seen as a mere changing of the subject from the traditional epistemological concerns. However, I argue that once we dispose of the challenge of global epistemological skepticism, we can successfully address these very same concerns by thinking about the question in terms of our social practices of justification.

Sponsor Approval:

Aron Edidin

Division of Humanities

Chapter One: The Quest for Common Ground

This thesis is about central issues in epistemology. Throughout much of Western Philosophy's history, epistemology was at the very heart of the discipline. However, in the past half a century or so, traditional epistemology, true to its Cartesian, Lockean, and Kantian roots, has faced some serious opposition. Perhaps the harshest critic of traditional epistemology was Richard Rorty who objected that epistemology so conceived must have a foundationalist, infallibilist basis—a basis which he took, drawing on arguments from Quine and Sellars, to be impossible. Rorty's concluding thought is that we should throw out the traditional epistemological project if we are going to adopt his broadly pragmatist picture. One of the results of this thought is that the epistemological work of the "new pragmatists" such as Donald Davidson and Robert Brandom, who generally agree with the Rortyan picture, has not been regarded as presenting possible answers to the traditional epistemological questions. This, I think, is a mistake. In this thesis I will attempt to tie together some of epistemological work of the new pragmatists, primarily Brandom and Davidson, and attempt to tell a convincing story about how what we believe must really be largely correct, how our justifications really are tied to the world, and how we are really justified in most of our beliefs. The goal of the project, as a whole, is to draw an epistemologically satisfactory picture of justification that draws from the practices of giving and asking for reasons as they're exhibited in actual discursive practices.

Although the tradition I am primarily working out of has included some of the strongest criticisms of traditional epistemology, I take myself to be attempting to answer the very same questions that concerned the traditional epistemologists. I will come at the questions from a different angle, but I am by no means trying to avoid them. As such, my

project is not an abandonment of the traditional epistemological concerns (as is often said of the work of the “pragmatists” just mentioned) but aims to answer the concerns in a way that would genuinely satisfy the traditional epistemologists. I will argue that answering these questions in a genuine and satisfactory way gives us a new platform to think about contemporary epistemological questions such as the analysis of justification and knowledge and how epistemological inquiries fit into a naturalist picture that have become detached from the traditional ones.

1. What Should the Goal of Epistemology Be?

The best characterization of how an epistemology true to the traditional aspiration ought to conceive of itself might have been given by Richard Rorty, one of its biggest opponents. Eulogizing what he takes to be the dead subject of epistemology, recollecting the failure of past attempts to answer the traditional questions, Rorty attempts to capture its essence: “To construct an epistemology is to find the maximum common ground with others. The assumption that an epistemology can be constructed is the assumption that such common ground exists.”¹ By “common ground” here, Rorty is speaking of an *epistemic* common ground, a “common ground for adjudicating knowledge claims.”² Drawing from the work of Kuhn and Feyerabend, he links this sort of common ground with *commensurability*. One of the criteria for two claims to be commensurable is that there must be a *principled way of coming to agreement*. “Common ground” is the set of beliefs and standards for forming beliefs that makes this principled agreement possible. Without such common ground, there seems to be an impenetrable gulf between two radically different claims arising from radically different ways of thinking about the world. To engage in the traditional task

¹ Richard Rorty, *Philosophy and the Mirror of Nature*, 318

² *Philosophy and the Mirror of Nature*, 317

of epistemology then, is to seek to demonstrate the maximum epistemic common ground that all epistemic agents share.

Rorty tells us, however, that an epistemology that aims to find the maximum common grounds among any agents capable of making truth claims only can function on the “assumption that such common ground exists,” and I think he’s right. If the traditional epistemological is to be possible, it must *already be the case* that we share a basic agreement both about the way the world is and about basic belief-forming norms with every epistemic agent with whom we might engage in discourse. Without a shared understanding of the way the world is, we have nothing to work from. We cannot justify an inference as an attempt to expand common ground if there is nothing mutually agreed upon from which such an inference could be produced. Likewise, without shared norms of belief-forming norms, we have nowhere to go. As Lewis Carroll observed in his charming “What the Tortoise Said to Achilles,” if Modus Ponens is not followed as a belief-forming *practice*, then, no matter which and how many propositions are accepted, no common ground can be gained from inference.³

So this common ground is what we must first look for if we want to get an epistemological enterprise up and running. Where do we look? Descartes sat in a room and engaged in the process of critical reflection to arrive at absolute certainty. If we can provide a ground for belief and inference that is absolutely logically indubitable, so went the thought, then of course it must be shared by everyone. Most epistemologists agree that Descartes’ project, in its specific form, is philosophically bankrupt, and I agree. In fact, I think Descartes had it backwards. Rather than looking inwards towards our own reasoning, we should be looking *outward* to the beliefs and belief-forming practices of epistemic agents

³ Lewis Carroll, “What the Tortoise Said to Achilles”

engaging in discursive practices in the world. If we're looking for common ground, we should look for it in actual discursive practice, where various sorts of beliefs and epistemic norms will be at play. What we need to do, if we want to answer the question of common ground, is to map out the space of epistemic norms and see how they actually function. Examining the pragmatics of our epistemic practices in the right sort of way, I argue, can illuminate the answers to the questions of how our beliefs are justified, the very same ones that concerned Descartes.

To embark on our investigation, we first need to clarify what we are *not* looking for. There is not going to be some set of explicitly articulated epistemic principles that all rational agents adopt. Most people aren't prepared to offer a very good answer to the question, "On what general principles do you form the beliefs you have?" Of course, some might say "logic, rationality, etc." Others might add in something like faith or justice. But we are wrong if we think we can look to any of these explicitly articulated principles for believing and find a common ground. The way we articulate our practices are far too different, and for all we know we might be wrong about the belief-forming principles we follow, or not even know what it would mean to follow the principle we say we follow. What we are looking for are epistemic norms as *embedded in the practices* of forming beliefs, granting beliefs as justified, and giving and asking for justification. Only here can we hope to find universally shared norms. I take my starting point from the account of discursive practices offered in Robert Brandom's landmark book, *Making It Explicit*. The central task of his work is to offer an account of meaning as derived from social practices and the normative attitudes implicit in these practices. I will not concern myself here with his overarching semantic agenda, but rather the subset of his project which consists of his cataloguing of the normative

framework of our epistemic practices and the discursive practices necessary to get our epistemic practices up and running in the first place.

If Brandom is roughly on the right track, then there must be some set of normative standards shared by all agents capable of discourse. Brandom's project is not simply to describe some set of *localized* norms that some discursive agents share, but also to describe the underlying set of norms that make discourse possible *at all*. Among this latter set of norms are *epistemic* ones, norms relating to our practices of acquiring knowledge ourselves, attributing knowledge to others, challenging the claims of others and justifying our own in the face of challenges from others. A central commitment of my project here is that there is a core set of universal epistemic norms. This is not to say that there won't be many localized differences in epistemic norms among different communities at different times and places. There almost certainly will be such differences, but the claim is that any divergence of epistemic practices can only occur against a backdrop of shared epistemic norms since we can only count a community as having genuinely *epistemic* practices if they all share a fundamental normative structure by virtue of which it makes sense to call them the epistemic practices at all. This shares a certain sort of transcendental element to epistemological inquiry with a traditionally Kantian outlook. However, once again, it is important to note that, rather than looking *inward* to experience we are looking *outward* to public social practices and identifying the norms imbedded in them.

2. Global and Local Justification

It is identifying this universal set of epistemic norms and working out their implications that I think can provide the answers to the fundamental epistemological questions. The traditional epistemological project, however, isn't simply a *descriptive* one, but a *normative* one, since knowledge and justification are themselves normative notions. Thus

my approach requires bridging the gap between *identifying* our actual epistemic norms and normatively *endorsing* them. It is not enough that these norms be universally shared in order to demonstrate the sort of “common ground” that traditional epistemology seeks to show; they must also be *correct* and lead to *true* beliefs in standard cases. It is for this reason that a response to the skeptical challenge is at the heart of this thesis. The classical versions of skepticism attempt to pose a challenge to our most widely held beliefs and basic belief-forming processes. If it is possible, as the skeptic claims, that there could be a massive divergence between our beliefs about the world and the world itself, then answering the descriptive question about how our epistemic practices actually go could not provide the answer to the normative question of how we *ought to go*. If on the other hand we can show contra the skeptic’s challenge that, in the standard cases, the central epistemic norms that we have systematically lead to true beliefs, then the descriptive project of cataloguing our central epistemic norms can provide answers to these normative questions.

Giving an answer to the skeptic is not only a main task of my project here but will also be contextualized within the project as a particular sort of justification in response to a global challenge to our beliefs. The practice of justification is the main epistemic practice that will be fleshed out in detail. Justification as a *practice* rather than a property attributed to beliefs is what I will call *discursive justification*, and it is a practice we engage in rather frequently. The aim of conversation isn’t always to come to an agreement on the way things really are and expand our shared understanding of the world. In fact, this sort of conversation is actually probably in the minority. We make small talk, we express mutual excitement about some thing or another, we request, we offer support and comfort, we flirt, and so on and so forth. In some sense, all of these activities have the upshot of expanding common ground of a related sort to what I am talking about here, but, centrally for the

present project, there are also conversations aimed at *settling disagreement*. I might ask a friend of mine, “You *really* believe in Bigfoot? How so?” Even though this is a friendly challenge, it is still a challenge to my friend’s belief. And if this is an appropriate context for such a challenge, there are some normative standards by which some sort of justificatory response ought to be provided. When we justify a claim there are various sorts of strategies that we employ. We might show how an inference can be made, to that claim, from some other claim to which we take the challenger to be committed (“See these pictures? There’s no way they could have been photo-shopped! Therefore, Bigfoot *must* be real!”). We might appeal to our own authority (“I’m the president of the Skeptic Society. Now, *of course*, there’s good reason to believe in Bigfoot if *I* believe in him) or someone else’s. We might appeal to direct perception (“When I was in Oregon, I saw bigfoot with my *own two eyes!* I know what Bigfoot looks like, and I saw him *clear as day!*”). Or we might attempt one of the many other possible justificatory moves which are appropriate in various contexts.

Brandom is not centrally focused on justification in *Making It Explicit*, but he does locate justification within larger discursive practices. Justification (as the sort of discursive justification just mentioned) is a move we can make in the game of giving and asking for reasons when a claim of ours is challenged. It aims at gaining *entitlement* for a claim, a sort of social *license* that makes holding that claim acceptable.⁴ On Brandom’s account, justification functions in the context of a *default and challenge* structure of entitlement. This will be explained in detail in the next chapter, but the basic notion is that the majority of claims are entitled by *default*, and claims requiring justification in the face of a challenge are *exceptions* to the standard case. While some might take this pragmatic structure as reason to dismiss the skeptical challenge, since according to the pragmatics of our everyday epistemic practices we

⁴ Brandom, *Making it Explicit*, 159.

don't *need* to give positive justification for the majority of our beliefs, I maintain that answering the skeptical challenge is first necessary for any reference to the actual structure of our epistemic practices to have any normative force. Thus, in the account of discursive justification I offer here, a *global justification* is a necessary component to an epistemologically satisfactory account of justification even with respect to local issues.

Ultimately, this allows me to offer a proposal of justification for belief that differs radically from traditional foundationalist or coherentist accounts. Responding to the skeptical challenge allows us to treat the actual practice of discursive justification as central in explaining the property of being justified. On the account of justification I am proposing here, discursive justification is explanatorily prior to justification as a property attributed to beliefs (the sort of justification that has traditionally concerned epistemologists). I propose that a belief is justified if and only if it can be *discursively* justified against appropriate challenges of which the believer ought to be aware. This entails thinking of some specific set of criteria that a belief must fulfill (be it foundationalist, coherentist, or something else) as having only secondary function since articulating one set of criteria or another may aid in giving a discursive justification in various circumstance. Thinking of justification first and foremost this way allows one to avoid many of the pitfalls that accounts of justification have traditionally succumbed to. For one, the epistemic regress problem that has often been taken as a reason to endorse a foundationalist conception of justification is a complete non-issue on the sort of account I offer. Since the practice of discursive justification functions within the structure of default and challenge in which the need to offer a justification for some belief is an *exception* to the standard case of default entitlement, an account of justification for beliefs which takes its basis on this practice sees the epistemic regress problem as simply confused about the nature of justification and thus not posing a real threat at all.

Of course, I don't attempt to identify *all* the discursive norms at play when we offer justifications. Such a task is far beyond my current resources or goal here. I also am not denying that there are some localized norms in the practice of justification that lead us astray or systematically get it wrong. For fallible agents, such localized pockets of epistemic failure are inevitable. My only commitment here is to the general structure of the practice, which I will identify here, and its positive epistemic status. Cashing out the details of particular justificatory practices, and sorting out the epistemically good bits of localized practices from the epistemically bad ones is a different project.

3. A New Path for Naturalized Epistemology

If we are able to draw epistemological conclusions from looking at the pragmatics of our core epistemic practices, a new path is opened for "Naturalized Epistemology," a way of thinking about answering epistemological questions proposed by Quine. According to Quine, once we give up the goal of epistemology as an a priori discipline to *ground* a posteriori natural science, there is nothing stopping us from calling on the empirical sciences in answering the traditional epistemological questions. On the traditional picture, where epistemology aimed to give a foundation for the empirical sciences, this move would seem viciously circular. However, once we have given up the foundational quest, the issue switches to a purely descriptive one. Quine writes, "The stimulation of his sensory receptors is all the evidence anybody has had to go on, ultimately, in arriving at his picture of the world. Why not just see how this construction really proceeds? Why not settle for psychology?"⁵ This suggestion was not greeted well by most epistemologists. In a popular discussion, Jaegwon Kim tells us that Quine wants to replace the normative question of justification with the nonnormative causal question of how we come to form beliefs and in

⁵ Quine, "Epistemology Naturalized" in *Ontological Relativity and Other Essays*. 83

doing this, justification, being an essentially normative notion, drops out of epistemology and knowledge along with it.⁶ This is to say that Quine’s “naturalized epistemology” cannot be seen as a “replacement” or “successor subject” for epistemology since the two disciplines are not even investigating the same thing, and I think this is a correct observation. Quine’s specific suggestion for epistemology is too strong if we want to retain anything worthy of the title “epistemology.”

The strategy I offer here, if fruitful, gives us a way to stay true to Quine’s urge to connect epistemology to the empirical sciences while avoiding Kim’s criticism. As I’ve said, I’m going to argue there that our actual epistemic norms must be by and large correct. If framed in this context, then the study of the epistemic norms embedded in our actual discursive practices can have a normative epistemological upshot. I do think the study of the norms embedded in our social practices is a broadly empirical one, and thus a study that can be undertaken by natural science. Even if the conceptual resources are not currently in place to make this study sufficiently systematic and rigorous, I am optimistic about the prospect.

Among contemporary epistemologists who adopt a more or less Quinian outlook of naturalized epistemology there is often little to no worry about the traditional concerns regarding skepticism. A response to skepticism, I maintain, is necessary if these naturalistic attempts are taken to be an extension or replacement of the traditional tasks and not a mere changing of the subject. If naturalized epistemologists want to claim that they are really doing the sort of epistemology that has concerned philosophers for the past few hundred years and not merely the psychology and sociology of knowledge, a different subject entirely, then answering the traditional questions head-on is in order. If these are questions that we

⁶ Kim, “What Is Naturalized Epistemology?” in *Philosophical Perspectives*.

ought to be able to discard, then we ought to be able to explain to someone who thinks they are particularly pressing *why* they can be discarded. Only once this is done, can the naturalized attempts to draw conclusions about knowledge and justification from empirical investigation be satisfactory to one concerned with traditional questions such as skepticism. Of course, it certainly need not be the task of everyone doing naturalized epistemology to respond to skepticism, and these efforts of naturalized epistemology are productive in themselves, but if someone *does* take on the task of responding directly to the skeptic with the aim of showing how we can draw normative conclusions about how we are justified in our beliefs from an empirical, descriptive investigation it gives a new sort of authority to these naturalized accounts of knowledge and justification as having genuine epistemological upshot.

Thinking about belief and justification in a Brandomian context in which such things are essentially discursive phenomena attributable in the full-fledged sense only to language-users like ourselves is a very different version of naturalism than the sort that is employed by many philosophers working in naturalized epistemology such as Hilary Kornblith and Ruth Millikan who draw heavily from the biological sciences and treat the knowledge of nonlinguistic animals as the paradigm case of knowledge from which we make conclusions about the sort of knowledge that we have. Drawing from Brandom, I work in the exact opposite direction. Still, I think both approaches are productive and in fact necessary if we are going to get a detailed and thorough naturalized epistemological picture. The distinction between favoring one of the two approaches, I think, is favoring either causal or conceptual priority. With Brandom, I am committed to the claim that understanding our own practices as linguistic discursive beings is *conceptually* prior to treating nonlinguistic animals as having beliefs and knowledge, and yet, in telling a *causal* story about how we come to have belief and

knowledge in the discursive sense that we have these things, we must start by attributing belief and knowledge to nonlinguistic animals. It is not my main task here to explain how these two pieces fit together, but it is important to note that I see both approaches as having an essential place in explaining how knowledge functions and comes about in the natural world.

4. The Upcoming Chapters

In the next chapter I will provide a sketch of the Brandomian framework of discursive practice and recent additions to this framework by Rebecca Kukla and Mark Lance, highlighting the practices that must be necessary for the epistemic common ground I am aiming for. I will outline the default and challenge structure of entitlement within which this whole discussion operates, as well as our practices of observing objects and events around us, recognizing each other as discursive agents capable of giving and asking for reasons, and the basic inferential structure in which reasoning operates. After this broad sketch of our discursive and epistemic practices I will turn my eye toward the skeptic in chapter three, offering a global justification of the majority of our beliefs in response to the skeptic's challenge. This will have the upshot of justifying an adherence to the default and challenge model of entitlement on normative rather than simply descriptive grounds. My task, in chapter four, is then to offer an account of discursive justification which functions for local justification and has this global justificatory element built in, ensuring that if something *can be justified* by this account then we have a way of saying it is grounded in the way things really are. This is accomplished by tying together both causal and inferential justificatory moves, connecting a belief causally to the world and inferentially to other beliefs. Finally, this allows me to in chapter five to formulate a condition of justification for belief based on the potential for that belief to be discursively justified.

Chapter Two: The Discursive Community

The preliminary task of this thesis, as articulated in the previous chapter, is twofold: to show that, given our epistemic practices, we are able to come to large-scale agreement, and that this agreement will be grounded in the way the world really is. This chapter will offer a basic sketch of the structure of our epistemic practices in the attempt to show that there is already widespread agreement among epistemic agents about what is true and the belief-forming methods for obtaining true beliefs, and that our practices, if they run smoothly, will inevitably widen this agreement. Although I will not give a detailed account in this chapter about how this widening of agreement actually occurs by the process of discursive justification (I will save that for chapter four), the epistemic practices discussed in this chapter will feature centrally in my attempts to answer the epistemological questions of the next two chapters.

1. The Discursive Community

We humans, social beings that we are, are community members by nature. There are many communities that we belong to, and they vary greatly in terms of the formality and rigidity of membership conditions. To be a member of a certain community is to be held to certain standards of behavior and thought. If I am a member of the Skeptic Society, and I believe in Big Foot, the Loch Ness Monster, and UFOs, there are certain community standards that I am violating. Likewise, if I am a Boston Red Sox fan, and I go to a game wearing Yankees gear, then I am also violating community standards. In both of these cases, I'm going to have some explaining to do if I want my membership in the community to continue to be taken seriously. As a member of these communities, I am held responsible for my transgression of their norms. One response is to acknowledge the transgression and

accept the criticism as a reason to change my beliefs or behavior. On the other hand, however (assuming my offense isn't perceived as too great), if I want to *maintain* commitment to my beliefs or behavior, I am allowed to give reasons *for* my apparent violation of the community standards. Perhaps some surprising new and credible evidence for UFOs and Big Foot just came out of which my cohorts are unaware, and perhaps the reason I was wearing Yankees gear was because it was national opposite day.⁷

For Robert Brandom, it is this responsibility to conform to standards and be subject to reasons that distinguishes us from mere animals. He writes,

“We are the ones on whom reasons are binding, who are subject to the peculiar force of the better reason. This force is a species of normative force, a rational ‘ought.’ Being rational is being bound or constrained by these norms, being subject to the authority of reasons.”⁸

Drawing from intellectual history, Brandom jumps on the phrase “rational animals” to mark out the sort of beings that we are, but another way of making the distinction to highlight the essential normative element is to say that we are “*responsible animals*.”⁹ Insofar as we are persons, we are held *accountable* for what we say and what we do in accordance with norms constitutive of the communities to which we belong. It is central to my explanation of epistemic norms here that by virtue of this basic responsibility and accountability that we all share, we enter into a community of the most broad and basic type: a community of responsible persons. Membership in this community is a prerequisite for membership into

⁷ The actual date of opposite day is January 25th, and this is not during baseball season, so if someone gave this as a reason, they'd have some explaining to do.

⁸ *Making it Explicit*, 5

⁹ Brandom takes the norms of rationality *proper* to be the sort of central norms governing. He takes the normative. According to Brandom, the possibility of being held responsible under the jurisdiction any norms comes down to being held responsible to the norms of rationality.

any other community, since to be a member is to be able to be held to normative standards at all.

This community could be called various things, setting the bounds by one characteristic of its constitutive members or another. In their recent work, Rebecca Kukla and Mark Lance use the term “discursive community,” to demarcate the boundaries as those beings who can engage in discourse, in practices of giving and asking for reasons.¹⁰ I think this is a good way of setting the boundaries, and I will follow their usage. We aren’t really holding someone responsible for their claims or actions unless we understand what would be necessary for them to justify themselves.¹¹ As we saw with the two examples earlier, we hold a transgressor of community standards responsible in the case that they *lack* appropriate reasons for their actions. In order to really hold someone responsible for their actions, they must be an agent capable of giving reasons. A dog might tear up the carpet, and, though we might be upset at the situation, we do not hold the dog truly responsible because we know there is no possibility for him to *account* for his action. We might say that there is a sense in which the dog is not the “owner” of his actions in that he can’t, to use the common phrase, “*own up*” to them, acknowledging his transgression of the norms and taking responsibility for it or offering justifying reasons for his action.

An important point for our purposes here is that only members of the discursive community can *claim* something to be true or false. This highlights the fact that making a true claim is not the same thing as uttering a true sentence. A parrot may squawk out true sentences all day and night, but he isn’t *claiming* that these things are true—he’s just spouting them out. Unless one is *responsible* for the sentences they are uttering, they are not actually

¹⁰ Kukla and Lance “Yo! and ‘Lol!” 196

¹¹ Although, depending on the weight of the transgression we might, in extreme cases, judge that there is *nothing* that could count as an appropriate justification.

claiming that the sentences are true. This means that, in order to have the ability to claim something true or false, one must be imbedded in all sorts of normative practices of giving and asking for reasons by which one can be held responsible for their claims. Making a claim enters one into a sort of contract. One is obligated to defend their claims if they are absurd or inflammatory, if the stakes are particularly high, or if there is just some good reasons to doubt them, since when one puts forward a claim, they are putting it forward *as something true*, and thus, something *others* should *take as true*.¹² The responsibility we have for our claims rests on the fact that when we are making claims we are engaged in a social practice in which we can lead others astray by claiming things that are false, misleading, or irrelevant.

2. The Game of Giving and Asking for Reasons

Brandom repeatedly uses the term “game of giving and asking for reasons” to describe this sort of discursive practice in which we are immersed. The game can be centrally articulated in terms of two sorts of normative statuses: *commitments* and *entitlements*. The basic sort of commitment on which belief can be modeled is an *assertional* commitment, a commitment we have to our claims that ties us to them as *responsible* in the sense just discussed. When we make an assertion we commit ourselves to taking it as true due to the sort of normative social practices we are involved in when we make assertions. These practices account for the absurdity, though not strictly logical contradictoriness, of the Moorean paradoxes such as “P’ and I don’t believe it.” On this model, belief is a socially situated *taking true*, and putting forward something as *true* (asserting it) while putting forward the assertion that one is not *taking it as true* (renouncing commitment/belief from it) is a discursive absurdity.

¹² *Making it Explicit*, 170

Correlated to the status of commitment is that of *entitlement*. Entitlement is a sort of social license for commitments, and to attribute entitlement to an assertional commitment is to grant it a sort of positive epistemic status. If someone is attributed entitlement to their commitment it is appropriate for them to stay committed to it, to continue to take it as true.¹³ We are attributed entitlement for the majority of our commitments by default. That is, the default-status of our commitments is entitlement to them. Without treating discursive agents as largely entitled to their commitments, there is no way to get assertional discursive practice up and running. For an agent to commit herself to something, to take it as true, is, in the paradigm case to entitle others to the same commitment. If I say, “The Red Sox won the game today,” it’s part of the pragmatic structure of my speech act, part of what makes it an assertion, that others can take it as true when they hear me say it. Asserting would lose all of its pragmatic force if, in the standard case, the speaker was not entitled such that the entitlement to the assertion could be passed on to those who hear it.

Even though commitments are entitled by default, it does not mean that they cannot lose that status. Other participants can challenge a participant’s commitments, and when a commitment is *appropriately* challenged (when a challenger is entitled to challenge) and *successfully* challenged (if the asserter is unable to “vindicate the commitment by demonstrating entitlement to it”) it will result in in the asserter *losing* entitlement to that commitment.¹⁴ Thus, Brandom calls his model the “Default and Challenge Structure of Entitlement.” In response to a challenge, a participator can rescind her commitment, challenge the entitlement of the challenger to the challenge, or, acknowledging the

¹³ “Attribute” is the term that Brandom uses here to refer the basic practice of treating one as entitled. Diverging slightly from the standard usage of the term, “attribute” is intended as not merely a descriptive action, but partly *constitutive* of the status itself. The sense in which the term functions is explained in more depth on page 21 with respect to the word “judge.”

¹⁴ *Making It Explicit*, 178

appropriateness of the challenge and maintaining the commitment, *offer a justification* for her commitment with the aim of securing entitlement for it. The result of this practice, if all goes smoothly, is a widening of commitments and entitlements shared by community members.

It is important not to over-intellectualize this whole practice to the point at which it becomes dubious whether or not everyone even engages in it. Hilary Kornblith raises this point with a skeptical eye to Brandom. He regards being a player in the game of giving and asking for reasons as being involved in a sort of *Socratic practice*. As a sort of reductio, Kornblith says that on the Brandomian picture, “Someone who never gives or asks for reasons, and is part of a culture that simply opts out of the practice, does not have beliefs at all.”¹⁵ However, what is centrally being talked about here is not the intellectualized practice of giving and asking for reason that we find in academic disciplines, but the general non-academic practice that permeates all areas of our everyday life. Of course, for some cultures, there will be localized pockets of beliefs for which the sorts of reasons given and accepted will diverge from our own. In some cultures there are sets of beliefs that are treated as default-entitled that we think can and ought to be challenged, and this may appear on the surface as opting out of the practice of giving and asking for reasons. But to think that some cultures might opt out of the practice *entirely* is to misunderstand the broad scope of what is meant by the phrase “giving and asking for reasons.” What is meant here is the base-level responsiveness to reason and the ability to account for claims and actions that separates us from non-linguistic animals like dogs.

3. Building the Epistemic Toolkit

¹⁵ Kornblith, *Knowledge and Its Place in Nature*. 86

If all goes smoothly in the game of giving and asking for reasons, it's possible for all agents making truth-claims to come to a principled agreement. But what makes us think the practice will go smoothly? The fact that the ideal projected by the way our discursive practice functions is agreement does not actually ensure that such agreement is possible. On the face of it, it seems possible that discursive agents might run into insurmountable walls in playing the game of giving and asking for reasons. What we need to show, then, to exclude this possibility, is that we have enough common ground both in beliefs and belief forming practices with anyone who we might engage in discourse with to come to agreement. The strategy I will employ to show this is to argue that each of our capacities necessary for the coming to agreement with other truth-claimers is also a requirement for membership in the discursive community. If a certain extent of failure to employ a particular capacity is sufficient to get "kicked out" of the discursive community, then we can say that employing that capacity is necessary for membership.

To kick someone out of a particular community is to say that they are no longer *entitled* to enter that social space, and to treat someone in such a way that we can take away entitlement and not *allow* them to participate in certain practices as a member of a certain community is to still treat them as a *person* who can be held to be *responsible* to certain normative standards. Let's imagine I get kicked out of a magicians' club for sharing secrets with non-members. The possibility that I can get kicked out of the magicians' club means I'm still treated as a responsible *person* that can be held to more general normative standards. If I keep going to the meetings, I will likely break any ties of mutual respect with my fellow magicians, and if they take my actions seriously enough, they might say, "If you keep coming

we are going to call the police.”¹⁶ Suppose, I go again, and a restraining order is issued. A restraining order assumes a certain amount of personal responsibility but, more officially, it substitutes a purely social norm of reasonable interaction among respectful people for a rigidly instituted norm with strong material (not just social) consequences for transgression (fine or imprisonment). Suppose I break the restraining order, and I am taken to prison by force. Even here I am held to certain normative standards. The officer might say, as I struggle to get out of handcuffs, “Come on. Cooperate.” Now, I am not necessarily obligated to follow the officer’s suggestion in order to continue to be held to normative standards, but I am held to *understand* what the officer wants me to do and to somehow *show* this understanding by acknowledging responsibility to a set of norms. I might exclaim, “I refuse!” or something of the like, acknowledging that there is a normative demand placed upon me, but it is one that I am rejecting. Suppose I don’t make such an exclamation, and instead I continue to struggle against the force of the law. Eventually I’ll find myself in an institution in full body restraints. How else might I disobey norms after this point? Well, I might break all the conversational norms, blabbering nonsense and spitting in everyone’s face, and, if this continues long enough, I’ll be treated as basically comatose, broken in some important sense, and can no longer appropriately be treated as a person.

There are two types of sanctioning occurring in the previous set of examples: material and normative sanctioning.¹⁷ In conjunction with often being *materially* sanctioned (such as being physically restrained, either in handcuffs or jail), our actions are *normatively* sanctioned in one way or another (such as being kicked out of a particular club or marked as

¹⁶ It is important to note here that I must be conforming to various other discursive norms in order for this action to be taken. That is, I must be still acting, in some sense, reasonably, displaying a sense of urgency to explain my radical behavior. Otherwise, I will be treated as *malfunctioning* in some important way.

¹⁷ Brandom discusses the distinction at length in the first chapter of *Making it Explicit*.

a convicted felon). Something is a normative sanction insofar as it is understandable as only involving a change in the normative attitudes that can be appropriately adopted towards an agent, independent of any material sanctioning. So, I might get kicked out of a club even if I simultaneously vow never to attend again, or I might get convicted as a felon even if the president commutes any prison sentence. Likewise, a child may do something wrong and his mother, without issuing any punishment, might simply say that she's just *really disappointed*. The ultimate normative sanction for an agent's transgression is for that agent to *no longer be treated as an agent capable of being held to normative sanctioning at all*. For this to happen is for a community to decide that adopting a particular normative attitude towards an agent *itself* can provide no sanctioning effect on that agent.

To get "kicked out" of the discursive community is to stop being a discursive agent. I use scarequotes for "kicked out" here, since it is a normative notion. We can't kick people out of the discursive community in the way that we can kick someone out of a magic club. Rather, we judge that it would make no sense to think of them as a member of the community of individuals who can be held to reasons. This is a two-sided action, both descriptive and constitutive. In one sense, our action of judging is a passive *realization* that this individual is not a discursive agent. On the other hand, this judgment forms a part of a set of actions *constitutive* of the status. Being a discursive agent cannot be understood apart from the status of being *treated* as a discursive agent. To treat someone as not responsible to normative sanctioning and only material sanctioning is to stop treating them as a responsible agent, to remove their membership from the discursive community entirely. In the example I've given, eventually, refusing to follow any conversational norms, I can no longer appropriately be treated as a person. Of course, this is an extreme case of the ultimate normative sanction being imposed, and there are less extreme examples. Traumatic

brain injury or extreme social conditioning might also make one impervious to any normative sanctioning.

Now, I want to propose a few capacities which I take to be necessary for epistemic success and show that the successive failure to employ these will result in removal from the discursive community entirely. That is, since any proposed agent lacking these capacities cannot be treated as discursive agent, all discursive agents must possess these capacities.

3a. Inferential Capacities

One essential element if we are going to have epistemic success is the capacity to form beliefs *rationally*, so as to come to agreement in a principled manner. Some psychologists have argued that we *really are not* “rational animals” as Aristotle and others such as Brandom have defined us, but rather have a strong sense of irrationality ingrained into us.¹⁸ Though, of course, there are important observations here on how humans tend to behave irrationally in certain situations, my claim here is that we can only understand ourselves as discursive agents if we are *essentially* rational.

In trying to articulate the nature and function of propositional attitudes such as beliefs and desires, Daniel Dennett has done a significant amount of work developing what he calls “intentional stance.” To adopt the intentional stance is to interpret something as an “intentional system,” a being that acts in accord with beliefs and desires. Central to Dennett’s view is the idea that, we can only treat something as having beliefs if we assume it is rational. He observes,

“one gets nowhere with the assumption that entity x has beliefs p, q, r . . . unless one also supposes that x believes what follows from p, q, r . . . ; otherwise there is no way of ruling out the prediction that x will, in the face of beliefs p, q, r . . . do

¹⁸ Jonathan Haidt, makes this case with respect to political orientation in *The Ritegeous Mind*.

something utterly stupid, and, if we cannot rule out that prediction, we will have acquired no predictive power at all.”¹⁹

In order to be believers, we need to possess these materially embodied inferential capacities.

We can imagine a “radical interpreter” attempting to interpret an animal (human or otherwise) as having belief with particular contents.²⁰ The only way she can do so is by interpreting it as *acting in accord* with its beliefs and desires, and this notion of acting in accord with beliefs and desires, can only be cashed out in terms acting *rationally*. Thus, in order to adopt the intentional stance to an organism, we must interpret it as a rational organism.

There is, of course, a difference between being rationally *interpretable*, like a chess playing computer or an amoeba, and being rational in the sense of rationally *responsible*, having the sort of responsibility characteristic of community members and the ability to respond in *explicitly rational* ways to challenges. Dennett’s intentional stance makes relatively little of this distinction. For Dennett, the intentional stance applies equally to nonlinguistic animals and things like chess-playing computers as it does fully-discursive humans. Peter Hacker has criticized Dennett along these lines, saying,

“In order to be able to be able to treat a being *as if* it believed that p or wanted to V, we must know what it *is* to believe or to want something. We must know under what circumstances it is appropriate to avow a belief or desire in our own case. Only when we have grasped this—that is, when we have understood the use of the verbs ‘to believe’ and ‘to want’ (in both the first and the third person)—can we come to understand what it is to treat a being *as if* it believed this or wanted that. But on Dennett’s account the adoption of the intentional strategy towards an intentional system, including human beings, is *never* other than an ‘as if.’”²¹

This same point has been made, much more constructively by Brandom. According to Brandom, our ability to adopt the intentional stance and attribute rationality to chess-playing computers and nonlinguistic animals alike “derives from the mastery of the richer practices

¹⁹ Dennett, “Intentional Systems,” 11

²⁰ Davidson, “Belief and the Basis of Meaning” and elsewhere

²¹ Hacker and Bennett, *Philosophical Foundations of Neuroscience*, 426

of giving and asking for reasons for the sort of doxastic and practical discursive commitments that are not attributed to these simpler folk.”²² The full-scale belief of the sort we are concerned about here is to be understood only as contextualized in the Brandomian practice of giving and asking for reasons described in this chapter, as linked inexorably to rational responsibility. As such, we cannot have beliefs apart from the ability to respond appropriately to challenges. And to respond appropriately to a challenge is to respond in a way that is treated by the discursive community as *rational*. So having a belief, in the full-fledged Brandomian sense, presumes, not just an implicit sort of rationality that Dennett deems necessary, but a rationality that can be seen in the *explicit articulation of reasons*. To say that we are “rational animals” in the full sense is to say that we are *responsible* for following norms of rationality, and we are accountable for the rationality (or irrationality) of our actions or beliefs.

We can only make sense of *explicit* inference rules and justifications for knowledge claims once we are fully immersed in the practice of inferring and treating appropriate inferences as entitled. Brandom writes, “One must start with a notion of taking or treating inferences as correct in practice. Without such a practice there is no game of giving and asking for reasons to bring inferences into in the form of explicit assertions.”²³ More will be said in the next chapter to argue that the way in which we treat inferences as correct in practice is *actually* conducive to true belief, but for now, it is enough to note that this practice of treating inferences *as* rational is necessary to getting an epistemic practice up and running.

3b. Recognitive Capacities

²² *Making it Explicit*, 276

²³ *Making it Explicit*, 205.

It is not enough, however, that we share basic inferential capacities. We must also share the basic ability to mutually recognize things in the world, giving us the friction we need to agree about the world around us. In their book, *Yo and Lo!* Rebecca Kukla and Mark Lance argue for the importance of making a distinction between normal declarative utterances which merely describe something one saw or that one saw it and *recognitives*, which function to “give expression to a speaker’s recognition of something.”²⁴ Unlike traditional declaratives, these speech acts are essentially indexed to the speaker—that is, they essentially incorporate the speaker’s first person voice. Two sorts of recognitive acts that K&L view as absolutely essential for membership into the discursive community are what they call “Yo” acts and “Lo” acts (hence the title of their book). “Yo” acts, of the sort “Yo, Hana!” are *vocatives*, which call a person to recognize the recognitive force of the act. Vocatives, when performed in appropriate situations, bring the one who is called into a normative relation with the caller.²⁵ If I say, “Yo, Hana!” and she doesn’t respond, she is transgressing a norm. It is important to note that this normative demand wasn’t there before I performed the communicative act; rather my vocative *creates* this normative relation. Hana had no obligation to discursively engage me before I called her with this vocative.

The epistemic importance of vocatives is obvious. Without the ability to recognize others and call them out to engage with us, there is no way of engaging in the practice of justification which essentially involves calling others to challenge their claims. If we want to find common ground between our views and anyone else’s we need to be able to discursively engage the person, bringing them into a second-personal relation that normatively binds them to give reasons. Jeremy Wanderer has argued that Brandom’s notion of a challenge

²⁴ Kukla and Lance, “*Yo!*” and “*Lo!*” p. 45

²⁵ “*Yo!*” and “*Lo!*”, 141

must be thought of as a second-person normatively binding speech act in which “the reciprocal act of acknowledgement of the address is an essential part of the act that it is.”²⁶

This reciprocity is also essential to K&L’s account of vocative “Yo!” acts. In order for the act to be felicitous, it is not enough for the speaker to recognize the person at whom the act is directed; the person at whom the act is directed must recognize me as recognizing them.

It is also clear that vocatives are essential to being a member of the discursive community. If

I can’t recognize others’ attempts to call me, I am simply unresponsive. In order to recognize anyone as a person, I must recognize them as having the ability to recognize me—to call me out and to acknowledge my calling out.²⁷ Without competence in performing and responding to vocative acts one cannot be a member of the discursive community. This is for the simple reason that vocative acts are essentially two-way acts. In order to fulfill its function, a vocative must be acknowledged by the one to whom it is addressed.²⁸

Something that can’t “Yo!” back is not a thing that can be “Yo!”ed in the first place, and thus not a discursive agent at all since no one could possibly discursively engage it.

In addition to the vocative recognitive, Kukla and Lance also highlight a recognitive act they call an *observative*. An “*observative*,” such as “Lo, a rabbit!” is a special type of recognitive that gives expression to a state of affairs *in observation*. An *observative* doesn’t simply declare that a certain state of the world is the case, or that one is seeing that it is the case, but *discursively marks* the receptive intake itself. If I say “Lo, a rabbit!” it might be appropriate for you to assert to someone else that there is a rabbit around, but unless you

²⁶ Wanderer, “Brandom’s Challenges” 102

²⁷ It is worth noting that this recursive acknowledgement between agents has been featured prominently in the empirical literature on the. Notably, Michael Tomasello (2008) regards this “recursive mindreading” as essential to the shared intentionality that distinguishes human communication from ape communication.

²⁸ “Yo!” and “Lo!” 146

actually see the rabbit, you can't exclaim "Lo, a rabbit!" yourself. A "Lo!" claim, like "Lo, a rabbit!" isn't just an observative, but is an *ostensive* observative. That is, it doesn't just mark one's recognition of something in observation (in the way a speech act like, "Oh wow, a rabbit!" might) but *calls others to recognize the object as well.*²⁹ The ability to produce "Lo" acts comes down to the ability to respond reliably to objects in the world, to take them *as* the objects that they are, and call another's attention to these objects.

On Kukla and Lance's account, the epistemological importance of acts of ostension is very clear: they "establish our responsiveness to a shared empirical world and make possible reasonable debate about that world."³⁰ Particularly, not just observatives in general, but *ostensive* observatives establish this *shared* contact with the world. For this reason, ostensive "Lo" acts are particularly important tools in our epistemological tool kits. In the end of chapter three, I will argue that the epistemological centrality of these sorts of ostensive acts suggests a divergence from traditional accounts of justification such as foundationalism and coherentism, but for now, it is enough to note the epistemological upshot of these acts. For our present task, however, it is not just enough to note that these acts are useful tools in the practice of gaining common ground. We must also show how they are *fundamental* to our status as discursive beings. Only then can we throw them in our kit.

Kukla and Lance argue for the discursive essentiality of "Lo!" acts by tying them directly with "Yo!" acts. All of our interactions with each other are *materially embodied* interactions. We can't interact with one another unless we can recognize one another, and "recognize" here (since all these interactions take place in a material world) must be a "literal

²⁹ It is, after all, a contraction of the imperative, "Look!"

³⁰ "Yo!" and "Lo!" 210

receptive encounter.”³¹ So, a “Yo!” act, which we have seen is essential to membership into a discursive community, simultaneously functions as a “Lo!”³² We can only agree or disagree with people whom we recognize as rationally responsible agents, and we can only do this if we there is a mutual recognition of another as materially embodied agents, material beings among other objects in a shared material world.

4. A Picture Forming

There is a close link here between the sort of move that Kukla and Lance make in connecting recognition of others as discursive to shared recognition of objects in the world, to the move Davidson makes in “On the Very Idea of a Conceptual Scheme” and elsewhere. On Davidson’s picture, in order to recognize someone as a discursive agent, we must attribute beliefs to them, and this requires that we must have *some* beliefs about the same objects. We must also, as Dennett notes, treat them as having the same inferential capacities as our own. The fact that shared observative and inferential capacities are necessary for our recognition of each other as agents capable of discourse shows that there cannot be *complete* incommensurability among any discursive agents whom we can recognize as such. This Davidsonian point is a crucial one for my present project. It takes full force in chapter four where I attempt to show how this impossibility of complete incommensurability allows us to deal with the worries of scientific incommensurability brought upon by the theory-ladenness of observation stressed by philosophers like Kuhn or Feyerabend. Before I do this, however, I will articulate the Davidsonian picture in greater detail and spell out how it allows

³¹ “Yo!” and “Lo!” 211

³² At least in the core cases, since there are circumstantial outliers in which this is not the case. Imagine giving a talk and thinking that there might be a hand of a person asking a question in the back of the room. You might say “Is that a hand back there?” calling out a person with a “Yo” act if there is one there, but not committing yourself to an observative “Lo.” Cases like these, however, are deviations from the norm in which a “Yo” also functions as a “Lo.”

us to dismiss the worry of global skepticism. That is the task I turn to in the next chapter, and it will serve as the cornerstone of this thesis.

Chapter Three: How to Tell the Skeptic to Get Lost

In the last chapter, I tried to show that discursive practice will move towards agreement. According to the preceding analysis, we are all players in the game of giving and asking for reasons. We are epistemic agents, responsible for our claims, we can call each other to engage in the practice of justification to settle disagreement, and we hold each other to certain shared epistemic norms in doing so. But the whole analysis offered in the last chapter rests on the notion that we *really are* materially embodied agents functioning in a shared world, able to mutually recognize objects that are *really there*. No matter how much analysis of the pragmatics of epistemic norms is conducted, an epistemological question remains: how do we know we're getting it right? Kukla and Lance make sure to clarify that, in answering questions about the pragmatics of our epistemic practices, they are not attempting to answer this epistemological one:

An epistemologist might ask: if everything is working smoothly, how can we tell if we are getting objects right or wrong? How do we know whether our social practices of disclosure and justification are in fact the appropriate practices for getting the world right? But we are trying to get at a pragmatic rather than an epistemological question here. We want to know: What sort of social practice involves us holding one another to getting objects right rather than merely coping with them effectively, and how can objects play a role in such a practice? How one would check whether that practice is successful is a different matter. If we mistakenly think that we are trying to answer the epistemological question, then the problem as we have set it out here can seem insoluble: after all, the actual norms we are committed to are the only norms we have for telling how things are. If our practices are working smoothly for telling how things are, from a material and social point of view, then it seems that we have all that we can have. How could we distinguish between practices that merely aim to get things right and those that succeed?³³

The threat of having practices that aim to get it right go smoothly, but massively fail is the threat of epistemological skepticism. Answering the epistemological skeptic in a satisfactory

³³ Rebecca Kukla and Mark Lance. "Intersubjectivity and Receptive Experience." *The Southern Journal of Philosophy*. Forthcoming.

manner is the task of this chapter. This is the task for two reasons. First, as I've just said, any investigation of the normative pragmatics of epistemic practices must take its starting point on the fact that we really are embodied agents interacting with real objects in the world. But second, and more centrally, the possibility of a complete divergence between our actual practices that aim at getting it right and practices that really do get it right would prevent us from drawing epistemological conclusions from descriptive work on our epistemic practices. If this massive divergence *isn't* possible, then answering broadly empirical questions about how the practices of justification and knowledge attribution most central to our epistemic practices *actually work* can shed light on the normative question what it means for a belief to *really be* justified or to count as knowledge.

Of course, there might be some situations in which there is a particular *localized* practice that is going smoothly but getting it wrong. Tarot reading, for example might be going perfectly smoothly by the epistemic norms of Tarot reading, but be coming out with wildly false beliefs, and so obviously I don't think studying the pragmatics of the epistemic practices involved in Tarot reading would have any epistemological upshot. But it seems like there is a big difference between these sorts of localized epistemic practices and our *core* epistemic practices, the practices that I identified in the previous chapter, which are *necessary* to get epistemic practices in general up and running. I argue in this chapter that the divergence isn't possible with respect to this latter set of core practices. According to this line of reasoning, the upshot of arguing against skepticism isn't simply the epistemological conclusion that our beliefs must be mostly true, but the *meta-epistemological* conclusion that, since our actual epistemic practices are getting it right, to do epistemic pragmatics is really to do epistemology proper. It is this latter conclusion that I am most interested in, and my final

argument against the skeptic here will be a (hopefully successful) example of getting epistemological mileage out of a descriptive account of our epistemic pragmatics.

My argument here will draw heavily from Donald Davidson's epistemological work. Davidson has argued that the very idea of having a belief requires that most of our beliefs are true, and so epistemological skepticism is fundamentally mistaken. In response, Rorty has said that this doesn't really amount to answering the skeptic, but rather, just "telling him to get lost" (something Rorty is entirely fine with). I will argue here that there is a way of legitimately answering the skeptic that *consists in* explaining to him why he really must either change his ways or "get lost" from the epistemological conversation. To do this, I first endorse Davidson's arguments against skepticism and attempt to strengthen his position by formulating it as a *transcendental* argument. Then, Drawing from an analysis of the Matrix scenario proposed by David Chalmers, I we must have mostly true beliefs *even if* any of the various Cartesian-style skeptical hypotheses are true. Finally, I argue that, if Davidson's right and I succeed in defending him, the epistemological consequence of the fact that the majority of our beliefs are true and our central belief-forming practices are truth-conducive is that we can take what we know about the *pragmatics* of our central epistemic practices to have genuine *epistemological* upshot. This gives us a final way of silencing the skeptic by pointing out, with newfound epistemological force, that he's breaking all the rules of epistemic discourse.

1. Davidson on Meaning and Skepticism

Davidson's arguments against skepticism are a direct consequence of his theory of meaning, so an explanation of the latter is necessary to appreciate the former. Consider the following example: I am shipwrecked and find myself on an uncharted island with an unknown native population. I stumble across one such native whose language is entirely

foreign to me. To understand what he is saying, and correlatively, what his beliefs are, I must take an interpretive stance towards him and see how he responds to stimuli in his environment. This three-part relation between interpreter, speaker/believer, and shared stimulus is what Davidson calls “triangulation.”³⁴ So, to use Quine’s example, I might hold up a rabbit, and he might (presumably understanding that I am trying to learn his language and not, for example trying to attack him with the rabbit or offer a gift to him), exclaim “Gavagai!” In trying to understand the language he is speaking, I also will attribute beliefs to him. Suppose that I take “gavagai” to mean “rabbit.” I then will attribute to the native the belief *that a rabbit is present*. I do this on the grounds that, since a rabbit is, in fact, present and clearly visible, this belief is appropriate. In order to make sense of the native’s beliefs at all, I must attribute to the native the beliefs that he *ought to have*, and, in doing this, I must attribute *mostly true* beliefs to him.

The insight from this radical interpretation argument can be extended to language *learning*, not just interpreting, where it takes its full force in highlighting the connection between the cause of an utterance or a belief and its meaning. On Davidson’s account of meaning, meanings of the most basic kind are correlated directly with the objects that cause them. He writes, “What stands in the way of global skepticism of the senses is, in my view, the fact that we must, in the plainest and methodologically most basic cases take the objects of a belief to be the causes of that belief.”³⁵ In the conditioning process of language-learning a child learns to say certain words and sentences in the situations for which they are appropriate. For simple object words, some of the first words learned by a child (and the

³⁴ Donald Davidson, “Three Varieties of Knowledge” in *Subjective, Intersubjective, Objective*. Oxford University Press: 2001.

³⁵ Davidson, “A Coherence Theory of Truth and Knowledge” in *Subjective, Intersubjective Objective*, pg. 151.

sorts of words we can associate with empirical beliefs about one's immediate environment), the appropriate situation in which they are to be uttered are when the object is *present*, as a response *to this object*. Though this is a rough and oversimplified picture as Davidson admits, a child learns the word "rabbit," for example, by positive conditioning whenever a rabbit is present and negative conditioning when one is not present. The same sort of triangulation that occurs in radical interpretation occurs in language learning. That is, there is a three-way interaction between a language-learner, a language-conditioner and a stimulus in the world. Davidson, following Wittgenstein, asserts that to know the meaning of the word is no more than to be able to use it appropriately. To know the meaning of "rabbit" for example, is to be able to say it, in the most basic cases, *as a response to rabbits*.

Despite Rorty's charge that Davidson is not answering the traditional epistemological worry, Davidson is by no means modest about the epistemological implications of his considerations regarding belief and meaning. He concludes, "The fallout from these considerations for the theory of knowledge is revolutionary. If words and thoughts are, in the most basic cases, necessarily about the sorts of objects and events that commonly cause them, there is no room for Cartesian doubts about the independent existence of such objects and events."³⁶ The idea here is pretty straightforward. The concepts in which my beliefs and thoughts are framed arise in a communal context which essentially involves another language user and shared objects and events to which language-speakers can both respond. In order for us to have beliefs about rabbits, for example, and for it to make sense to talk about these

³⁶ Davidson, "The Myth of the Subjective," in *Subjective, Intersubjective, Objective*. 45.

beliefs, there must *really be* rabbits and these rabbits must *really be* the things causally responsible for our beliefs about rabbits.³⁷

2. Making Davidson's Argument Explicit

To many, these arguments, if proposed as a direct refutation of skepticism, seem lacking. Since Davidson's arguments seem to rely on broadly *empirical* observations about the way agents interacting with objects can be interpreted and about the way the language learning process works, the skeptic would assert that it already *assumes* that we are basically right about the way the world is. Accordingly, all it shows is that we must be largely in agreement about this assumption, but not largely right about the way the world actually is. Relying on the way language actually works to argue that most of our beliefs must be true, so says the skeptic, is to take too much for granted, for it is to already assume the very conclusion it is attempting to argue for. The traditional responses to global skepticism have often attempted to provide an a priori *foundation* for knowledge, and Davidson is clear that his anti-skeptical arguments suggest that "empirical knowledge has no epistemological foundation, and needs none."³⁸ In this case, however, how can Davidson's argument really be a response to global skepticism?

The answer, I believe, is to think of these arguments as *transcendental* ones that entirely *erase the need* for an epistemological foundation from which we can infer the existence of the objective world. This clearly seems to be what Davidson is up to when, he says, for example:

If one can show, as I think is possible, that in order to have a thought, even a doubt, one must already know that there are other minds and an environment we share with

³⁷ At least in the paradigm cases, since, of course, a rabbit-shaped rock might sometimes be causally responsible for a belief about a rabbit.

³⁸ Davidson, "Myth of the Subjective" 46

them, then this amounts to saying that it is impossible seriously to doubt these things—we cannot give a coherent content to such doubts.³⁹

In putting things in this way, where we move from the fact that one is thinking to the existence of one's environment, it appears on the face of it as if Davidson is following the anti-skeptical strategy of Descartes. However, if Davidson really thinks that empirical knowledge needs no foundation, then we cannot think of this strategy as one involving an *inference* from inner certainty, outward to the world. Rather, we should think of the move as *elucidatory* rather than inferential. That is, once we lay out what the concepts of thought and belief *really mean* and illuminate the sorts of things that must be in play for these concepts to *even make sense*, we understand that there must be an objective world about which we are mostly right.

We might construe the argument as having two essential points: (1) belief and understanding only makes sense within the context of our public intersubjective discursive practices, and (2) within this context, there is no possible way to deny the *objectivity* of the world. The first point is inherited from latter Wittgenstein and we can think of it in terms of the normative nature of understanding. We cannot think of things like belief and understanding as divorced from essentially social discursive norms. This broadly Wittgensteinian point is a core commitment of this thesis. The second point is the crucial one that provides the transcendental response to skepticism Davidson can offer. Our epistemic practices are essentially *world-involving* practices. That is, it we cannot make sense of ourselves as having epistemic practices at all if we do not also regard there as being objects in the world that *constrain* these practices. The very notion of unconstrained epistemic practices is nonsensical. It is this constraint which constitutes the *objective* dimension of our

epistemic practices. In order to make sense of ourselves as believers at all, we must regard of our beliefs being *responses* to objects. This is the crucial third point of triangulation. There is no way to take our practices *away* from the objects in the world they are about, and still make sense of them. To try to take this objective third point of triangulation away is to verge on incoherence. As an illustration of this point, consider the following exchange:

SKEPTIC: The external world does not exist!
BELIEVER: Ok, have a seat on that couch across from me and I'll try to change your mind.
SKEPTIC: (Sits down.) So, let me hear your argument.
BELIEVER: (holds out one hand in between them.) *Here is a hand!*
SKEPTIC: Oh, no! I won't fall for that one! You're simply begging the question.
BELIEVER: Am I? So, *I* see my hand right now, and I presume *you* see my hand as well. It's right there in front of our faces.
SKEPTIC: No, but if my argument is right, we could both be deceived, and it could be the case that you don't *have any hands*.
BELIEVER: Ok, if you don't want to call this thing that I am holding out a "hand," that's just fine with me. But "hand" seems like a perfectly suitable word for this thing, and that's what we've *been* calling these things throughout the years.
SKEPTIC: But that's not an argument!
BELIEVER: Nope. It's better than an argument. It's a hand!

Perhaps it is a failure of imagination, but I'm not quite sure what the skeptic's next response would be. The reason why an ostensive move like this is so powerful against a skeptical challenge is because employing our ostensive capacities, as I've said in the previous chapter, is a necessary requisite of our status as discursive agents. We must engage in ostensive practices and these practices necessarily commit us to understanding the world that we are all responding to as *objective*. The crucial

In this context, I believe we must see Davidson as committed to a form of direct realism regarding perception, even though he never explicitly endorses such a view. Direct realism, as I use it here, is the thesis that the things we are directly aware *objects in the world themselves*, and, in standard cases, we form our perceptual beliefs in virtue of our awareness of these objects rather than awareness of some mental phenomena. "In virtue of" here, is

meant in an *epistemological* sense rather than a *causal* sense. Davidson says, “Although sensation plays a crucial role in the causal process that connects beliefs with the world, it is a mistake to think it plays an epistemological role in determining the contents of those beliefs.”⁴⁰ There are all sorts of causal intermediaries between the object of an utterance or belief and the belief or utterance itself (such as the set of light rays bouncing off an object, the stimulation of the eyeballs, the electrical signals going down the optic nerve, etc.), but Davidson insists that the only causal connection of serious *epistemological* significance for us is that between the object itself and the belief.

However, at least in some places in Davidson’s writing, it seems as if he not only wants to restrict *mental sensations* from being the epistemic basis of our perceptual beliefs, but *objects* as well. In “A Coherence Theory of Truth and Knowledge,” he endorses *coherentism*, the epistemological view that “nothing can count as a reason for holding a belief except another belief.”⁴¹ Such coherentism, according to John McDowell, leaves us with the threat of “spinning frictionlessly in the void.”⁴² If Davidson to actually endorsed this view it would be rather mysterious, since the sort of “frictionlessness” McDowell thinks coherentism permits seems *wholly other* to Davidson’s approach which insists on the primacy of triangulation.

As we might have expected, Davidson later rejected the use of the term “coherentism,” saying, “My emphasis on coherence was properly just a way of making a negative point that ‘all that counts as evidence or justification for a belief must come from the same totality of belief to which it belongs.’” If this was Davidson’s main point, then,

⁴⁰ Davidson, *Myth of the Subjective*

⁴¹ Davidson, *A Coherence Theory of Truth and Knowledge* in *Subjective, Intersubjective, Objective* pg. 141

⁴² John McDowell, *Mind and World*. Harvard University Press: 1994

rather than saying, “nothing can count as reason a reason for holding a belief except another belief,” he *ought to have said*, “Anything’s counting for a reason for holding a belief has to be understood as having this normative significance within the context of our totality of beliefs.” If he had said the latter, it would not preclude objects from counting as reasons. Within the context our general set of beliefs about the world, and our norms of belief-formation, objects themselves become reason-giving. This is the way that Kukla and Lance propose thinking of the reason-giving nature of objects in the world. They write,

The *world* can demand nothing of us and holds us to nothing, but our normative practices are, generally speaking, intersubjective and collaborative, and *other people* with whom we engage in practices are indeed agents who hold us to all sorts of things. Add to this the fact that, as we explored above, our practices are typically essentially object-involving, and we earn the result that other people ensure that worldly objects and events have specific normative significances for us that we cannot simply choose to ignore.⁴³

As an illustration of this point, consider the example of the popular arcade game Whack-a-Mole. In this game, animatronic moles come out of holes, and the player’s objective is to strike them with a foam mallet as they come out. Of course, the player is reliably responsive, on causal grounds to whack the moles when they come out of the holes. But there is also a norm at play. The player *ought* to whack the moles when they come out. This is because she is playing Whack-a-Mole. Her whacking of the moles is *justified* (in a non-epistemic sense) given the context that she is in. Of course, there is nothing *intrinsic* about the moles that exert this normative force; they only attain this normative significance through their contextualization in a game in which it is *correct* to whack them. We say the same sort of thing about how objects in the world provide a rational constraint for belief.

Like the act of whacking in Whack-a-Mole, the act of observing, generally is an act couched in normative social practices. The suggestion then is this: given a certain shared

⁴³ Kukla and

understanding of what objects are and what the standards of rationality are, certain norms for dealing rationally with our encounters with objects are put in place, and as such we can say that objects provide rational constraint for our beliefs. On this account, I might be aware of a rabbit, but the rabbit itself is not what exerts the normative force on me to believe that a rabbit is present. Rather, this normative force is exerted by my particular location in a community with certain epistemic norms. There is of course, a difference between the Whack-A-Mole norms and perceptual norms. I might opt out of playing Whack-A-Mole, and so there would be no sense in which I would be bound by the norms of the game. But, given that abiding by perceptual norms is essential to our status as discursive beings, we cannot opt out of this normative commitment in the same sort of way.

Once we've laid this framework out, it becomes clear that the objection that Davidson problematically *assumes* the existence of the external world to make his arguments against skepticism is misplaced. Our very existence as discursive agents commits us to such a world, and to try to take oneself away from that necessary commitment is absurd. We are necessarily immersed in a world that together we understand and interact with. The very notion of making sense of anything at all requires triangulation, and this requires holding each other normatively responsible to objects in the world that we all receptively encounter.

3. A (not so) Real World Example

But hold on a second, says the Cartesian skeptic. What happens if one of the skeptical scenarios *actually obtains*? What if, for example, we *really are* brains in vats? In this case, it seems unclear how the Davidsonian response is supposed to work. In order to engage in radical interpretation, an interpreter must be in the presence of an *embodied agent* interacting with the world (the *same world* the interpreter inhabits) in various interpretable ways.

Triangulation can only take place in this context. How could one interpret the beliefs of a

brain in a vat that is merely being fed electrical signals? Presumably, if I (living in the real world) come across a brain in a vat, whatever sort of beliefs it may have, I will discern what they are, not by interpreting the brain's actions and applying charity (since the brain isn't an actor at all, this isn't even a possibility for me), but by simply seeing what stimuli the brain is being fed. It seems the possibility remains that the brain could be massively deluded, and so it seems possible that *I* could be in that circumstance and thus massively deluded.

I regard the most detailed and concrete form of the Cartesian skeptical scenario to be brain-in-vat skepticism or Matrix skepticism in which we are immersed in a virtual world by way of a computer program stimulating our brains in certain ways, so let's entertain this hypothesis. For lack of a better term, let's call the virtual network, the simulated "world" that we might be immersed in, "the Matrix."⁴⁴ As I continue, I will talk about what people who are located within the Matrix say and do rather than talking about what "brains in vats" say and do, since I think it ends up being more coherent. I'll argue now that there is a way in which, even if we are in the Matrix, the Davidsonian line of thought regarding triangulation still holds, avoiding Cartesian skepticism.

One particularly famous attempt to address the skeptical worry of being a brain in a vat, drawing from considerations on meaning related to those of Davidson, has come from Hilary Putnam.⁴⁵ The thought goes like this: Suppose a person in the Matrix says "There's a cat there." If their statement is to come out true and there are no causal connections between whatever they are calling "cat" and physical cats in the physical world, then they must be talking about something other than the things we call "cats" here on Earth. What

⁴⁴ We can think of it as a virtual world called "the Matrix" as in the movie *The Matrix*. Of course, if it is a virtual world we are in, "the Matrix" as I am using it, is not the same world as the one in the movie, since that world is fictional.

⁴⁵ Putnam, "Brains in a Vat"

could they be talking about if they're not talking about a cat? Putnam gives three suggestions as to what the word "cat," as spoken by someone in the Matrix, is referring to which do have a causal connection to their utterance. The word could refer to (1) cats "in the image," (2) electrical impulses in the brain that cause cat experiences, or (3) features of the program that are responsible for those electrical impulses.⁴⁶ A "cat in the image" is presumably something like a sense impression. I will take the "features of the program responsible for the electrical impulses," to be something like the raw input from the program into which causes the electrical impulses in the brain. This may be analogous in the physical world to the stimulation of the visual receptors on one's eyeballs.⁴⁷ Putnam's suggestions seem unsatisfactory, however. All of these suggestions are very far conceptually from the things we (presuming we are not in a vat-network) are referring to when we are using the word "cat," namely *cats*, small, usually furry mammals of the family Felidae.

If we take any of Putnam's suggestions to be the meanings and causes of the words of people in the Matrix, the account of triangulation that Davidson gives, which is purportedly conceptually necessary, is lost from the picture entirely. A central aspect of triangulation and the direct realist extension of it that I've presented here is the notion of shared *perception* of objects in the world and shared *response* to these objects. Since we don't, in fact, perceive electrical impulses in the brain (whether it's a brain in a vat or in a human head), nor do we perceive the computer program accounting for these impulses, these things cannot be the shared stimuli of triangulation. It is also the case that, if we want to keep

⁴⁶ Putnam, 394

⁴⁷ There is another way that one might want to understand "program feature" which will accord with what I want to say a person in the Matrix's words refers to, but, since I am pretty sure it is not what Putnam has in mind, I will leave it aside for now

epistemological intermediaries out of our theory of perception, we must not perceive, in the central instance, sense data or anything like “objects in the image.”

It is important to recall that Davidson states, “The causal connections between thought and objects and events in the world could have been established in entirely different ways without this making any philosophically significant difference to the contents or veridicality of perceptual belief.”⁴⁸ Certainly, if we are in the Matrix, there are extra causal elements to a full story between beliefs and the world, and so the full causal story is quite different than we currently believe it is, but, if Davidson’s right, this shouldn’t be catastrophic either epistemologically or semantically. On Putnam’s view, however, it is quite catastrophic. Consider the sentence, “My aunt has three cats.” If we take any of Putnam’s suggestions as to what the word “cat” refers to, this statement is going to come out absurd. My aunt undoubtedly does not own “cats in the image” or sensory impressions. When she tells her neighbor, “Take care of Fluffy,” she is not telling her neighbor to take care of her sense impression which she has strangely named, nor is she telling her neighbor to take care of her electrical impulses. She is telling her neighbor to take care of *Fluffy*, something which, if not a cat, ought to be *very much like one*.

The solution to this issue is to find some set of objects that *do* exist in the Matrix and that people in the Matrix would have beliefs about, such that, if *we* were these people and our beliefs were about these objects, we wouldn’t be massively wrong about the sorts of things our beliefs are about. Offering us something just along these lines, David Chalmers has made the case, quite convincingly I think, that the objects of belief for people in the Matrix are *virtual objects*.⁴⁹ That is, when a person in the Matrix has the belief “I am petting a

⁴⁸ Davidson, “The Myth of the Subjective” 45

⁴⁹ David Chalmers, “The Matrix as Metaphysics.”

cat,” they have a belief about a virtual cat. Virtual objects are, in fact, very similar to non-virtual objects but differ in the fact that ultimately they are not composed of tiny bits of matter but are produced by computer code.⁵⁰ So, when someone in the Matrix says “cat,” it refers to a virtual cat, and, since there *are* virtual cats in the Matrix, the sentence “My Aunt has three cats,” may very well be true when uttered by someone in the Matrix. On this line of reasoning, even if we are in the Matrix, our beliefs about the objects around us are still mostly true, and we still have a basic grip on what sorts of things these objects are, even if we’re wrong about the underlying metaphysical structure. Chalmers, accordingly, calls the Matrix hypothesis a *metaphysical*, rather than skeptical hypothesis. He sees the hypothesis as not categorically different from fundamentalist religious hypotheses, a “creation myth for the information age.”⁵¹

Looking at things in the way Chalmers proposes, rather than the Putnamian way, allows us to think of the causal relations between the objects our beliefs are about and our utterances in the way Davidson’s account of meaning demands. In order for triangulation to make sense in a virtual world, it has to be the triangulation *within* the virtual world, and not the triangulation within the material metaphysical ground of the virtual world, and this is the mistake that Putnam makes in thinking that the meanings of words in the Matrix might be something like electrical signals in the brain. However we want to cash out the metaphysics of virtual causation, even if one wants to say that it’s some form of “pseudo-causation,” it is *this* sort of causation that accounts for meaning in a virtual world. Supposing I grew up in

⁵⁰ The distinction between virtual and physical objects might be a bit imprecise, since, there is some sense in which the virtual objects of the matrix can still be conceived of as physical objects. They share all sorts of properties with the objects we call physical objects, such as casting a shadow and having mass. Perhaps the better term to contrast virtual objects with would be *terminally* physical objects, in that, for a non-virtual physical object, it’s physical stuff *all the way down*.

⁵¹ Chalmers, “The Matrix as Metaphysics”

the Matrix, I learned to say the word “cat” in response to virtual cats and the same for all other objects. Through this process, and having functional (though virtually instantiated) perceptual capacities, I ended up having very reliable perceptual abilities when it comes to perceiving virtual objects. Thus, whether or not I am in the Matrix, I have a basic grasp about what’s going on in the world around me, in much the same sense.

Stopping the argument here, and saying, *even if we are in the Matrix*, we still have mostly true beliefs, gives us everything we need to avoid global skepticism without taking the further step, as Putnam does, of concluding that we definitely aren’t in a Matrix scenario. I do not think that Putnam’s stronger anti-skeptical conclusion is plausible, but I won’t argue against it here, since it seems clear that it is *unnecessary* as a response to global skepticism. Epistemologically, what concerns us about the skeptical suggestion is not simply that we could be brains in a vat, but the conjunctive statement that we could be brains in a vat, and *if we were, all of our beliefs (or the vast majority of them) would be false*. If that latter conditional is false as I have just argued it is, then the worry of global skepticism goes away.

This isn’t to say, of course, that whether or not we’re in the Matrix is just *no big deal at all*. I’d be quite worried if I had reason to believe we were all in the Matrix. But the worry here is *metaphysical*, and perhaps *existential*, not a global *epistemological* worry. If we had some evidence to believe that it’s just as likely we’re in the Matrix as not in the Matrix, we’d have to be agnostic about the *metaphysical nature* of the world we’re all living in, but we wouldn’t have to be agnostic about whether or not all of our beliefs about this world are false. Global metaphysical agnosticism does not entail global epistemological skepticism. Fortunately, however, once we have dispelled the worry of global skepticism, we have enough epistemic friction to reasonably judge the probability, based on what we *do* know about the world, of our being in the Matrix. Some, notably Nick Bostrom, has argued that the probability of us

not being brains in vats, but being *completely* virtual, is in fact quite high. There are of course many responses to Bostrom's argument, and this debate has many interesting implications for computer science and probability theory. Given what I've said here, however, our *epistemological* outlook will not have to wait on the way this debate ultimately turns out. I'm relatively confident it will not turn out with everyone conceding that we are probably in a virtual world, but I won't give my reasons for that confidence here, since, for the purposes of this thesis, I don't need to.

It is now clear how an argument from radical interpretation can make sense as a response to a skeptical concern. The very fact that interpretation is *coherent* with respect to our beliefs shows that there is *some world* (however the metaphysics end up working out) about which we have mostly true beliefs. And if we reflect a bit, *just look at us!* We are embodied agents interacting with the world in various interpretable ways. Now that we've gone through how this sort of ostension to ourselves and the world is a necessary requisite of our status as discursive agents and our ability to make sense of anything at all, and we've shown how this ostension *still* makes sense in the various skeptical scenarios, it no longer simply begs the question against the skeptic.

5. A Shift in Skeptical Strategy

But the skeptic isn't done yet. At this point, he might shift strategies, moving from a *Cartesian* skepticism, which attempts to positively put forward doubt by imagining skeptical scenarios, to an *Agrippan* skepticism which attempts to cast doubt on our epistemic practices generally.⁵² He might say "But how do you know that any of *that* is true? What reason do you have to believe that anything at all you've just said is correct?" The goal of the Agrippan

⁵² These two forms of skepticism have been distinguished in this way by Michael Williams. See *Problems of Knowledge* and "Responsibilism and Reliabilism."

skeptic is to push us back into the *epistemic regress problem*. We might set up the problem as follows:

1. I want to say my belief that P is justified.
2. If this is the case, then there must be some other belief (call it Q) which gives me reason to believe that P.
3. But then I need *another* belief (call it R) that gives me reason to believe Q and so on ad infinitum.

Given this problem, according to the Agrippan Skeptic, one is forced into a trilemma. The three horns of the trilemma, as Michael Williams puts them, are *The Mode of Assumption*, in which one dogmatically assumes that he cannot be wrong about a given belief, *The Mode of Circularity*, in which one goes circles back to a claim in attempting to justify it, and *The Mode of Infinity* in which one keeps trying to offer new justifications ad infinitum.⁵³ We might correlate biting the bullet with respect to each of these horns with the following three theories of justification:

1. Foundationalism: the idea that at some point this regress terminates with basic beliefs for which we do not need to give reasons. The problem was originally posed as the main motivation to hold a foundationalist picture of justification.
2. Coherentism: the idea that reasons can loop back on themselves, that justification isn't completely linear, and a belief is justified just in case it fits into the most coherent justificatory *network*.
3. Infinitism: the idea that justifications of this sort can unproblematically go on forever.

While people have defended versions of all of these three views, I do not think we need to pick one.⁵⁴ Rather, I want to point out that there is another line of thought that rejects the formulation of regress problem as having basically gotten things backwards. The idea is this:

⁵³ Michael Williams, *Problems of Knowledge*. pg. 62.

⁵⁴ As far as I know Peter Klein is basically the only person who defends infinitism. See Klein, "Human Knowledge and the Infinite Regress of Reasons."

the very notion of justification only makes sense in light of our *practices* of justification, and at a certain point, any justification will just come down to what our practices actually are. This is a point that we can find at some places Wittgenstein's writing. In the *Philosophical Investigations*, Wittgenstein writes, "Once I have exhausted the justifications, I have reached bedrock, and my spade is turned. Then I am inclined to say: '*This is simply what I do.*'"⁵⁵ There are various ways of reading this passage, but one way we might read Wittgenstein here as saying that, after a certain extent of justification, all we ought to do is just refer to what the practices that we are mutually engaged in, which make knowledge and justification possible at all, actually are. We cannot, of course, get outside of our practices in order to justify them. But we can, from inside our practices, explain *why* they must be mostly truth-conducive, and that's what I've tried to do with the arguments Davidson. Even stronger, we can say that we *must* think of our practices as mostly correct in order for them to make any sense at all. This would be to say that, by the very pragmatic structure of how our reason-giving practices work, our inferential practices and accordingly our beliefs must be treated as mostly be right. Thus, in response to the regress problem, we can say that there is no reason to think that every belief *needs* an explicit justification, since that's just not the way our practices of reason-giving work.

Recall now Brandom's default and challenge model of entitlement that I discussed in the previous chapter. It was argued that we are attributed entitlement for the majority of our commitments by default. That is, the default-status of our commitments is entitlement to them. Without treating discursive agents as largely entitled to their commitments, there is no way to get assertional discursive practice up and running. For an agent to commit herself to something, to take it as true, is, in the paradigm case, to entitle others to the same

⁵⁵ Wittgenstein, *Philosophical Investigations* pg. 91e

commitment. If I say, “The Red Sox won the game today,” it’s part of the pragmatic structure of my speech act, part of what makes it an assertion, that others can take it as true when they hear me say it. Asserting would lose all of its pragmatic force if, in the standard case, the speaker was not entitled such that the entitlement to the assertion could be passed on to those who hear it. Assertions can be challenged in this practice, but a challenge to a default-entitled assertion must positively do something to secure *its own* entitlement. If entitlement wasn’t the default condition for treating claims, and if we doubted *literally everything everyone said*, there’d be no common ground from which we could make sense of any discussion at all. Let me give an example to make this clear. In order to play soccer, for example, it is not enough that everyone else just be playing soccer; we have to *think* everyone else is playing soccer as well. Only then will following the rules of soccer make any sense to us. Likewise, in order to engage in rational discourse, it is not enough that everyone else just be rational; we have to *think* everyone else is basically rational. To try to make someone give reasons for their beliefs without presuming that they are basically rational is to verge on incoherence. It’s like trying to enforce a hand-ball penalty without presuming that the person is playing soccer!

The Davidsonian line of argument that I’ve defended here supports the notion that Brandom’s picture of entitlement can be viewed as *more* than a description of the pragmatics of our epistemic practices to which we must necessarily conform; if Davidson’s right that our beliefs must be mostly true, then this justifies Brandom’s picture as a source of *genuine epistemological insight*. According to Brandom’s pragmatics, most of our beliefs are, and must be, by the very nature of discursive practice, attributed entitlement. But of course entitlement is not simply a *descriptive* claim about the attitudes taken by the members of the discursive community, but a *normative* claim about the epistemic standing of a claim. We can,

of course, *wrongfully* attribute entitlement. For example, a student might raise his hand in a math class, saying “I know the answer,” and, by luck, get the answer correct by performing the wrong operation. In this case, he will be *attributed* entitlement to the claim. However, he *ought* not to be attributed such entitlement since he performed the wrong operation and the likelihood of having come about the right answer is quite low. Here, we might say that, though entitlement has been attributed to his claim, he is not *really* entitled to it. On the solely pragmatic grounds of Brandom’s view, we are only in a position to make the descriptive claims that the majority of beliefs, by default, must be *attributed* entitlement, not the normative claim that believers are *actually entitled* to the majority of their beliefs. However, Davidson’s argument that the beliefs of any language-speaker must be largely true shows that this prima facie entitlement *attribution* is *actually appropriate*. It turns out, if Davidson’s right, that the vast majority of beliefs *attributed* entitlement in our existing epistemic practice are *actually entitled*. Thus, not only are the majority of our beliefs true, but the majority of our attitudes towards these beliefs are normatively appropriate. All of this amounts to the fact that our epistemic practices, as the stand, are generally good ones.⁵⁶

The conjunction of Wittgenstein, Brandom, and Davidson now gives us a legitimate answer to the Agrippan skeptic. We can sum up the response as follows: 1.) We cannot jump outside of our practices in order to justify them, since the very notion of justification only makes sense within the context of our practices (Wittgenstein). 2.) The structure of our epistemic practices is such that, in order to engage in the practice of giving and asking for reasons we must treat each other as basically right about the way the world is (Brandom). 3.)

⁵⁶ I’m not the first to employ Brandom’s default and challenge model of entitlement as a response to skepticism. Notably, Michael Williams has adopted it in his approach to skepticism, and Miranda Fricker has expanded on his attempts. However, without some independent reason to dismiss the Cartesian challenge, as I’ve just attempted to give, I don’t think this sort of response has any weight.

We can justify from *within* our practices, the only sort of justification we can possibly hope for according to (1), why treating each other as having mostly true beliefs is not only pragmatically necessary but actually correct (Davidson). And thus, the regress problem cannot even get off the ground.

6. Finally Telling the Skeptic to Get Lost

Suppose the skeptic *still persists*, now, saying, “Well how do you know any of *that* is true?” Though this is a bit frustrating, we can now respond to this further challenge as well, silencing the skeptic for good. Now we can tell the skeptic that, by abandoning the default and challenge model of entitlement and asking for justification, even after justification has been given, he is abandoning the way the norms of justificatory practice actually work. We might think of epistemic practice as a game, and if you break the rules enough, you’re no longer really playing it. Unless the Agrippan skeptic plays by the epistemic rules and accepts justifications which connect a particular belief in question to our default-entitled set of beliefs, the very notion of justification, and thereby challenge, becomes meaningless. Since he breaks all the rules of epistemic practice, particularly the default and challenge structure of entitlement, the Agrippan’s speech acts will eventually lose all of their normative force.⁵⁷ A genuine challenge, in actual epistemic practice *does* have normative force, and an agent whose commitment is genuinely challenged has an obligation to defend that commitment, but the

⁵⁷ For a particularly interesting real-world example of this happening, consider the popular Christian apologist Sye Ten Bruggencate. Sye argues that that the alternative to belief in God is the inability to know anything at all, and his argumentative strategy is to simply repeat the skeptical challenge ad nauseam. As a result, especially against his more articulate interlocutors, you can literally see his words starting to lose their force as his attack persists. If you watch any of Sye’s interviews you can quite vividly see this happening. His interview with the scientist P.Z. Meyers is a rather good example, which can be seen here: <http://www.youtube.com/watch?v=oIp9OGnnfrA>

skeptic's speech acts stop being treated as genuine challenges, and thus his interlocutor eventually feels no normative obligation to answer them.

Given that we've just seen how the default and challenge structure of entitlement is justified, we can now say that they *ought* to lose their normative force as they do, since there is no coherent challenge to be made if an appeal to our default-entitled set of beliefs will not be accepted as an answer to this challenge. We can make all this *explicit* to the relentless skeptic, and say that's why his challenges *stop being the speech act he intends them to be*, no longer counting as genuine challenges.⁵⁸ Thus, in telling the relentless skeptic to get lost, we are saying to him that he is contributing nothing to conversation but *mere words* with no more normative force behind them than the barks of a dog. By violating all the epistemic norms, he is subjecting himself to a sort of *epistemic sanctioning* in which he ends up no longer being treated like a player in epistemic practice.

Once this is made explicit, I don't think we should be surprised with this result. In fact, the idea that our core epistemic practices might *not* be good ones is utterly incoherent, since the very idea of justification and warrant only arises with respect to our actual practices. What would a sort of justification be like that doesn't connect, even in principle, to our understanding of the world? For this reason the fact that we cannot get out of our actual epistemic practices in order to justify them should not be a concern. Since we can explain why our epistemic practices of belief-formation must, by their very structure, be truth-generative and our practices of entitlement attribution must be generally correct, we have all the grounds we need to be secure that we're not massively deluded.

⁵⁸ This is to say that, if the skeptic continues endlessly attempting to challenge any justification, his speech acts not only lose their *perlocutionary* force (in making one offer up a response), but their *illocutionary* force as well. The speech act is infelicitous *as a challenge* in the same way that running up to two random people on the street and saying "I now pronounce you man and wife" is infelicitous as a marriage pronouncement.

7. Getting a Bit Meta

What sort of justificatory strategy was employed here in responding the skeptic? Is it a version of foundationalism, citing direct perception, a version of coherentism, relying on all of our justifications hanging together in a coherent picture, or is it an externalist response, emphasizing a causal connection between beliefs and objects in the world? Rather than picking one of these labels, my claim here is that no particular label is sufficient to neatly categorize the different sorts of moves at play here. The best label, I believe, is that I've simply tried to *discursively justify* the claim that our beliefs are largely correct against the skeptic's various challenges, and several different justificatory strategies are employed to justify our standard everyday justificatory strategy, and externalist, coherentist, and foundationalist elements are appropriate at various places in the response to skepticism offered here. It is important to note that the sort of multi-layered justificatory strategy going on here isn't the standard way we justify commitments in everyday discursive practice, since, in an epistemological context, we are attempting to respond to challenges that are not usually appropriate. Discursive justification in everyday practice, following the default and challenge model of entitlement usually ends up looking quite a bit like a form of modest foundationalism. In normal discursive practice, we have a level of default-entitled beliefs that do not demand any sort of inferential justification from other beliefs but from which we can justify other beliefs inferentially. But the argumentative structure in this chapter ought to show that the justificatory picture we are working with here is more malleable than modest foundationalism, since it allows us to answer a challenge to our default-entitled beliefs, inferentially justifying the majority of the set as a whole. And, on the picture I am advocating here, *only by* performing this explicit global justification can we be

epistemologically entitled to the modest foundationalism that is normally at play in discursive practice.

I argued in chapter two that our ability to engage in shared ostension is essential to our status as discursive beings. Given the arguments in this chapter, we can now say that, regardless of whatever metaphysical status the objects to which we are able to ostend might have, the very fact that this ostension is possible indicates that we *do* occupy a shared world about which we can have various true or false beliefs. The ostensive “Lo!” acts that Kukla and Lance have called *observatives* play a central epistemological role in the account of justification that I will lay out in the next chapter. In targeting these sorts of speech acts which essentially rely on a social intersubjective context as having this central role, the account itself relies on a social intersubjective context as well. Once again, without the arguments in this chapter, this way of thinking about justification could not satisfactorily replace the traditional first-person internalist line of thought, but, with the threat of skepticism dispelled, there is nothing stopping us from thinking of justification in this way. Articulating an account of discursive justification which is epistemologically satisfactory through and through is my task for the next chapter, which I turn to now.

Chapter Four: The Seesaw Model of Discursive Justification

In the previous two chapters I argued that our epistemic practices systematically lead us towards agreement on the way the world really is. I also argued, at the end of the previous chapter, that this suggests a move from traditional foundationalist or coherentist theories of justification to thinking of justification first and foremost as it actually takes place in discursive practice. In this chapter I offer an account of discursive justification with respect to empirical knowledge. This account attempts to systematically tie together the central sorts of moves that one can play in justifying, and to show how agreement which is grounded in the way things really are can be reached by these justificatory moves. I call it (in lack of a better metaphor) the “Seesaw Model of Justification.” The approach attempts to connect causal justificatory moves with inferential justificatory moves and thus explain how we can expand and revise our understanding of the way things are such as to find common ground with others, all while staying connected with the way things really are—never losing our touch with the world.

I argue for two fundamentally different sorts of justificatory moves: one which ties the *physical and social state* of having a particular belief *causally to the world*, and one which ties the *content* of a belief *inferentially to other beliefs*. These two sorts of justifications fit together like the two sides of a seesaw, and our noninferential perceptual beliefs about the world around us function, in the paradigm case, as the *fulcrum* of the seesaw, the point at which justification switches from *causal* to *inferential*, and vice versa. Justification, on this model, like a seesaw, moves in two directions. The first direction assures that our beliefs from which we infer and *expand* our knowledge are connected to the world. The second direction assures that the beliefs which we take to be causally connected to the world are always revisable in light of new beliefs. The beliefs which we take to be our ties to the world can

always be thrown through the inferential web of belief, and thus we have the tools to *alter our starting point* to find common ground with others and avoid the worries of incommensurability attributed to Kuhn and Feyerabend.

1. Two Kinds of Justificatory Moves

My account here owes a great deal, especially terminologically, to Susan Haack's *Evidence and Inquiry*. In her book, she presents a dual-aspect account of evidence in which casual states and relations, as well as logical propositions and relations, play a necessary factor. This chapter reformulates how these two aspects fit together to form a comprehensive picture of *justification*. Haack's account is an account of *evidence*, not justification, and so, though the two notions are very closely related, our approaches are quite different. Nevertheless, I will employ her distinction between state and content senses of belief, though I think that a slight change in terminology can clear up some confusion. Haack makes the distinction between the "state and content senses of 'belief'". She introduces the terms S-belief and C-belief to make this distinction between "someone's believing something and what they believe."⁵⁹

Talking of state and content beliefs as two senses of belief seems to me to paint a misleading picture. It seems much more intuitive and clear to me to, rather than speaking of state and content *senses* of belief, to speak of state and content *aspects* of belief. We can say then that all beliefs have both an S-aspect and a C-aspect. One aspect of a belief is the actual physical, mental and social situation *of believing*, that a believer is in when holding a belief, which is causally affected by events in the world and which affects behavior. We can describe this state of having a belief in various ways, with various degrees of normativity built into the description. Quine, whose account of justification will be discussed shortly,

⁵⁹ Susan Haack. *Evidence and Inquiry*. Prometheus Books: (1993). p. 118.

keeps it purely in physiological terms, and identifies the belief-state with a physiological disposition to behave in a certain way. This has the benefit of being able to stand without worry in causal relationships with objects and events in the world, but the more we try to keep the state of having a belief in purely physiological terms, the harder it becomes to see this as *actually the state of having the belief* and not merely a *causal necessity* for that state. Since, on the Brandomian picture we are working with here, belief is essentially a normative notion, we ought to keep this normativity built into what we identify as the state of having a belief. Thus, the state of *being committed* to the truth of a particular proposition, the sort of commitment I discussed in chapter two, is what I will primarily identify as the S-aspect of a belief.⁶⁰ The other aspect is the *content* of the belief, which can be expressed propositionally and which the belief-state marks a commitment to. So, if I believe that Paris is in France, the S-aspect of this belief is the physically and socially instantiated mental state of my being committed the truth of Paris being in France and the C-aspect is the proposition “Paris is in France” itself.

Corresponding to these two aspects of belief, we can discern two sorts of justification that one can offer for a belief: one which latches onto the propositional content of that belief (the C-aspect), and one which makes reference to the mental, social, and physical state of that belief (the S-aspect). First, when asked to justify a claim or belief, the standard strategy is to show how that claim can be inferred from a different claim (or set of claims) that does not need to be justified (at least not in that present situation). This is usually how a defense of a claim will work in the situations that first come to mind, such as scientific or political issues that are called into question by another person. I claim X,

⁶⁰ Of course, this is not to deny that certain physiological states that must obtain for this commitment to obtain.

someone challenges X, and I go back to P, which is unchallenged at that point, and show how, from P, X can be inferred. If I am defending a scientific claim, for example, my task is to show how the claim in question can be inferred from experimental evidence or some other scientific claim that both my challenger and I agree upon. If I am successful in showing how the inference can be made, then my challenger must either change her initial position of agreement with regard to the experimental evidence, or rescind her challenge of my belief. In this standard sense, what is being called into question is the way the C-aspect of the belief stands in relation to the C-aspect of related beliefs. The justification here is an appeal to the way the content of the belief in question, expressed propositionally, stands in logical relation to other propositions which we take to be true. This sort of justification, we can call, using the same terminology, a C-justification.

On the other hand, instead of looking at the relation between the propositional content of a belief and the contents of other beliefs, we can look at the situation in which a given belief is produced, the environment and the faculties producing it, and say that the belief is justified if a true belief is likely to come about in such a situation. This sort of justification rests on the fact that we are physical organisms in a physical world and there are causal factors that go into belief-production. A certain state of the world will reliably cause a belief state with a certain true content that goes along with that state. It is still the C-aspect of the belief getting justified here. Belief-states *themselves* do not get justified because it only makes sense to justify something with propositional form. However, in this sort of justification, the C-aspect of the belief is not being tied to related contents, but tied to its related S-aspect, and this S-aspect is in turn being tied, causally, to other states in the world. A justification of a belief which follows roughly along these lines, I will call an S-justification.

Traditionally, the internalist theories of justification paradigmatic of the epistemological tradition have dealt primarily with C-justifications. Classical foundationalist theories, starting with the Cartesian project, have started with belief-contents that they took as indubitable and concerned themselves with how the contents of other beliefs could be inferred from those foundations. Coherentist theories have concerned themselves not with a one-way justification of contents, but how the set of belief-contents as a whole fit together. The psychological or physiological S-aspect of the beliefs in question and the causal story of belief-production have traditionally not played a central role in such internalist theories of justification. On the other hand, S-justifications, which tie the truth of a belief with causal factors regarding its production and the environment in which it is produced, is primarily what externalist theories of justification have latched onto. Perhaps the most widely acknowledged of these externalists theories is Alvin Goldman's reliabilism, but Alvin Plantinga's Proper-function account of warrant is also a good example of an externalist theory of justification. The general principle behind the account can be put as follows: If you're designed such that in situation S you produce true belief P, then if you are in situation S and produce true belief P, and you're working in the way that you're designed (if you're not broken in some relevant way), then you know P.⁶¹

To play the role of S-justification in this account, I will employ a naturalized version of Plantinga's account where "designed" can mean designed by evolution and, just as importantly for the present view, *designed by a specific community by conditioning through language-learning (or special training)*. "Design" is kept particularly loose here, but an important aspect of design is that it is a fundamentally normative notion, and this is essential if we are retaining the normative notion of *commitment* in targeting what the S-aspect of a belief is. I

⁶¹ Alvin Plantinga, *Warrant and Proper Function*. Oxford University Press: (1993).

have identified the S-aspect of a belief as the mental state of commitment, modeled directly on the social status of being committed, and, though it is the state of having a certain normative status, it can stand in causal relations with physical objects and events unproblematically, in the same way the normative status of owing my Mom twenty dollars is causally related to the physical act of me taking the money off the counter. Still, there must be a normative context in which this physical act can be causally connected to the normative status. In the example of owing money, it is the general rule under which I am bound that if you take someone's money you ought to pay them back. Plantinga's proper-functionalist account gives us this normative context by which we can causally connect the physical state of an object in the world to the normative mental state of having a belief. We are "designed" through evolution and language learning to form certain commitments about the various objects and events in the world in various circumstances. If I form the perceptual belief that there is a red object in front of me, and there in fact is no red object in front of me, *something went wrong*, since I'm "designed" such that I only form this sort of belief *as a response* to red objects.⁶² So, noting this, in conjunction to the claim I *am* functioning as I ought to, would function as a causal S-justification.

⁶² Often an appeal to reliability or authority given in a discursive justification will walk the line between being a real S-justification, a genuine line of reasoning to believe that something is true, and being an attempt to reject a challenge to entitlement status. Consider a situation in which I am at the nature park, and the park ranger says "That's a young bald eagle," and I ask him, "How do you know that?" He responds, "I've been doing this for years, I'm an expert, and I'm never wrong." On the one hand, we can see this as an effective S-justification, giving logical reasons to believe that the statement is true. If this guy reliably produces true beliefs in situations like these, there's good reason to believe that he's right here. Even though his justification is externalist, it is still evidentialist in that he is giving me good evidence to believe that he's right (assuming his premise of reliability is true). On the other hand, we can see this as a shrugging off of a challenge to his entitlement status, a sort of putting me in my epistemic place, saying, "I shouldn't need to give a justification for my belief in order to be epistemically entitled to it." Either way, his statement is aimed at the same thing: entitlement status. There is no strong line between these two sorts of discursive

2. Indispensability and Interconnectedness of S and C-Justification

One of my central claims is that only together can S-justifications and C-justifications provide a satisfactory account of empirical justification. First, why is S-justification indispensable? Couldn't we justify all of our beliefs on logical grounds, as opposed to appealing to the causal factors involved in the production of our beliefs? I don't think we could. It seems clear to me that a causal component is a necessary addition to the logical component of justification. We are physical animals living on a physical world. We form our beliefs in response to the physical events in the world that leave an imprint on our sense organs. Without making explicit note of this connection between our beliefs and the world, our epistemological picture would have no contact with the world.

In the previous chapter, in response to skepticism, I endorsed the Davidsonian position that meaning and causation are intimately connected. We have been linguistically trained to be able to report the objects and events around us, forming linguistic commitments. According to Davidson's line of reasoning, this sort of linguistic training and the sort of interpretation that goes along with it can only make sense if the vast majority of our beliefs are true. This argument functions as a *global S-justification*. If someone proposes a global challenge to our beliefs, saying that they may be radically disconnected from the way the world really is, we can respond to it in this way, and it roughly fits the model Plantinga sets forward. Insofar as we actually are speaking a language and forming beliefs, we are functioning as we've been "designed" as language users, and, functioning in this way, we're guaranteed to have mostly true beliefs. The sort of global S-justification that I gave in the last chapter also enables us to offer another sort of global S-justification. Once we

moves, and for my present purposes, I will regard such moves as S-justifications, keeping in mind that this ambiguity does exist.

established an objective world about which we must have mostly true beliefs, we can also make sense of the reliability of this connection between our beliefs and the world in an evolutionary way. Dennett has often gestured towards this line of reasoning and it is in fact the explanation he gives for the success of the intentional stance. Once again we can put this S-justification in Plantinga's terms (in fact, it is the sort of proper-function explanation that Plantinga takes to be open to the naturalist) and say that our beliefs are globally justified since evolution "designs" us to have largely true belief and we're functioning in accordance with our biological design-plan.

Employing either this sort of linguistic or evolutionary line of reasoning we can say that our belief states, the S-aspect of our beliefs, are connected to the world. It is the S-aspect of belief alone that makes this connection, and the causal S-justification which latches onto the physical state of the belief, and ultimately ties it to the events in the world that we want to talk about. It is this causal connection that a belief with the content "X" has with X, which ensures that the arising of a belief with content "X" is a good indicator that X actually is the case. Without this causal link, no epistemological picture can be complete. It would be a picture where our beliefs would be floating, unconnected to the world or the things our beliefs are about. In keeping with the default and challenge model of entitlement, global S-justifications usually remain *implicit* in everyday discursive practice, only becoming explicit in discursive justification in theoretic, often philosophical, discussions. However, it is necessary that a large-scale S-justification *can be* made explicit in order for any justificatory picture to be complete.

If our S-justified beliefs can go on to C-justify other beliefs, our theoretical commitments will be tied to the world. This is the first direction of justification, where we make the move from state-justified beliefs to content-justifications, and I will explain, with

reference to Quine, how this can go in the next section of this chapter. But why do we need C-justifications in the first place? An S-justification can only be a good S-justification with the condition that one is designed to produce the belief “P” in situation S, and, in situation S, “P” is true. We have learned to respond to stimuli with the production of beliefs, but this learning is done in accordance with what a linguistic community takes to be the right set of beliefs, and there needs to be a way to show that these beliefs are true. Attesting to the fact that these are in fact true beliefs cannot simply be done with more S-justifications, but requires an inferential C-justification. As such, a mass S-justification of the majority of our beliefs, either of the evolutionary sort that Dennett offers or Davidson’s linguistic sort will get its force on C-justificatory grounds. We come to justify these arguments inferentially, inferring from the content of the things we already know to these conclusions. This is, after all, what I attempted to do in the previous chapter. The global S-justification I gave had to be C-justified *itself*, drawing inferentially from arguments about the nature of meaning and discursive practice.

This second direction, where the move is made from C-justified beliefs to new S-justifications, also provides the way in which our knowledge *grows and is revised*. Our web of beliefs, tied together with broadly logical relations, can form the backdrop of knowledge by way of which we can be conditioned (through ordinary language learning in the standard case or through special training in a scientific community) to reliably and noninferentially form new beliefs. New S-justified noninferential beliefs arise out of what was once inferential and contested C-justifications. In this way, our knowledge can grow. We also can *revise* our previously S-justified noninferential beliefs by subjecting them to C-justification. The latter portion of this chapter will provide a sketch of how this works.

3. The First Direction: State-Justified Beliefs to Content-Justifications

There are many different types of cases in which we move from S-justified beliefs to C-justifications. For example, I might believe that Bob Dylan is in town because a friend told me that he is, and, on these grounds, come to believe that he is not in the Bahamas. Though testimony is one way that a belief can be S-justified, the sort of S-justified beliefs that I am centrally concerned with in explicating this first direction are *observational* or *perceptual* beliefs.⁶³ Ultimately, if we keep trying to give C-justifications for our beliefs, it will come down to perception, what we observe in the world with our five senses. We need a way to tie justifications to both the propositional *content* of an observation (the inferential force of an observation) and the *state* of making that observation (the causally affected aspect of an observation). Thus, we need observations to function as a *fulcrum* of our seesaw, a point in which an epistemologist can identify the transition from causal physical factors to inferential ones, where S-justified observational beliefs, connected with the world, can go on to inferentially C-justify theoretical beliefs, and, in this way, connect theory to the world.

In his late epistemological work, Quine illustrates just how such a state/content fulcrum can work with his talk of the “observation sentence,” an occasion sentence that all members of a linguistic community can agree outright upon when witnessing the occasion. On Quine’s account, if we think of the “observation sentence” as a holophrastic noise, rather than in terms of the meanings of the words it employs, we can see it as something to which we are conditioned to assent to outright when certain sensory stimulation is present. In this way, thinking of it in purely physical and physiological terms, it is casually connected (and reliably connected) through conditioning to the physical events in the world. However,

⁶³ We must note that, strictly speaking, observations *themselves* are not beliefs, but they do commit us to what we might call “observational beliefs,” the beliefs that we noninferentially form about the things we observe when we observe them. This functions in the same way that on Kukla and Lance’s account, an observative like “Lo, a rabbit!” is not *identical* to the declarative claim “There is a rabbit there” but it does commits us to that declarative.

if we think of the sentence term by term, its content connects logically to theory. For this reason, Quine calls the observation sentence “Janus-faced.”⁶⁴ It is both a causally-influenced physical event and an inferentially productive proposition.

Though it serves its role on Quine’s account, thinking of an observation purely in physiological terms is not epistemologically essential. What is essential is that this fulcrum does not require an explicit inference and can be causally connected directly to the world, but that it is something *from which* one can make an explicit inference. This is why I have described it as a fulcrum; it is a point that marks the transfer, the shift in weight, from causal factors to logical ones. If we use this Quinian strategy of focusing on linguistic items and their connection to purely physiological states for the S-aspect of the fulcrum to be connected to an S-justification, it makes it difficult for us to get the normativity we need for our proper-functionalist S-justification to work. Thus, rather than Quine’s take of the “observation sentence” for the role of this state/content fulcrum, I will take a more pragmatically and normatively robust observatory speech act for my fulcrum. In chapter two, I introduced Kukla and Lance’s notion of an “*observative*,” such as “Lo, a Rabbit!” that doesn’t just express that there is a rabbit present, but also functions “to discursively mark and communicate the event of recognition itself.”⁶⁵ Unlike traditional declaratives, these speech acts are essentially *indexed to the speaker*—that is, they essentially incorporate the speaker’s first person voice. Kukla and Lance base their account on distinction between the *input of a speech act*, what would entitle one to make the speech act, and the *output of a speech act*, the normative statuses the speech act strives, as part of its function, to bring about. They argue that, while a declarative (in the traditional sense) has both an agent neutral input and

⁶⁴ “In Praise of Observation Sentences.” p. 109.

⁶⁵ Ibid 46

output, the input of an observative is agent-relative, and the output is agent-neutral.

Entitlement to an observative, the *input* of the speech act, is “*inherently one’s own*,” since we “come to be entitled to a belief on the basis of observation because we *encounter* the world in a certain way.”⁶⁶

Because of these two aspects, observatives as described by Kukla and Lance preserve the “Janus-faced” nature of Quine’s observation sentences while retaining the normativity that I take to be central to the explanation. Thus, thinking of an observatives as Janus-faced we can tie the agent-relative aspect of it causally to the world with an S-justification and the agent-neutral aspect inferentially to other beliefs with a C-justification.⁶⁷ To make an observative like “Lo, a rabbit!” is to make clear that (unless something is abnormal) we are *receptively encountering* a rabbit. Our first-person encounter with the world in a certain way is what entitles us to make certain observatives. Thus, an S-justification, which latches on to this aspect of an observative, ties our situated way of encountering the world with objects and events of the world itself by explicitly articulating a proper-functioning *recognitive ability*. On the global scale, we can say that we know our recognition of things is the recognition of things *as they really are* because we have been “designed,” by the conjunction of evolution and language learning to recognize things as they really are. This is making use of the *global* sort of Davidsonian/evolutionary S-justification, but we can also employ *local* proper-

⁶⁶ Kukla and Lance, p. 50.

⁶⁷ This fits neatly with the distinction between S-justifications and C-justifications as I’ve laid them out. C-justifications themselves are essentially agent neutral and S-justifications are essentially agent relative. If anyone knows all the facts involved in a C-justification, they are entitled to give that justification. The same is not the case with S-justifications, and this can be shown quite simply with appeals to authority, one form of S-justification. I might come to believe a certain fact about the Barn Swallow because the world’s leading expert on Swallows has said it, and citing this as a reason for holding this belief is perfectly legitimate for me, but it’s not a legitimate epistemic move if the world’s leading expert cites this (the fact that he, the world’s leading expert, has said it) as the primary reason for his belief.

functionalist S-justifications. For the sort of basic observatives that all relatively competent language users can and are expected to reliably make (such as, “There’s a cat,” “That’s a chair,” “Here’s a hand”), a global S-justification with the assumption that the case is not an exception to the standard case is implicitly at play. But for observatives of the less basic sort which require some sort of special training that we are not default-entitled to make, or if a potential challenger has some reason to think that we *are* in an exceptional situation, we can give *explicit* S-justifications for our observatives. If we are justifying an observative these S-justifications must make reference to our *recognitive ability*, the agent-relative aspect of an observative.

To make this point clear, let’s look at an extended example. Consider Jimmy, a freshman in a biology class, a class in which he pays very little attention (let’s say he’s a literature major and needs the class for a general education requirement). He looks through a microscope in which a slide is prepared of cells undergoing mitosis. Jimmy has no clue what mitosis is, nor does he have any solid idea of what a cell is, other than perhaps for knowing that animals have them. His classmate comes up to him as he is looking through the microscope and asks him what’s on the slide. “Apparently,” Jimmy says to his peer, judging from his recollection of the label on the slide at which he is looking, “there are cells undergoing mitosis here.” This is a true report, and Jimmy is justified in making it. Now consider Jimmy’s professor, a highly trained microbiologist, who is asked what is on the slide in the microscope. She doesn’t know, so she looks through the microscope, and says, “Oh cool, there are cells undergoing mitosis here.”⁶⁸ Looking only at content, this is the same assertion that Jimmy makes, but when she makes it, the normative function of her speech act

⁶⁸ If these contextual bits, such as “Apparently” and “Oh cool” concern you, note that they are simply to describe contextual factors and could easily be replaced by a non-verbal description of their reactions.

is quite different. Jimmy's speech act is a declarative, whereas his professor's is an observative. When the professor looks into the microscope and says "There are cells undergoing mitosis here," she commits herself, not just to the fact that there *are* cells undergoing mitosis, but that they are *sensibly present to her*, she is expressing her receptive encounter with these cells. By contrast, Jimmy who, while looking through the same microscope, says the same sentence, commits himself to no such report. Judging on the basis of authority that the slide in fact contains cells undergoing mitosis, he simply makes that fact available.

Imagine if the respective claims were to be challenged. The hypothetical challenger knows the circumstances of Jimmy and his professor and asks them "How do you know?" When asked, Jimmy can call on whatever resources are available to him. He may make reference to, as I have set up the scenario, the fact that the slide is labeled "cells going under mitosis," or he could say that the professor told him that was what was on the slide, or he could refer to something else that would justify the claim. The professor finds herself in quite a different circumstance. She can do one of two things. She can justify the fact by making reference to her observational competence. Perhaps she'll say something like, "I'm a trained biologist. I know what it looks like for cells to undergo mitosis, and I'm not going to mistake something else for it." She may also say how the fact can be inferred from an observation considered by the challenger to be more basic. She might say something like, "The squiggles are going to opposite sides of the blobs and the blobs are separating from each other—that's what cells undergoing mitosis looks like." But any justification which makes no reference, at least implicitly, to her perceptual recognition of the cells undergoing mitosis is inappropriate. What she *cannot* do, when her statement (remember how she expresses it: "Oh cool, cells undergoing mitosis!") is challenged is say, "The label on the

slide says it.” If she says this she’s going to look silly, and her original report will have looked like a fraud. The fact that the label on the slide says “cells undergoing mitosis” may justify the agent-neutral fact that there are cells undergoing mitosis as much as anything else, but it is simply not an appropriate justification here. This is because a challenge to the professor’s report isn’t a call to justify the agent-neutral claim that there are cells undergoing mitosis on the slide; it is a call to justify her *recognition* of the cells undergoing mitosis, and that’s why any justification she offers must make some sort of reference to her ability to recognize these things. This seems to make it clear that, whereas Jimmy’s report is simply a report of the fact that cells are undergoing mitosis, the professor’s report is a report of this fact which also, as Kukla and Lance says “bears its recognitive character on its sleeve.”⁶⁹

Both Jimmy’s report and the professor’s report commit us to the same agent neutral fact that there are cells undergoing mitosis on the slide, a fact can be linked inferentially to other beliefs in a C-justification. Only the professor’s observative, however, functions as a fulcrum which can ground empirical knowledge, being S-justified with respect to her recognitive capacities. Suppose there is a very well-supported theory that if you cut off a single layer of an onion root-tip, stain it, and look at it under a microscope there will be cells undergoing mitosis. Of course more evidence can be gained for any theory, no matter how well supported. Jimmy wants to gather more evidence for this theory, so he follows the procedure exactly, and at the end of the procedure he looks at his slice of onion root-tip and, judging by the fact that the theory is well supported, reports, “There are cells undergoing mitosis here.” This obviously provides no more evidence for this theory, but what is wrong here? Jimmy’s assertion is certainly a report of a fact about the world, and it is a well-justified report. The reason Jimmy’s report doesn’t count as additional evidence is because

⁶⁹ “Yo!” and “Lo!” 75

isn't reporting any *new* empirical data; he's simply made an inference from the already existing data.

A causal S-justification roots the content of the observative causally to the world. This causal path that we recognize in our justification as connecting the propositional content "X" with the event X, moves from X, the event happening, through the stimulation of the sensory receptors (as Quine rightly notes), and to the perceptual judgment with content "X". Without an S-justification of the truthfulness of the content of our perceptual judgments, we have no reason to rely on them justify our inferences. Inferential justification has to stop somewhere, and it must make a connection to the world where it stops. If we stop at perception (which, as the paradigm case of non-inferential knowledge, seems like a good place to stop) we need to put the content of perception in truthful connection with the world with an S-justification. This, as I've said, is the first direction of justification, which roots the grounds of our inferences to the world. The second direction of justification, which I turn to now, ensures that these grounds are in fact correct and gives us the possibility of revising and expanding them.

4. The Second Direction: Content-Justified Beliefs to State-Justification

In Plantinga's account of proper function, an essential element is that the design plan be a *good one*. That is, the cognitive faculties have to have been designed, and designed well, to produce *true* beliefs. On a naturalized account, this is quite a big issue that must be dealt with. An S-justification can only do any justifying work if the justifier knows that the causal relations involved lead to *truth*. But if our observational beliefs provide the basis for our inferences, *how can we have the knowledge necessary to S-justify them?* It seems like we have stepped right back into the epistemic regress problem that Plantinga's account strives to avoid. But, because of this second direction, we haven't.

Even though our S-justified noninferential beliefs ground us to the world, it doesn't mean they can't be tested by being thrown into the logical web of beliefs themselves; it just means that if we do this we need a different foothold for the time being. An observative can be called into question on logical grounds, requiring a C-justification rather than its usual S-justification, if different beliefs in the belief web call for it to be challenged. Consider a Kuhnian revolution in which, because a whole set of background beliefs change, basic observational beliefs must change as well. This encourages Kuhn to make the claim the basic experience of scientists changes with paradigm change.⁷⁰ Even if we are taking observatives to be the grounding aspect of our epistemology (connected to the world through an S-justification), this isn't necessarily epistemologically problematic; it merely means that, if a certain observative tendency is being called into question, scientists can't use *that* observational belief as a grounding foothold at that time. This is exactly the point Sellars makes when he says, "For empirical knowledge, like its sophisticated extension, science, is rational, not because it has a *foundation*, but because it is a self-correcting enterprise which can put *any* claim in jeopardy, though not *all* at once."⁷¹ The ground of our justifications is not a "foundation" which gives us an unshakable, indubitable ground for our inferences; it ideally only provides the best conjunction of security and informativity with respect to our belief-structure.

Our cognitive abilities are conditioned to fit with the best set of beliefs that we have at the moment, the set of beliefs which we think are most likely to be true. If we have reason to change this set, we can learn to see the world differently, and, though at first correct observation may require some thought and inference, eventually our immediate way

⁷⁰ Thomas Kuhn. *The Structure of Scientific Revolutions*. University of Chicago Press: (1962).

⁷¹ Wilfrid Sellars. *Empiricism and the Philosophy of Mind*. Harvard University Press: (1997). p. 78.

of experientially encountering the world becomes filtered through the correct set of beliefs. Imagine a child who learns the word “fish,” and recognizes any swimming animal with flippers as a fish. He then sees a whale and calls it “fish,” and, though this recognition may have corresponded to some set of beliefs at some point in history (perhaps before sufficiently developed biology), the best set of beliefs, the one most internally and externally consistent, that we have now classifies this animal as a “whale” which is not a fish but a mammal. So, though it may take some errors and extra thought at first, eventually his basic observative tendencies become conditioned (with the help of his parents) to this set of beliefs. Our perceptual encounters get no more basic than the sort which includes things like “fish”, and there is no epistemological foundation more theory-neutral than our everyday basic understanding of the world which includes all sorts of these judgments. This is unproblematic, however, insofar as we note that this basic understanding is not independent of our understanding of the world, and so the way we perceptually encounter the world is constantly reflecting a more comprehensive and empirically adequate set of beliefs.

Using C-justifications to revise which beliefs are to be S-justified, a growing understanding due to C-justifications expands the scope of our S-justified noninferential beliefs. The complexity and informativeness of our experience grows as we acquire a new theoretical understand of the world. What we have once inferred to from an experience formed by a less theoretical worldview now provides the new bundle of background beliefs that go into a new experiencing of the world. The sort of theoretical commitments that once been at the forefront of research now form the filter through which one sees the world and forms new theoretical commitments. The possibility of this top-down revision traditionally has raised the worry of incommensurability. However, given that my account

here takes Davidson's arguments about meaning and interpretation as a key element, this worry cannot be raised. Though we can revise *some* of our observational beliefs, we cannot revise *all of them*; the vast majority of them must stay intact. This top-down revision, then, is what allows us to *avoid* incommensurability, since it allows our noninferential beliefs to be flexible in response to the world, a world which, if we share beliefs about *any* elements of it with other discursive agents, we must be largely in agreement about.

Chapter Five: From Discursive Justification to Justified Belief and Knowledge

Justification, as I have been urging here, is, in the central instance, not simply an abstract property attributed to beliefs, but a distinct, materially embodied *social act*. To distinguish this act from the traditional conception of justification, I have been specifically using the term *discursive* justification. In this chapter, I want to present an account of justified belief that plays the role justification has played in traditional Justified True Belief (JTB) accounts of knowledge.⁷² This account stems from the considerations of the previous chapters, but, as an account of justified belief, it is importantly different than an account of discursive justification as I've given in the last chapter. Discursive justification is something which we *do*, rather than some property which our beliefs have. Giving a discursive justification is a move we can make in the game of giving and asking for reasons when a claim of ours is challenged. I'll argue here that for a belief to be externally recognized *as justified* is for it to be granted positive epistemic status on the grounds that it *can be* discursively justified against appropriate challenges, even if the believer herself has not performed such a justification.

1. Being Entitled and Being Justified: An Important Distinction

Brandom is quick to note the distinction between discursive justification as an act and justification as a property of beliefs, but he makes little distinction between how the *entitlement* to a claim works and the *justification* that a belief might have:

“In one sense, to call a belief justified is to invoke its relation to the process of *justifying* it. To *be* justified in this sense is to *have been* justified—exhibited as the conclusion of an inference of a certain kind. In another sense, to call a belief justified is to attribute to it what might be called *positive justificatory status*. Positive

⁷² In light of Gettier problems, most accounts of such a sort have also included a fourth condition which excludes Gettier cases from counting as knowledge. I will address these cases at the end of this chapter.

justificatory status is just what has been talked about here in terms of *entitlement* to a claim.”⁷³

On Brandom’s social articulation of knowledge, when someone attributes knowledge they are doing three things: (1) attributing a commitment to an agent, (2) attributing entitlement to that commitment, and (3) undertaking that commitment themselves.⁷⁴ The first corresponds to the criterion of belief, the second to justification, and the third to truth. On this account, it is not discursive justification, but *entitlement* that plays the role traditionally occupied by “justification” in JBT accounts of knowledge. However, it seems unlikely that these two things can be so easily equated. Anthony Quinton writes, “For a belief to be justified it is not enough for it to be accepted, let alone merely entertained: there must also be good reason for accepting is not simply to undertake a commitment and be attributed entitlement for it, but to have *good grounds* to accept that belief.” However, if the default and challenge model of entitlement is the pragmatic structure in which we are working, we face a problem. The majority of our beliefs, in the first place, must not be entitled by our having *good reasons* for them, but by default. And yet, if we want to treat ourselves as really *justified* in holding them and not merely having been attributed entitlement to them, it seems we must have some way of *demonstrating* that these beliefs have positive epistemic status.

In light of this concern, Brandom jumps on the externalist stance of rejecting the traditional line of thought highlighted by Quinton. However, he adds that the *attributer* of knowledge must have good reasons to endorse the reliability of the knower. He gives the following example:

“Suppose that Monique has been trained reliably to discriminate hornbeams by their leaves. As a result of the training, she is often disposed to respond to the visibility of leaves of the right sort by noninferentially reporting the presence of a

⁷³ Making it Explicit 204.

⁷⁴ Brandom, *Articulating Reasons*

hornbeam. She understands what it means to claim that something is a hornbeam and, in circumstances appropriate for such reports, actually comes to believe that there is a hornbeam present. She may still be uncertain of her discriminatory capacity long after she has in fact become reliable. In such a situation she may have a true belief that there is hornbeam in front of her, yet be completely unable to justify that claim (for instance, by citing features distinctive of hornbeam leaves), and even deny that she is a reliable noninferential reporter of hornbeams.

Yet, the reliabilists point out, it can be entirely in order for one who does take her to be a reliable reporter of them, not only to come to believe that there is a hornbeam present on the basis of her report, but to cite her report (at least deferentially) as what warrants that belief. This is treating the claim as authoritative in just the way that is required for knowledge. Someone who thus takes her to be reliable can accordingly attribute to Monique the knowledge that there is a hornbeam in front of her, in spite of her protestations to the contrary. What makes her claim knowledge (according to the attributor) is the fact of her reliability (according to the attributor), regardless of her attitudes toward that reliability.”⁷⁵

Brandom takes this claim to be relatively straightforward. However, it does not seem clear that we *do* want to attribute knowledge to Monique here, except perhaps in a stretched sense. This example is strikingly similar to examples that Laurence Bonjour gives which he takes to be *counterexamples* to externalist theories of knowledge. Consider the following example involving clairvoyance that Bonjour gives:

Case IV. Norman, under certain conditions that usually obtain, is a completely reliable clairvoyant with respect to certain kinds of subject matter. He possesses no evidence or reasons of any kind for or against the general possibility of such a cognitive power, or for or against the thesis that he possesses it. One day Norman comes to believe that the President is in New York City, though he has no evidence either for or against this belief. In fact the belief is true and results from his clairvoyant power, under circumstances in which it is completely reliable.⁷⁶

It seems clear in this case we would not want to say that Norman is justified in his belief.

From his perspective, there is *no reason* to believe that the President is, in fact, in New York.

In response to the reliabilist point that the external relation of reliability should give the belief its justification Bonjour asks, “*Why* should the mere fact that such an external relation obtains mean that Norman’s belief is epistemically justified, when the relation in question is

⁷⁵ *Making it Explicit*, 219

⁷⁶ Laurence Bonjour. “Externalist Theories of Knowledge.”

entirely outside his ken?”⁷⁷ Thus, he concludes, “Norman’s acceptance of the belief about the President’s whereabouts is epistemically irrational and irresponsible and thereby unjustified.”

There are certainly differences between Monique’s situation and Norman’s. For one, as Brandom presents his example, Monique at least has *some* reasons to think that she might be a reliable reporter of hornbeams. She has, after all, been training to identify them, and she might simply be unreasonably unconfident in her abilities. It is relatively common for individuals to underestimate their epistemic abilities like this. In such cases, when one underestimates a reliable epistemic capacity of theirs even though they have good reasons to be confident in this capacity and is aware of these reasons, an argument might be made that they really are justified in their beliefs formed by this capacity. However, the account Brandom gives puts no weight on whether or not Monique has any reason to be confident in her abilities. The only thing that Brandom cites as a reason to attribute knowledge to Monique is the fact that *someone else* who knows that she is a reliable reporter could cite her report as a reason to believe a hornbeam is present. This line of reasoning opens Brandom up to Bonjour’s counterexample. If someone *knew about* Norman’s cognitive power of clairvoyance and its reliability, it would be appropriate for that person to both believe that the President is in New York and to cite Norman’s report as what warrants that belief. But this does nothing to warrant the claim that *Norman* knew the President was in New York, at least in the standard sense of know. Though we might want to say that Norman is *entitled* to the belief in the sense that his having that belief entitles others, if they are aware of his cognitive capacity, to that belief as well, this sort of positive epistemic status that the belief

⁷⁷ Lawrence Bonjour. “Externalist Theories of Knowledge.”

has is not equivalent to nor sufficient for the status that we might attribute to Norman of *being justified*.

2. Being Entitled and Being Justified: A Conceptual Story

Even if this notion of justification for belief is explanatorily *derivative* on the notion of entitlement for a claim, and entitlement is conceptually prior to this notion, it does not mean that justification can be *reduced* to entitlement. Justification is neither necessary nor sufficient for entitlement. As Brandom notes, there might be good reason to treat someone as entitled on the basis of some external relation even if they aren't aware of that relation such that they have no reason to treat themselves as entitled. So being justified (in the sense that Bonjour cares about) is not a necessary condition for being entitled. In addition, being justified in believing a particular claim is not sufficient for being entitled in asserting that claim.

Suppose I perform a set of experiments which show a particularly surprising result. I check my calculations several times, and come to believe with a high degree of confidence that the result is correct. I then burn all the results and evidence of the experiment having ever taken place. I will be justified in believing the result, but if I go about claiming the results of the experiment I will not be entitled to that claim (assuming I don't have a strong enough reputation of reliability to still secure entitlement for the claim about the results). Being entitled to the claim requires the extra step of actually discursively justifying it, but this extra step is not something that I need to do to be justified in believing it.

At this point, it will be productive to make some sense of the distinction by telling a bit of a genealogical story. Miranda Fricker gives one such story, synthesizing Brandomian epistemic pragmatics (as articulated by Michael Williams) and Edward Craig's genealogical conception of knowledge. According to her story, in the epistemic "state of nature," the primordial state of our epistemic enterprises, there was a "fundamental practical pressure to

stand in co-operative epistemic relations with fellow inquirers.”⁷⁸ Many sets of eyes and ears collecting and sharing information are much better than just one set. Along with this practical pressure came a need to recognize good informants from lousy ones. A good informant is someone who both is able to communicate accurate and helpful information and can be *identified* as someone able to communicate accurate and helpful information. On Fricker’s genealogical story, in the first instance, good informants were recognized primarily on externalist grounds (Are informants of this type usually reliable in these sorts of circumstances? What is the informant’s track record? Was the informant looking in the right direction?). Some of this *might* be available to the informant but it not *essential* that it is.

Eventually, the ability for informants to *give reasons* for the information they are providing, to discursively justify their claims to information-seekers, emerged as a particularly good way of distinguishing good informants from bad ones. I now want to add to this story a bit. With the advent of discursive justification as a particularly good way of identifying oneself as a good informant a new intellectualized status emerged: the status of being justified. This status is importantly different than being recognized as a good informant and thus as entitled to a claim where the ability to discursively justify that claim is a good indicator but not a necessary condition of that status. With this new intellectualized status, the deontological status of epistemic *responsibility* explicitly emerged as directly linked to the ability to *account* for ones claims. An epistemically responsible agent will only claim something if they hold themselves *accountable* for claiming that thing. While there might be good indicators to an *external observer* that an informant is reliable on the grounds of some recognizable external factors, if these factors are not cognitively accessible to the informant, then there is no way of the informant knowing whether or not they are reliable. The rise of

⁷⁸ Fricker 38

internal justification and epistemic responsibility put it on the informant to do the work of identifying itself as a good informant, saving inquirers the effort. It also allowed informants to build up records of being epistemically responsible, giving justifications when called for them which ensure that they are good informants, and gaining more epistemic authority.

Fricker, in her story, does not stress the reciprocity of this whole process. It is not that there are some informants and some inquirers, but rather that all informants are *also* inquirers and vice versa. This is centrally important since, on my view, an inquirer can only make sense of an informer's status of being justified if the inquirer is also versed into the practice of justification. On my story here, the property of being justified that a belief or a believer might have is conceptually derivative on discursive justification. Before we can make sense of recognizing another believer as justified, we must first have mastered the practice of justifying *our own* beliefs. This is not to say that discursive justifications are conceptually prior attribution of *entitlements*, and that is the key genealogical difference between the two statuses. The story to tell here of the conceptual relation between entitlement and justification is not completely straightforward. On the one hand, we might think of entitlement as conceptually prior to discursive justification since, as Brandom says, "One must start with a notion of taking or treating inferences as correct in practice. Without such a practice there is no game of giving and asking for reasons to bring inferences into in the form of explicit assertions." On the other hand we might think of *full-fledged* entitlement as (to use Heidegger's term) equiprimordial with discursive justification, requiring a sort of responsibility for one's claims and thus the ability to discursively justify a claim if challenged. In truth, the move from primordial state-of-nature entitlement, to the fully intellectualized property attributed explicitly to beliefs is a gray and shady one with the advent of discursive justification as a practice playing the central connecting role.

I do not have the space here to give any sort of detailed account of how this move from entitlement to justification might have occurred, nor do I really know what the details such an account would look like, but the progression in broad outline is something like this: (1) Implicit “communicators” who are implicitly recognized within a communicatory system as good or bad informants with entitlement interpretable from fully discursive agents outside of the practice. We see pictures of this in biological accounts such as Dennett’s and Millikan’s. (2) Quasi-explicit entitlement attributions by discursive practitioners, treating as entitled and challenging when appropriate, giving discursive justifications, without actually articulating that structure. This seems to be the minimum conceptual requirement for the sort of normative practice Brandom describes in MIE. (3) Fully articulated and intellectualized justification *criteria*, explicit standards for when a claim is justified. This is the sort of project that we get in traditional epistemological accounts of justification. Seeing the whole genealogical picture of what that particular project is a temporal slice, however, can help articulating a successful account of justification.

3. A Principle of Justification

We are now in a position to formulate a principle for justification, as epistemologists have a tradition of doing. For this role, I propose principle AJ for agent justification:

AJ: An agent S’s belief that P is justified iff, for any *appropriate* challenge to P that S *ought to be aware of as such*, S can, given what she knows, in principle discursively justify P against this challenge.

This principle for justification might seem trivial, but its intuitiveness and simplicity ought to be seen as a virtue, and the reason that it is able to function as a genuine account of justification is the work that was done throughout this thesis which culminated in the account of discursive justification offered in the previous chapter. If it were not for the arguments against skepticism presented in chapter three and the continuation of those

arguments in the previous chapter which made the case that our justifications are rooted in the way things actually are, this principle of justification would be too weak, for there could be things that could be discursively justified in an existing social practice and yet *not actually* be justified. That is, only if we could be sure that our existing justificatory practices are good ones could we get a normative claim that something is actually justified out of the descriptive claim that it could be discursively justified in practice. The arguments of the previous two chapters, if successful, give us the ability to collapse this distinction between the normative claim that a belief is justified and the descriptive claim that it *can be* discursively justified since they attempt to show that we have good reason to endorse our actual epistemic norms (and indeed must *endorse* them), making principle AJ available to us. On the other hand, if it were not for the pragmatic considerations for strong commensurability in the second chapters and the account of justification in the previous chapter, this principle of justification would run the risk of being too strong, excluding things that may be justified, but can't be justified against some challenges from some members of the discursive community because of issues of incommensurability.

I want to now look more closely at this principle, at some of its intuitive appeal and some of the issues that it might face. To answer the question "am I justified in my belief that P?" one can ask herself "Could I justify my belief that P if one were to challenge it?" and subsequently, "what possible appropriate challenges to this belief are there? Can I respond to these challenges in a sufficient manner?" But what does it mean to respond to a challenge in a sufficient manner? Sufficient for *whom*? The first thought is that we ought to be able to provide a justification that would be sufficient *to the challenger*. That is, we answer affirmatively if and only if we think that, if we were given ample time and patience and our challenger was intellectual honest, *they would end up agreeing with us*. Of course, however, it is

not this simple. There are going to be some limitations to this actually happening in practice, and that is why the “in principle” clause is important for AJ. Some people will be consistently unreasonable, some people, like the persistent Agrippan skeptic of chapter three, will continually take themselves to be offering legitimate challenges when they are not and even when it has been explained to them why this is the case. As I said in the opening chapter, I’m not likely to be able to go to the local fortune teller and convince her that Tarot reading isn’t a truth-conductive way of predicting the future. In actual practice, some people will surely refuse to come to agreement, but it is necessary here that it can be made completely clear that the challenger *ought* to come to agreement. It is, of course, not just *actual* challenges raised by other members of discursive practices that we ought to be aware of but also challenges that would be appropriate if they were made. Suppose I am the only one who knows that I don’t own a pair of blue socks. If I form the belief that I am wearing a pair of blue socks right now, there are no actual members of the discursive community that can raise an appropriate challenge against this belief, but I am certainly not justified in believing it. There is an appropriate challenge that can be raised, given what *I* know.

One of the classic issues for accounts of justification has been whether the account is an *internalist* or *externalist* one. The principle I offer consists of a hybrid of both elements. Since the sort of justification that drives this principle of justification is *discursive* justification, a justification that I can *offer* for a particular claim in discursive practice, it is obvious that any pure justificatory externalism must be ruled out. I could not offer in my discursive justification, of course, any fact about my belief-forming process which I am unaware of. If I *am* aware of relevant features of my belief-forming process such as the reliability of a given perceptual capacity I have, then I *can* appeal to this reliability in giving a discursive justification (an S-justification as I defined in the last chapter). For this reason a belief like

Norman's being justified on purely externalist grounds is ruled out. Since clairvoyance is not a default-entitled epistemic practice, there are appropriate challenges to the beliefs Norman forms which have the form "On what grounds did you come to assert that?" and, in absence of a good S-justification in response to this challenge, Norman cannot be justified by AJ.

On the other hand, discursive justification cannot be thought of as wholly internal. The justification that I offer will rest on it being located in a particular discursive context where particular commitments are shared, and only within this external context can my discursive justification have any force. Whether or not I have discursively justified myself depends on whether or not other people do and ought to come to agreement with me. The main externalist element of this account is the use of the notion of *appropriate challenges*. It is this aspect that allows us to avoid the "over-intellectualization" of justified belief that Michael Williams worries about for purely deontological accounts of justification.⁷⁹ For the majority of one's beliefs, asking the question of justification is not usually something that needs to be done. We don't see a book on a table, in normal circumstances, and ask ourselves "am I justified in my belief that there's a book there?"⁸⁰ The fact that this question need never arise reflects the fact that, for most of our everyday beliefs, there isn't an appropriate challenge which compels us to ask this question. In these cases, the criterion of justification AJ still applies but somewhat vacuously so. These beliefs are justified, not because we all have answers to the appropriate challenges, but rather just because there aren't any appropriate challenges.

⁷⁹ Michael Williams, "Responsibility and Reliability"

⁸⁰ Even epistemologists concerned with the question of overarching justification for everyday beliefs don't question particular everyday beliefs in their actual lives. When this does occur in the context of an exchange with a skeptic, and an actual everyday belief is called into question, it is meant simply as a single instance of beliefs being generally unjustified.

Principle AJ is at least partly in line with the deontological conception of justification, since a central feature of it is fulfilling one's epistemic duty by being prepared to give a justification in responses to challenges of which one *ought to be aware*. But it is not *simply* attempting to fulfill one's epistemic duty that will qualify one as justified; one must also be epistemically *competent* in offering justification. Alvin Plantinga, in proposing a difficulty for the deontological conception of justification gives an example of a psychiatric patient who strongly believes some absurd truth and attempts to justify it to everyone he meets. Still, according to Plantinga, though he's confused and his faculties aren't properly functioning, we can't *fault him* for abandoning any epistemic *duty*.⁸¹ Plantinga is right here, but, though justification is a partly deontological concept, he won't in fact be able to discursively justify his claims. He may *think* he has discursively justified them and that his justifications have been misunderstood or ignored, but he hasn't *actually* justified them. Of course, it is not *simply* the generation of agreement in practice that qualifies one as justified, though more often than not this is a good indicator. It's always *possible*, however, that the psychiatric patient is correct and he *is* being misunderstood or ignored due to the epistemic sanctioning that occurred as a result of his situation. The "in principle" clause covers deviant cases like this, and, since we can't possibly achieve epistemic perfection in practice (though we ought to try our best), cases like this are somewhat inevitable.

There is one more conceptual issue here that needs to be addressed. Suppose you put me in the Amazonian rainforest and I have to discursively justify my belief that the theory of evolution is true to a tribe that has no understanding of modern biology (let's assume we speak the same language). A justification offered here would not provide me the

⁸¹ Plantinga, *Warranted Christian Belief*. 101

same sort of entitlement status as a justification offered in the context of, say, an academic debate. So, even if I offer the same justification content-wise in both settings, it may be the case that I have discursively justified myself in one setting but not the other. Still, we say that, even when I am thrown in the Amazonian rainforest and unable to discursively justify my belief to the natives, I am still *justified* in my belief that evolution is true. This seems to pose a challenge to the present account in favor of a purely internalist account, at least with respect to intellectually high-level beliefs like the belief in evolution. I want to say that this challenge does not imply that it's the internal understanding that matters and not the context in which a discursive justification can be produced, but rather it means that we are members of a community (a sub-community within the broad discursive community) in which I *can* discursively justify my belief and it is this community which we are judging my belief with respect to. The treating of my belief as justified, then, still rests on the external context and the attitudes of the surrounding community members.

Given this explanation, however, to avoid the worry of incommensurability at the level of communities, there must be bridges of possible lines of justification connecting the communities. Even if there are many individuals in one community justified in a belief but unable to justify it across radically different community lines (like myself with respect to evolution) there must be *some* means by *some* set of individuals by which a belief can be justified across any community lines. This path will likely include several intermediary steps of communicating, translating, and sharing information, and once again, with many particular beliefs, may not be possible in actual practice given various contingent constraints of time and resources.

One of the most straightforward upsides of this account of justification that should be clear by now is that it fits nicely with the pragmatic response to the epistemic regress

problem was offered in chapter three. The regress problem has long been thought of the central motivation for holding a foundationalist account of justification, but on the default and challenge model of entitlement this problem does not get off the ground, since reasons than can be employed in discursive justification are only needed when there is an appropriate challenge. Combined with the argument against skepticism in chapter three, the idea is that, by the very pragmatic structure of how our discursive practices work, our inferential practices and accordingly our beliefs must mostly be on track, and to thus there is no reason to think that every belief needs to have an explicit justification in the bag. Now, as I argued in the last chapter, we can still allow any claim or inferential process to be challenged, grabbing another foothold for the time being, but once again, we do this all against the backdrop of treating inferences as largely correct and beliefs as largely true.

4. Undefeated Agent Justification, Gettier Problems, and Knowledge

Since Edmund Gettier's famous paper, "Is Justified True Belief Knowledge," the traditional analysis of knowledge as justified, true belief has been called into question.⁸² Though there is no clearly agreed upon response to the problem raised in Gettier's problem, one of the common responses is to propose a fourth condition and say that knowledge is *undefeated* justified true belief.⁸³ A defeater is something that, as the name suggests, defeats the justification or positive epistemic status that a belief would have in absence of such a defeater. As a way of extending AJ this way of thinking about knowledge, I propose principle UAJ, for undefeated agent justification:

UAJ: An agent S's belief that P is undefeatedly justified iff, for any appropriate challenge to P by any real or possible challenger, S can, given what she knows, in principle discursively justify P against this challenge.

⁸² Gettier, "Is Justified True Belief Knowledge"

⁸³ One of the early and influential accounts of this sort was Lehrer and Paxon's "Knowledge: Undefeated, Justified, True Belief."

Thus to say explicitly “I know P” is to explicitly commit oneself to the idea that if someone were to present an appropriate challenge to P, P could be justified against this challenge.

The important distinction between UAJ and AJ is that the challenges being considered are not only the challenges that the agent *ought to be aware of*, but *all* appropriate challenges. This accounts for why justified false belief is not knowledge. An agent with a false belief could satisfy AJ, but not UAJ, since, for any false belief, there must be some way, in principle, to *point out* that it is false and this would be an appropriate challenge to that belief, even if there is no actual agent *aware* of that fact.

This principle also accounts for Gettier cases in which there are justified true beliefs that are not knowledge. In Gettier problems, the agent isn’t epistemically responsible for being aware of the challenge that would require them to revoke their belief (at least given what they know at the moment). These cases usually involve strange circumstances on two ends: (1) a strange circumstance that defeats a justification that would, in the standard case, be successful, and (2) an unlikely coincidence that makes the belief true for reasons other than those cited in the justification. Let’s look at one of Gettier’s two cases with this principle of undefeated justification in mind. Gettier’s second case features Smith, who forms the belief that Jones owns a Ford on very good evidence. Jones, a usually reliable person, has told him that he owns a Ford, and Smith has seen Jones driving around in a Ford. Since Smith is justified in believing Jones owns a Ford, he is also justified in believing an inclusive disjunctive belief with “Jones owns a Ford” as one of the disjuncts, so he forms the belief that Jones owns a Ford or Brown (another friend of his whose whereabouts Smith isn’t sure of) is in Barcelona. Turns out that Jones has lied about owning a Ford, and he in fact rented the car he is currently driving, but, against all odds, Brown is in fact in Barcelona.

So Smith's belief comes out to be true, it is justified, but we don't want to say that it is knowledge.

In this case, there is an appropriate challenge to be made to Smith's belief—the challenge that Jones is, in fact, lying. Jones could make this challenge, or could any hypothetical person who is aware of the fact that Jones is lying. This challenge is certainly appropriate, and, in response to this challenge, Smith would not be able to discursively justify his belief (since he has no grounds on which he could justify the belief that Brown is in Barcelona). If Smith was presented with this challenge, he would be epistemically obligated to give up belief. Still, Smith has no obligation to be aware of the appropriateness of this challenge since he has no evidence that Jones is dishonest, and so he only fails to satisfy UAJ and not the weaker principle AJ.

Of course, examples can be concocted which complicate the picture, but looking at the conditions for knowledge roughly in this way gives us an interesting solution which allows us to flexibly deal with them. First, there can be no coherent attempt to provide a counterexample which aims to demonstrate that UAJ plus true belief is too weak. This is because, for any such proposal, we can simply ask why we aren't willing to attribute knowledge in this case, and any reasons which one can offer for this judgment serve as a challenge to the belief in question that would result in UAJ not being satisfied. The reasons one cites might be any of the traditional accounts of justification or knowledge that epistemologists have proposed:

- Does it contradict immediate experience?
- Does it make an otherwise consistent system of beliefs inconsistent?
- Was it formed by an unreliable belief-forming process?
- Was it formed by cognitive faculties that are not functioning properly or in a different environment than the one for which they were designed?
- Is there a true proposition such that, if the believer knew it, they would revoke their belief?

Depending on the circumstances, any one of these reasons may serve as a legitimate challenge. On the account I've offered here, there's no need to have to pick one and only one. Perhaps some ways of thinking about it systematically provides more compelling and comprehensive reasons than others, and attempting to articulate these ways is fruitful epistemological work with great practical upside. But we shouldn't think that we don't know what knowledge is until we find the best and most comprehensive way of presenting challenges.

There is also the worry that UAJ, like many indefeasibility theories, is too *strong*, excluding things from being knowledge that we should really count as knowledge. Usually the sorts of examples that raise this worry involve misleading defeaters. To give one example that belongs to Marshal Swain, suppose that I attend the wedding of two friends of mine. The wedding ceremony is performed by the priest without any errors, and I am certainly justified in believing that they are in fact married. Entirely unbeknownst to me, however, that the Archbishop has gone insane at the time the ceremony is performed and, in his insanity, denounced the priest as a fraud who is not authorized to marry anyone. In this case, there seems to be a true proposition that would defeat my belief that my friends are married, the proposition that the Archbishop denounced the priest who performed the ceremony as a fraud. If I was informed of this proposition, the justification I had, that I've just seen the wedding, would be no good. Still even though this is case, it still seems clear that I know my friends are married.

When the principle of undefeated justification that we are looking at is contextualized in discursive terms as I have done here, however, we can make sense of these examples. We can entertain UAJ by imagining hypothetical appropriate challengers and

seeing if the subject could justify the belief in question. Of course, we shouldn't allow into this consideration challengers that are intellectually dishonest or presenting false facts. Likewise, we shouldn't allow into this consideration challengers that are *misled*. For *any* of our beliefs we can imagine a hypothetical challenger who is a master of rhetoric and twister of facts who would force us to rescind that belief if we actually encountered them, so none of our beliefs could satisfy the principle. So, to take Swain's example, any hypothetical challenger who offered the Archbishop's denouncement as reason to give up my belief is either intentionally being misleading or they themselves are misled themselves, so we do not take them into consideration when thinking about UAJ. Of course, if I *actually spoke* to someone who was aware of the Archbishop's denouncement (but not of the fact that he had just gone insane), then I would not satisfy UAJ, since I would not satisfy the weaker AJ. I would have an epistemic obligation to take his challenge into account and, unable to respond to it (presuming I don't know he's insane), I would no longer be justified in my belief.

5. Tying Loose Ends

Now I certainly don't take what I've said here in these last sections to be the final word on the analysis of knowledge. All I'm proposing is that if we think of justification centrally in discursive terms as I have done throughout this thesis, we have new and effective tools with which we can analyze knowledge, drawing constructively from the way we actually interact. If what I've argued throughout in this thesis is roughly correct then most of our beliefs that satisfy AJ also satisfy UAJ. If our epistemic practices are by and large getting it right, then when we can respond to all the challenges that we ought to be aware of, there's usually no challenge that would defeat our commitment that we're not aware of. We are justified in holding the vast majority of our beliefs, and most of our beliefs do in fact constitute knowledge. Though of course all of our sets of beliefs are different due to our

varying standpoints in the world, most of the difference in belief that exists is complementary and completely unobjectionable. Where there *is* conflicting difference in belief, I've given some arguments why we should be optimistic that there are no *epistemological* reasons that these differences cannot be resolved (though of course there may be *practical* reasons).

My primary concern here has been *empirical* knowledge, since I think that this comprises most of our knowledge. There are many sorts of truth claims that I did not talk about in this thesis. The scope of what I've talked about was already wide enough as it stands. I did not talk about a priori mathematical or logical claims, nor did I talk about moral claims which, depending on one's view, may or may not end up being similarly a priori. However, given the Brandomian picture that I'm working with here it should be relatively straightforward to see how a priori mathematical and logical claims fit in as explicit expressions of universal rules of inference. I have assumed a naturalistic picture here, but I did not dive deeply into the nature of religious claims. If these claims end up functioning in the same way as empirical claims as some philosophers of religion think they ought to, then what I say here about empirical knowledge should carry over. But perhaps religious claims should be regarded as an aspect of quite different sort of discourse, and if this case, the questions of how common ground ought to be found with respect to these claims is not one I have addressed here.

This thesis has been an attempt to expand Brandom's incredibly large-scope project to explain our normative practices generally to the field of epistemology primarily by connecting it to Davidson's response to skepticism. I've attempted to explain justification in such a way that is a significant elucidation of our actual epistemic practices while also trying to answering some of the traditional epistemological problems. As answers to these

problems, the conclusions here are, of course, epistemological ones, but, also serving as an example of a particular *way* in which epistemological problems can be addressed, the conclusions here are as much *meta-epistemological* as they are epistemological. I suspect that the same strategy here of extending Brandom's broad picture of normative practices will be fruitful not just in epistemology but in metaethics as well, offering a sturdy conceptual ground for naturalized ethics in the same way the conclusions here ought to provide some grounds for naturalized epistemology. Of course, much of explanatory potential here rests on a commitment to the fruitfulness of Brandom's project. The final word about how fruitful this project really is will rest on how the empirical and philosophical landscape continues to develop.

Works Cited

- Bonjour, Laurence. 1980. "Externalist Theories of Empirical Knowledge." *Midwest Studies in Philosophy* 5:1, pp 53-74.
- Brandom, Robert. 1994. *Making it Explicit*. Harvard University Press.
- Brandom, Robert. 2001. *Articulating Reasons*: Harvard University Press.
- Bostrom, Nick. 2003. "Are You Living in a Computer Simulation?" *Philosophical Quarterly*. 53:211 pp. 243-255.
- Carroll, Lewis. 1895. "What Achilles Said to the Tortoise," *Mind* 4:14 pp. 278-280.
- Chalmers, David. 2005. "The Matrix as Metaphysics." in *Philosophers Explore the Matrix*. Oxford University Press.
- Davidson, Donald 2001. "A Coherence Theory of Truth and Knowledge" in *Subjective, Intersubjective Objective*. Oxford University Press. (Originally read at a colloquium in 1981, with the afterthoughts appearing in 1987.)
- Davidson, Donald. 2001. "The Myth of the Subjective," in *Subjective Intersubjective, Objective*. Oxford University Press. (Originally read at a conference in 1986.)
- Davidson, Donald. 2001. "On the Very Idea of a Conceptual Scheme" in *Inquires into Truth and Interpretation*. (Originally delivered in 1973 at the Eastern Meeting of the American Philosophical Association.)
- Davidson, Donald. 2004. "The Problem of Objectivity," in *Problems of Rationality*. Oxford University Press.
- Dennett, Daniel. 1981. "Intentional Systems," in *Brainstorms: Philosophical Essays on Mind and Psychology*. MIT Press.

- DeRose, Keith. 1999. "Contextualism: An Explanation and Defense." in *The Blackwell Guide to Epistemology* ed. J. Greco and E. Sosa. Blackwell.
- Fricker, Marilyn. 2008. "Skepticism and the Genealogy of Knowledge." in *Philosophical Papers* 37:1 pp. 27-50
- Gettier, Edmund. 1963. "Is Justified True Belief Knowledge?" in *Analysis* 23: 6 pp. 121-123.
- Haack, Susan. 2009. *Evidence and Inquiry: A Pragmatist Reconstruction of Epistemology*. Prometheus Books. Originally Published in 1993.
- Hacker, P.M.S. and M.R. Bennett. 2003. *Philosophical Foundations of Neuroscience*. Wiley-Blackwell.
- Haidt, Jonathan. 2012. *The Righteous Mind*. Pantheon Books.
- Haselton, Martie, Daniel Nettle, and Paul Andrews. 2005. "The Evolution of Cognitive Bias." In *The Handbook of Evolutionary Psychology*, pp. 724-746. Ed. David M. Buss. Hoboken, NJ, US: John Wiley & Sons Inc.
- Kim, Jaegwon. 1988. "What is Naturalized Epistemology?" in *Philosophical Perspectives* 2 pp. 381-405
- Klein, Peter. 1999. "Human Knowledge and the Infinite Regress of Reasons." In *Philosophical Perspectives*, 13 pp. 297-325
- Kornblith, Hilary. 2002. *Knowledge and Its Place in Nature*. Oxford University Press.
- Kuhn, Thomas. 1962. *The Structure of Scientific Revolutions*. University of Chicago Press.
- Kukla, Rebecca and Mark Lance. 2009. *"Yo!" and "Lo!"* Harvard University Press.
- Kukla, Rebecca and Mark Lance. (Forthcoming) "Intersubjectivity and Receptive Experience." *The Southern Journal of Philosophy*.

- Lehrer, Keith and Thomas D. Paxson Jr. 1969. "Knowledge: Undefeated Justified True Belief." *The Journal of Philosophy* 66:8 pp. 225-237.
- McDowell, John. 1994. *Mind and World*. Harvard University Press.
- Plantinga, Alvin. 1993. *Warrant and Proper Function*. Oxford University Press.
- Plantinga Alvin. 2000. *Warranted Christian Belief*. Oxford University Press.
- Plantinga, Alvin. 2010. "The Evolutionary Argument Against Naturalism" in *Science and Religion in Dialogue*. ed. Melville Y. Stewart. Wiley Blackwell. pp. 324-333
- Putnam, Hilary. 1981. "Brains in a Vat" in *Reason, Truth and History*. Cambridge University Press.
pp. 1-21.
- Quine, W. V. 1993. "In Praise of Observations Sentences." *The Journal of Philosophy*: (1993).
- Quine, W.V. 1969. "Epistemology Naturalized." in *Ontological Relativity and Other Essays*. Columbia University Press.
- Quine, W.V. 1981. "On the Very Idea of a Third Dogma." in *Theories and Things*. Harvard University Press.
- Quinton, Anthony. 1973. *The Nature of Things*. Routledge
- Rorty, Richard. 1979. *Philosophy and the Mirror of Nature*. Thirtieth Anniversary Edition. Princeton University Press.
- Sellars, Wilfrid. 1997. *Empiricism and the Philosophy of Mind*. Harvard University Press (Originally published in 1956).
- Sosa, Ernest . 2003. "Davidson's Epistemology." In *Contemporary Philosophy in Focus: Donald Davidson*, ed. Kirk Ludwig. Cambridge University Press.
- Swain, Marshall. 1974. "Epistemic Defeasibility" in *The American Philosophical Quarterly* 11:1 pp. 15-25.

- Tomasello, Michael. 2008. *Origins of Human Communication*. MIT Press.
- Wanderer, Jeremy. 2010. "Brandom's Challenges" in *Reading Brandom: Essays on Making it Explicit*. Ed Jeremy Wanderer and Bernard Weis. Routledge.
- Williams, Michael. 2008. "Responsibility and Reliability." in *Philosophical Papers* 31:1 pp. 1-26
- Williams, Michael. 2001. *Problems of Knowledge*. Oxford University Press.
- Wittgenstein, Ludwig. 2009. *Philosophical Investigations*. Fourth Edition trans. G.E.M. Anscombe, P.M.S. Hacker, and Joachim Shulte. Blackwell Publishing. (Originally published in 1953)