

9-2015

Social Epistemology and the Project of Mapping Science

Kamili Posey

Graduate Center, City University of New York

[How does access to this work benefit you? Let us know!](#)

Follow this and additional works at: https://academicworks.cuny.edu/gc_etds

 Part of the [Philosophy Commons](#)

Recommended Citation

Posey, Kamili, "Social Epistemology and the Project of Mapping Science" (2015). *CUNY Academic Works*.
https://academicworks.cuny.edu/gc_etds/1097

This Dissertation is brought to you by CUNY Academic Works. It has been accepted for inclusion in All Dissertations, Theses, and Capstone Projects by an authorized administrator of CUNY Academic Works. For more information, please contact deposit@gc.cuny.edu.

SOCIAL EPISTEMOLOGY AND THE PROJECT OF MAPPING SCIENCE

BY

KAMILI POSEY

A dissertation submitted to the Graduate Faculty in Philosophy in partial fulfillment of the requirements for the degree of Doctor of Philosophy, The City University of New York.

2015

© 2015
KAMILI POSEY
All Rights Reserved

This manuscript has been read and accepted for the
Graduate Faculty in Philosophy in satisfaction of the
dissertation requirement for the degree of Doctor of Philosophy.

Nikolas Pappas

Date

Chair of Examining Committee

Iakovos Vasiliou

Date

Executive Officer

Samir Chopra

Linda Alcoff

Robert Sinclair

Supervisory Committee

THE CITY UNIVERSITY OF NEW YORK

ABSTRACT

SOCIAL EPISTEMOLOGY AND THE PROJECT OF MAPPING SCIENCE

BY KAMILI POSEY

ADVISOR: SAMIR CHOPRA

One area of debate in naturalized epistemology is how to best interpret the relationship between naturalism and traditional analytic epistemology. This is particularly the case for epistemologists who commit to methodological forms of naturalism:

First, I contend that methodological naturalists must commit, at minimum, to the idea that the best methods of knowledge-production in science ought to determine the best methods of knowledge-production full stop. I also contend that the best methods of knowledge-production in science are determined by and performed within scientific communities. Thus a methodological naturalist, to minimally call herself such, ought to consider knowledge-production within “communities of inquirers” utilizing community-determined methods (*truth-conducive or not*) as the fundamental framework of the method of science.

Second, I contend that “communities of inquirers,” more broadly, including the paradigmatic “scientific community,” ought to be the primary terrain of naturalized epistemologists committed to methodological naturalism. And, because of this, I claim that all naturalized epistemologists ought to be *social* epistemologists. I agree that this may seem undesirable: social epistemology, following Rorty (1979), has a history of focusing on groups, institutions, and collective doxastic agents in consensus-based models that are either non-*veritistic* or radically deflationist about truth and objectivity. This has led to the exclusion of social epistemology from the traditional canon of epistemology, and the rejection of Peirce’s (1877) “communities of inquirers” view as too entangled with social and political “noise” to be conducive to finding truth. But I think this is an unnecessary interpretation. I claim, following Goldman (1999), that social epistemology *can* be truth-directed or “*veritistic*” in practice – even if not in the way that Goldman

(1999) imagines. Ultimately, I argue, inspired by the American pragmatists, and particularly Dewey and Peirce, we do not have to reject truth and objectivity in order to maintain a social account of knowledge-production.

Lastly, I suggest that one benefit of this view of epistemology is that it relocates a crucial area of philosophical inquiry from within the walls of “academic philosophy” and situates it where it actually belongs – *in the world*. One benefit of doing this is that epistemologists can make meaningful use of emerging tools in information science to visualize and theorize the activities within various “communities of inquirers:” this includes such tools as large-scale data mining, scientific citation mapping, topic and/or cluster mapping, and data visualizations.

PREFACE

“A philosopher is like a fly buzzing around in a fly bottle, according to Wittgenstein. A theory of the bottle, it doesn’t need. What it needs is to be shown the way out. Kant has an almost complete opposite metaphor. He imagines a dove resentful about air resistance; it could fly better, surely, if the air would just get out of the way. A lot of philosophy is these two metaphors battling it out. You are always trying to break free of something. This makes sense if you’re the bottled-up fly. But not if you’re the ungrateful dove. It can be very hard to tell.”

-Stephen Yablo

TABLE OF CONTENTS

Title Page.....	i
Copyright Page.....	ii
Approval Page.....	iii
Abstract.....	iv-v
Preface.....	vi
Table of Contents.....	vi-ix
List of Illustrations.....	x
1. The Introduction	
1.1. The Proposal.....	1
1.2. The Motivation.....	14
1.3. The Payoff.....	16
1.4. The Methodology.....	17
1.5. Project Outline.....	18
2. On Naturalism	
2.1. Methodological Naturalism and the Scope of “Science”.....	25
2.2. Methodological Naturalism and The Pragmatist Conception of Science.....	35
2.3. Social Epistemology and a More Honest Naturalism.....	38
3. Arguments for Social Epistemology	
3.1 Introduction.....	40
3.2 Radical Social Epistemology.....	45
3.2.1 Against Radical Social Epistemology.....	50
3.2.2 Can RSE Be Saved?.....	56
3.3 Consensus-Based Social Epistemology.....	63
3.3.1 Background.....	63
3.3.2 Rortyan Pragmatism.....	66

3.3.3	Feminist Philosophy of Science and Feminist Epistemologies.....	72
3.3.4	Solomon (2001) and Standpoint Epistemologies of Science.....	73
3.3.5	Against (CSE).....	76
4.	Veritistic Social Epistemology.....	79
4.1	Goldman (1999) and the Social Epistemological Project.....	79
4.2	<i>V-Values</i> , Non-Epistemic Interests, and Questions of Interest.....	84
4.3	Extrapolating from Testimony: Does Goldman (1999) Give Us a <i>Bayesian Social Epistemology</i> ?.....	90
4.3.1	Goldman (1999) and the “Problem of Priors”.....	92
4.3.2	Replies: Kitcher (2002) and Fallis (2002).....	95
4.3.3	Concluding Remarks.....	98
5.	Veritistic Social Epistemology in a Pragmatist Framework.....	99
5.1	Pragmatism, Realism, and Convergence.....	99
5.2	Pragmatism and the Convergence Thesis: A Bayesian Approach.....	101
5.2.1	Scoping Methods: A Subjectivist-Objectivist Approach.....	102
5.2.2	Dutch Books, Diachronic Coherence, and Revising the Reflection Principle.....	104
5.2.3	General Non-Partitioned Reflection.....	110
5.2.4	Scoping (GNR): Defeater Beliefs, Inert Beliefs, and “R-Reflection”	112
5.3	A Note on Bayesian Aggregation.....	113
5.4	Epistemic Disagreement and Epistemic Peers.....	116
5.5	What’s Next?: Mapping The Epistemic Terrain.....	120
6.	A Network-Level Epistemology: Representing Large-Scale Data.....	122
6.1	Systems-Oriented Social Epistemology and the Problem of Scale.....	122
6.2	A Network-Level Epistemology.....	125

6.2.1	Scientific Citation Studies and Scientific Citation Mapping.....	126
6.2.2	“Open” and “Closed” Network-Mapping for Social Epistemology	129
6.2.3	Can “Closed” Systems Contribute to Knowledge?.....	134
6.2.4	Network-Mapping and Emerging and Waning Clusters.....	135
6.3	Representing Large-Scale Data: Concluding Remarks.....	136
6.4	A Final Note on the “Social” Part of Social Epistemology.....	137
	Bibliography.....	139
	Network Analysis Maps – Bibliographic Information.....	148

LIST OF ILLUSTRATIONS

Figure 1, Map of Scientific Paradigms.....	139
Figure 2, The Structure of Science.....	140
Figure 3, Science-Related Wikipedian Activity.....	141

Chapter I

The Introduction

1.1 The Proposal

I assume that a good number of working philosophers subscribe to some variety of naturalism; even though it does come riddled with problems. Although I do not assume that most working philosophers subscribe to ontological naturalism, or that we need be dogmatic about what we mean by the “science” part of naturalism (See: Timothy Williamson’s 2011 NYTimes article, “What Is Naturalism?” for a good discussion of this).¹ It could very well be that philosophy is called upon by evidence to adopt both ontological naturalism and some variety of methodological dogmatism. But I aim to leave both of those possibilities aside for now. Generally speaking, a modest naturalist will be at least a methodological naturalist. The methodological naturalist claims, loosely, that our methods in philosophy (and inquiry in general) ought to be aligned with the best methods of science. Now, I also take it that much of the work of naturalism is involved with understanding how “science,” and more specifically, working scientists, generate and justify knowledge claims. Here the philosopher often works with loaded terms, like, “truth,” “objectivity,” and “veritism.” The aim is to give an account of how truth, etc., are arrived at via scientific inquiry and too often this ignores the following underlying irritation: At the root of scientific practice, scientists are all, of course, horribly fallible. This is compounded by the fact that, in practice, scientists also work as collective (or collaborative) inquirers. It is here that the social epistemologist meets the methodologist naturalist with deep concerns about a theory/practice divide.

The problem is to situate a truth-oriented theoretical understanding of scientific practice with the practice itself—and with all of its very real prejudices and complexities. This leads many sociologists of science and social epistemologists to claim that giving a truth-orientation to scientific theory is an unworkable “philosopher’s problem” (See: Solomon (2001)). Social epistemology, particularly of the sort expressed by Rorty (1979), Stich (1990, 1993), Longino (1990, 1993, 1994), Nelson (1993), and Solomon (1994, 2001), aims to meet the theoretical criteria of “truth” and “objectivity” in scientific practice by linking

¹ Williamson, Timothy. “What Is Naturalism?” *The New York Times*. [New York] September 4, 2011, Online Edition, The Opinionator: <http://opinionator.blogs.nytimes.com/2011/09/04/what-is-naturalism>.

truth criteria with criteria for justification—or, in some cases, with acceptance by a given community of inquirers. Additionally, they try to move away from ideal accounts of truth and individualistic or reductionistic accounts of ideal knowers in favor of approaches that prioritize the communal nature of scientific practice. The social epistemologist insists on something that the scientist naturally assumes—that is, however you conceive of scientific truth, it is most definitely a social product.

Not all social epistemologists agree on this, however. Goldman (2010) argues that “social epistemology” does not count as “real” epistemology in the cases where the epistemological goals are revisionist—or where they aim to completely undermine the goals of traditional epistemology. In his answer to the question: “What is social epistemology?” He offers a way to conceive of social epistemology as being in line with the goals (if not the content) of traditional epistemology. For Goldman (2010), this means that the social epistemologist’s goals still need to be truth-oriented, focused on objective rationality, and strict epistemic criteria for rationality. For this account, the importance of being aligned with traditional epistemology also serves as an important criterion for a viable social epistemology. This is primarily because of critics concern that the projects of social epistemology—particularly the epistemics of social institutions—might rightfully belong to other disciplines (See: Olsson’s (2006) reply to Goldman).

There are three things that we may be concerned with answering here concerning the nature of social epistemology. First, we’ll need to come to some understanding about what the unit of analysis should be, i.e., how do we answer the normative question about the unit of epistemic analysis? Are we dealing with individual epistemic agents or collective agents or both? Second, we’ll need to know how well the projects of social epistemology square with the projects of traditional epistemology, i.e., what is the relationship between social and traditional epistemology? And, third, we have to be able to say something about the truth-directedness, or truth-orientation, of inquiry. Specifically, we need some way of addressing the issue of truth in inquiry, particularly in light of the “social” nature of knowledge-production. How is it that groups, aggregates, and/or institutions can be truth-directed? Although these are interesting questions for analyzing problems internal to social epistemology, they are, for the most part, unnecessary for arguing for social epistemology in general. Instead, if we can sufficiently answer the question about the relationship between traditional and social epistemology, we may have no need to

answer these other questions. Further, and perhaps more importantly, we may have no need to (1) justify social epistemology as a project distinct from epistemology in general, or (2) make social epistemology answerable to mainstream, or traditional epistemology (See: Goldman (1999, 2010)). Thus, the remaining question about the nature of truth in inquiry will just become a problem of epistemology proper.

I prefer to think of “social epistemology” as a cluster of views about how social practices, behaviors, and institutions, etc., inform both knowledge-production and the “truth-making” process. However, I am not committed to this particular interpretation of social epistemology full stop. Because I hope to show that all epistemology is essentially social epistemology, it matters little whether or not we come to a consensus about how to best understand what we mean by the “social” part of social epistemology. If we can show that social context is a necessary part of answering epistemological questions of any variety, then all other discussions become matters of emphasis or degree. I’ll explain: If all epistemological questions must appeal, in some sense, to the social context in which they arise, then all epistemology is social epistemology. The “social” in this case will take on different meanings dependent on the particular context. As I see it, there is no benefit in outlining what will count as “social” in order to provide epistemic—as opposed to practical—constraints on knowledge-production. The social contexts that I have in mind take group inquiry (e.g., juries), aggregated inquiry (e.g., lab science, lab practice), and institutional inquiry (inquiry conducted by social, political and/or cultural institutions) as paradigmatic cases, not limiting cases of social knowledge-production. What we mean by “social” or “social context,” however, should remain flexible. This is particularly the case because of ongoing technological considerations (e.g., digital social spaces) that require that we cast the net of what we mean by “social” pretty wide. So the question is now: what justifies the assumption that all epistemology is inherently social epistemology? And how can social epistemology be veritistic? It is not enough to say that philosophical naturalism lends itself to a socially oriented epistemology (via naturalism’s commitment to the methods of “good” science), nor is it enough to say that the methods of “good” science are fundamentally veritistic. These claims must be justified by some other means. So here I look to the work of the classical pragmatists.

The classical pragmatists insist that there is no separation between minds and nature, or knowledge and the external world. They argue instead that we need to rid ourselves of what Dewey calls

our tendency toward “pernicious dualisms” that separate inquirers from the environment of inquiry, and take the project of philosophy to be fundamentally transcendental—or concerned with the nature of things “behind” what is empirically given.² Philosophy “done right” aligns itself with the methods of scientific inquiry and gives up speculative metaphysics and pointless epistemological questions about how we gain access to the “external world.” (The pragmatists never questioned that the world we experience is the world as it is). Generally speaking, I think that the classical pragmatists argued for a non-transcendental epistemology that used the methods of science as a guide to both formulating epistemological questions and as a method for belief-formation.³ They wanted a way to fully conceptualize the idea of epistemological methods with decidedly “scientific orientations,” to establish a type of scientific epistemology that preserves both truth and objectivity. This does not entail either an instrumental or relativistic account of epistemology—or its abandonment altogether. Our methods may be in line with science without being reduced to current methods in science, and may be value and interest-laden without being a mere matter of instrumental values full stop.

Although some philosophers disagree with the claim that there is a coherent “pragmatist epistemology,” as opposed to philosophical views belonging to Peirce, James, and Dewey, respectively—I would argue that there is a clear, delineated path from Peirce’s theory of truth as convergence, to James’s “pragmatist conception of truth,” to Dewey’s theory of truth as presented in the *Logic*. Additionally, I would argue along with R. A. Putnam (2010) and Pihlström (2004) that those who deny a

² See: Dewey (1929).

³ There is an interesting point to consider here concerning James’s “Will to Believe” and the idea that, concerning some types of beliefs, we are entitled to believe *without* Clifford’s “sufficient evidence” requirement as the cost of foregoing the belief outweighs the requirement that we do our due diligence in justifying it in the first place. As James argues, there are beliefs that we cannot possibly justify in our lifetime—for example, religious beliefs—and the cost of resting faithless, or abstaining from belief, is too high. Additionally, he argues there are types of beliefs where the “faith” in the belief must come before the “fact” of the belief (and, in some cases, faith can actually *create the fact*), as this is how we come to believe in things like currency and political institutions and ideals, like democracy. Although James’ view here appears to delimit how thoroughly he can support the Peircean “method of science” as our primary mode of inquiry, I think it’s important to note – as Turrisi (1997) suggests of Peirce’s “fixation” of belief and Furhmann (2006) reiterates nicely – that the “method of science” exercises more of a psychological or normative command on the inquirer to keep the *process of inquiry* sufficiently alive. That is, there is a normative command on inquiry that requires the inquirer to make certain that one does not rest on the *feeling* that a belief is closed merely because one feels emotionally confident, or satisfied, with the result of their findings. If this is the important methodological difference between the “method of science” and other methods of inquiry, then I see no conflict with James’s analysis of belief in the “Will to Believe,” unless one takes James to be saying that faith as a *wholly unexamined* sentiment or feeling is suitable for a justified belief and I do not take this as a reasonable characterization of James’s argument.

consistent “pragmatist epistemology” (Levi (2002), Margolis (2002)) may be focusing on differing metaphysical emphases amongst the classical pragmatists as opposed to epistemological differences. As differences in metaphysical theses do not (and cannot) influence the induction-based pragmatist conception of truth and objectivity, I would argue that all arguments along this line would ultimately fail. I agree that there are still substantive problems with the pragmatist conception of truth and objectivity (See: Levi (2002) and Hookway (2004)), but I don’t agree that any of these problems are fatal for pragmatist epistemology in general. The criticisms of the pragmatist conception of truth and objectivity (Again, see: Levi (2002) and Hookway (2004)) are useful in rethinking how to tool the approach, but in general they are too superficial to take seriously. Fuhmann (2006) is correct in saying that the convergence thesis—the contentious notion at the foundation of the pragmatist conception of truth—is actually more necessary to pragmatist epistemology than it initially appears.

This is because the convergence thesis is wedded to a formal apparatus for belief-change: a frequentist theory of probability tied to a community of inquirers using the same method. A group of inquirers iterating the same experiments should be able to objectively determine any hypothesis via a frequency probability-function carried out in the “long run” (as Dewey likes to say). However, there are two problems with this view: (1) the “end of inquiry” may in fact be “many headed,” or result in many (many) final “truths,” and (2) the frequency view does not give us a clear way to constrain our present beliefs.⁴ Fuhmann addresses the first of these problems quite nicely. He writes:

There is a simple way of guaranteeing the eventual coincidence of different inquiries...the threat of divergence stems from the possibility that X and Y are differently tempered: X and Y may resolve differently choice situations that arise in episodes of belief change. If X and Y could be brought to exercise the same choices—if, in other words, they would eventually adopt the same selection function...then given sufficient exposition of evidence, their states of belief will indeed

⁴ As Fuhmann (2006) writes: “What seriously interferes with the ideal of convergence is something James called *temperament*. Different inquirers, when confronted with the same evidence, may have different dispositions to react to that evidence. How one might adjust one’s beliefs when realizing a certain tension between them or when faced with new evidence may depend not only on one’s present beliefs, but also on how deeply entrenched they happen to be. It is mainly this second aspect that James has in mind when he writes of temperament.”

converge in the limit. So as to achieve convergence of beliefs we need convergence of dispositions to believe or, as James puts it, of temperaments.⁵

However, as Furhmann notes this is precisely where we find the “pragmatic role of absolute truth in inquiry.”⁶ Absolute truth “exercises a normative command” on inquiry by aligning our individual preferences (as inquirers) with a demand for truth.⁷ This is also where the pragmatist emphasis on a community of inquirers generates its real epistemic weight. If we only answered to ourselves as inquirers “every inquirer would be free to converge on his own personal limit of inquiry.”⁸ It is the community that compels the individual—in the name of truth in inquiry—to examine and eliminate any prejudices or biases that interrupt the pursuit of truth. This process, however, specifically is aimed at the long run. What it does not yet tell us is how to “truth-direct” our epistemic practices in the short run.

The problem that arises with the pragmatist conception of truth as convergence does so because of the pragmatist tendency to establish a theory/practice divide. Pihlström (2004) argues that this has happened, in part, due to James and Dewey’s adoption of a Peircean view of inquiry in conjunction with decidedly more pronounced views about human behavioral psychology. The primary differences here in the views of truth are not in kind but of degree of emphasis. Although the pragmatists focused on the method of science (“pure science” or “frequency-distribution”) to arrive at truth in the abstract, they did not explicitly show how we arrive at truth in the case of the concrete or everyday “particulars.” Thus it does seem likely that the root of this discrepancy actually lies in a metaphysical debate about universals (abstractions or “generals”) and particulars. In particular, on Peirce’s fierce refusal to accept nominalism even as an instrument to make our behavioral practices more conducive to explanation—as James and Dewey used them.

⁵ See: Furhmann (2006).

⁶ See: Furhmann (2006), p. 48.

⁷ See: Price (2010): He argues for truth as a necessary normative requirement of assertoric dialogue. “Truth,” in this sense, is one of three ways (along with “sincerity” and “justification”) that we have to govern and analyze our linguistic practices. Due to his thorough rejection of representationalism, truth is in no way tied to a metaphysical understanding of the external world—and thus, as Price admits, is not philosophically related to the classical pragmatist conception of truth. Contrastingly, what I suggest here, and what I take Furhmann (2006) to be making a case for, is that truth exercises a “normative command” on inquiry *because* of its connection with a formal, and potentially *metaphysical*, truth-making process. And it is this *process* that justifies our “truth-talk.”

⁸ *Ibid.*

What is called for here is better account of how we can arrive at truth in everyday inquiry in a way that (1) does not invoke the unnecessary metaphysical debate about universals and particulars, and (2) puts epistemic constraints on everyday inquiry in a way similar to what is called for in the “method of science.” To this end, the best candidate for epistemic constraint is an “operational,” or “bounded” account of truth (that comes to us via a bounded account of reality, one that, following McDowell’s pragmatism in *Mind and World*, requires the use of concepts) and a more constrained, decision-theoretic method of inquiry.⁹ Here I suggest that an operational account of truth tied to an objective Bayesian decision procedure will likely generate the kind of epistemic constraints that are needed in everyday inquiry. This method is inspired by Goldman’s (1999) use of Bayesian inference in testimony.

Goldman (1992) argues for a view he calls “social epistemics,” where the aim of epistemological questions is not only normative and prescriptive, but also dependent on the entire social, cultural, and political factors that influence the beliefs of individuals and communities. Goldman discusses four general evaluative positions for the social epistemologist: relativism, consensualism, expertism, and veritism. The two most prominent discussions concern relativism, which he attributes in part to Rorty (1979) and Kuhn (1970), and expertism, which he attributes to Stich and Nisbett (1980).¹⁰

⁹ For this account, objective likelihoods are constrained by what we know at *present* about the probability (frequency—or “chanciness”) of any given hypothesis. This account is “operational” in that all likelihoods will inevitably change (even if the probability of a hypothesis—say, the sun will rise tomorrow, is .99, it will still converge toward 1) as inquiry continues.

¹⁰See: Goldman (1992), p. 185. Here, he writes: “Relativism consists of three theses: (A) There are no universal, context-free, super-cultural, or transhistorical standards by which to judge different belief-forming methods. (This thesis might be called *nihilism*, or *anti-universalism*.) (B) Whatever methods a group accepts are *right for them*. (C) An individual’s belief is *socially warranted* (or rational) if and only if this belief is formed (or sustained) by methods that are acceptable by the individual’s group. I am not sure this position is fully endorsed by an identifiable theorist. Some theorists, such as Barry Barnes and David Bloor (1982) and Richard Rorty (1979), clearly endorse thesis (A)...They also appear to endorse (B). But it is debatable whether they endorse (C). Rorty, however, favors an account of rationality and epistemic authority in terms of ‘what society lets us say’; he identifies ‘the community as source of epistemic authority’ (Rorty 1979, pp. 174, 188). So perhaps he accepts the entire version of relativism I am describing. There are also affinities with Thomas Kuhn (1970), who sees the scientific community’s current paradigm as the only available instrument by which to appraise a member scientist’s beliefs” (185). *Expertism*, on the other hand, is the view that a group’s “belief-profile” is rational if the “whole profile [non-coincidentally] reflects the opinions of its constituent experts” (188-189). Additionally, of Stich’s (and Lehrer and Wagner’s) “expertism,” Goldman writes: “The source of rationality is each person’s (repeated) revision of his subjective probabilities in accord with weights he assigns to others and their subjective probabilities. Now the weights a person assigns represent his assessments of the competence, reliability, or expertise, of the members in question...There is no guarantee, then, that the weights (or degrees of ‘respect’) have any correspondence to genuine competence or expertise” (191-192). Note: the original reference for Goldman’s (1992) discussion of Kuhn comes from Kuhn (1962).

In brief, Goldman claims that the problem with Rorty and Kuhn's relativism and Stich's expertism is that both views are non-veridical. He argues instead for a position he calls "veritism," where "the goal of truth is the common denominator of intellectual pursuits, whatever methods or practices are championed as the best means to this end."¹¹ And because "knowers" are best understood in (multiple) relation(s) to other "knowers," one of the goals of veritistic social epistemics is to conceptualize how communities of knowers can arrive at truth. This is not to say that communities of knowers do not have other, also legitimate, ends in mind (e.g., political, cultural, legal, conversational, etc.), but insofar as their ends are epistemological and related to social knowledge, the goal of truth is the right one. Where I deviate from Goldman (1992, and more clearly in 1999) is in his use of subjective Bayesianism as the primary formal apparatus of truth-making.¹² Specifically, I will argue that Goldman (1999) cannot use the Bayesian model in the way that he would like to.¹³ Goldman does not offer (1) any justification for why a "reasoner's" subjective credences ought to match objectively likelihoods (although Lewis' (1980) Principal Principle is available to him) and (2) he does not show how in a subjective model a reasoner's credences will be sufficiently constrained at the outset of inquiry. Goldman makes the following case for subjective Bayesianism:

In particular, what I shall show (roughly) is that when a reasoner starts with accurate likelihoods (analogous to true premises), it is objectively probable that Bayesian inference will increase his degree of knowledge (truth possession) of the target proposition. More precisely, after using Bayes' Theorem, the objectively expected degree of truth possession (or V-value) associated with the reasoner's posterior DB [degree of belief] will be greater than the objectively expected degree of truth possession (or V-value) associated with the reasoner's prior DB...If a probabilistic reasoner begins with inaccurate likelihoods, she cannot expect the Bayesian method to improve her V-value...So let us ask what the Bayesian method will do if the reasoner has accurate likelihoods, that is, if her subjective likelihoods match the objective likelihoods. Here is where we

¹¹See: Goldman (1992), p. 192.

¹² See: Goldman (1999): He does address the obvious concern that Bayesianism of both subjective and objective stripes are individualistic models of belief-updating, but extrapolates the general model from a Bayesian analysis of testimony, which requires the inputs of multiple "knowers."

¹³ See: Kitcher (2002), Fallis (2002) and Talbott (2002) for additional critiques of Goldman's approach here.

locate our mathematical results. If subjective likelihoods match objective likelihoods, use of Bayes' Theorem leads to an objectively expected increase in degree of truth possession (V-value). In other words, under these conditions a Bayesian practice exhibits positive V-value.¹⁴ Goldman's view is in fact an unjustified account of how a Bayesian model could work for "social epistemics." Instead, he needs a suitably objective model that is also revisable, like the one that I suggested above. What I do like about Goldman's model is that it does attempt to account for and constrain the truth-making process in social environments (groups, aggregates, etc.) in ways that more consensus-based approaches cannot.¹⁵ Thus it is more in line with the account of social epistemology that I will defend here, i.e., one that emphasizes the pragmatist notion of convergence.

There are a few issues remaining here. First, I have suggested that we need a "bounded" or "operational" account of truth tied to an objective decision procedure for scientific inquiry. But (1) why doesn't the pragmatist "method of science" or frequentist account of probability work? And, (2), how does the model that I'm suggesting resolve the issues with the theory/practice divide, and the "method of science" that the frequentist view does not? Lastly, (3) if the model that I'm proposing is successful, how can we aggregate data in a way that diverse scientific communities can benefit from the successful, truth-directed research of other communities?

¹⁴ See: Goldman (1999), pp. 116-117.

¹⁵ See: Goldman (1992), p. 186-188. Goldman spends little time on *consensualism* because he believes it is too "plagued with problems." This is primarily due to complications similar to those with *relativism*, i.e., how exactly to demarcate a "socially warranted" belief or how to establish a clear notion of "social warrant" full stop. However, the most damaging critique of consensualism is that we can too easily imagine generating a "tainted" consensus, and this could put the entire community's "belief-profile" at odds with common-sense rationality. Goldman notes the following: "Consensus *per se* is not a reliable sign of rationality. It depends on how consensus is reached. All sorts of methods can yield consensus: brainwashing, the threat of the rack or burning at the stake, totalitarian control of the sources of information. Consensus reached by these means does not guarantee rationality." I think the problems with consensus approaches to social epistemology can be seen in some of the projects in feminist naturalized epistemology; and, in particular, in the work of Lynn Hankinson Nelson (1993). Nelson argues that one of the primary aims of feminist naturalized epistemology is to reconceive the role of the individual knower as the "agent of epistemology" to conceiving of *agents* within "epistemological communities." The unit of epistemological significance therefore is the *entire* community. As Nelson claims: "My arguments suggest that the collaborators, the consensus achievers, and, in more general terms, the agents who generate knowledge are communities and subcommunities, not individuals". She is equally skeptical of the claim that there is "one truth," or a singular conclusion (or entity) that will fully verify the "truth" of our inquiries. Rather than help us in the process of knowledge-acquisition this assumption often is stultifying, making dogma out of the "conclusions" of a single (of a few) epistemological communities. The interesting point of tension here is not with Nelson's non-veridical approach, but with the lack of a constructive apparatus to prevent "tainted" consensus-building. Also see: Nelson, Lynn (1993): "Epistemological Communities," in *Feminist Epistemologies*, eds. Linda Alcoff and Elizabeth Potter. New York: Routledge, p. 123-124.

The argument for (1) comes mostly from a pragmatist account of scientific inquiry and the convergence thesis. The formal apparatus that supports pragmatist epistemology is a frequentist account of probability. What makes the frequentist account appealing is that it is most directly tied to the pragmatist convergence thesis; but, as I intend to show, other objective probability views also lend themselves to the convergence thesis as well, and without many of the problems common to the most well-known interpretations of frequentism. First, however, we might ask where the frequentist account of probability lies in terms of the distinction between epistemic and objective interpretations of probability?¹⁶ As it might be a fair criticism of classical pragmatism to say that an objective interpretation of probability, e.g., frequentism, requires making unjustified ontological claims, and this is an inherently problematic position for the pragmatist. Rather, the most natural probabilistic interpretation for the pragmatist seems to be a subjective account that conditionalizes on subjectively determined probability assignments, as this account does not make any serious (non-revisable) metaphysical assumptions. Given the problems with Goldman's (1999) subjectivist account, how do we determine which interpretation of probability is the correct one for our purposes here? To clarify, this is not to ask the normative question: which interpretation of probability should we/can we justify using? In this case, it is best to agree that different interpretations of probability serve different purposes given what they are being used for, i.e., to remain agnostic in terms of the normative question. For our purposes, we want to know how the naturalist—and the naturalist's commitment to the entities of science—and the pragmatist—and the pragmatist's commitment to an inductive metaphysics—can be squared probability wise.

Some contemporary critics (Haack and Kolenda (1977), Gillies (2000)) claim that the classical pragmatists, and, in particular, Peirce, moved from a frequentist theory of probability to a propensity theory of probability in light of a shift in his metaphysical commitments. Peirce's dislike of nominalism led him away from the frequentist theory of probability perhaps because of the ontological problem that I

¹⁶ Millstein (2003) brings up a similar issue in terms of deciding what probability measure is best for understanding evolutionary theory, and I think her use of Gillies (2000) distinction between "epistemic" and "objective" probability interpretations is the right place to start when deciding what probability interpretation is best to use. She writes: Philosophers generally divide interpretations of probability into two basic kinds: (1) epistemic (or epistemological) probability—probability that is concerned with the knowledge or beliefs of human beings, and (2) objective (or ontological) probability—probability that is a feature of the world (like the sun, the earth, etc.), independent of the knowledge and beliefs of human beings (Gillies 2000, 2). Using these (not entirely uncontroversial) definitions as a starting point, we can now ask whether the transition probabilities in ET are epistemic or objective...."

suggested above—it is difficult to strike a sound philosophical balance between a commitment to the physical entities that frequentism quantifies over and the pragmatist commitment to remain agnostic toward such metaphysical commitments. Haack and Kolenda (1977) write:

The more specific version of Peirce's argument for the self-corrective character of inductive reasoning has its roots in his account of probability. His early account of probability was straightforwardly frequentist; subsequently (as he became increasingly anti-nominalist) he added a propensity element, identifying the tendency of a ratio to a limit with a "would be," of, e.g., a die, (Fisch [1967] suggests this is one symptom of Peirce's conversion from an initial nominalism to a mature realism).¹⁷

The "would be" of Peirce's account essentially concerns the probability of a given event, say, the toss of a fair coin, should the event be carried out in an infinite series of tries. The relative frequency of the coin coming up tails will converge around 0.5, and will become increasingly determinate as the series of fair coin tosses moves toward infinity. Thus we will eventually be able to say that the "propensity" or "dispositional quantity" of the fair coin is that the objective probability of its coming up tails is 0.5.¹⁸ The problem with this interpretation, as Gillies (2000) notes:

It depends on the conditions under which the die is thrown, as is shown by the following two interesting examples of Popper's. Suppose first we had a coin biased in favour of 'heads'. If we tossed it in a lower gravitational field (say on the Moon), the bias would very likely have less effect and prob(heads) would assume a lower value...For the second example we can use an ordinary coin but this time, instead of letting it fall on a flat surface, say on a table top, we allow it to fall on a surface in which a large number of slots have been cut. We no longer have two outcomes 'heads' and 'tails' but three, viz. 'heads', 'tails', and 'edge', the third outcome being that the coin sticks in one of the slots. Further, because 'edge' will have a finite probability, the probability of 'heads' will be reduced. This example shows that not only do the probabilities of

¹⁷ Haack, Susan and Konstantin Kolenda. (1977). "Two Fallibilists in Search of the Truth," in *Proceedings of the Aristotelian Society, Supplementary Volume*, Vol. 51 (1977), pp. 63-104.

¹⁸ Gillies, Donald. (2000). "Varieties of Propensity" in *British Journal for the Philosophy of Science* 51 (4): 807-835.

outcomes change with the manner of tossing but even that the exact nature of the outcomes can similarly vary.¹⁹

The successful take-away from Peirce's account of probability, as Gillies (2000) suggests, is that we gain an account of "would-be" probability that does not require that the full sequence of an event is needed to have a(n) (approximate) probability measure. More than that, Gillies (2000) argues that a propensity interpretation of probability would give us a workable objective account without the ontological baggage of frequentism. So, (2), why not adopt a propensity model of probability, similar to Popper's or Peirce's later work? Why am I suggesting an operational account of truth tied to objective Bayesianism instead?

The bottom line is the following: Propensity views are far too ill defined to do any useful work here. The classical propensity view comes from Popper's (1957) early work, but, as Gillies (2000) suggests, the propensity view seems to have been developed by many philosophers and scientists simultaneously in order to handle the frequentist problem of single case probabilities. The single case problem is perhaps the biggest criticism against frequentism and it runs something like this: there is a large class of cases where the event that we want to assign a probability to is non-repeatable, for example, the last game of the 2011 World Series or the assassination of Abraham Lincoln, but a frequentist account of probability cannot tell us anything about how to assign probabilities in these cases.²⁰ This, critics (See: Hájek (1996)) argue, is the fatal flaw in the frequentist account of probability:

Certain statements are 'single case' in virtue of the very logical form: for example, universal generalizations and existential claims. Some people think that non-trivial (objective) probabilities attach to such statements—as it might be, 'the probability that all ravens are black is 0.9', or 'the probability that there exists tachyons is 0.1'. If there is sense to be made of such probabilities, then it is not the frequentist who can make it, for such statements only get one opportunity of being true or false...So the reaction might be: 'frequentism was never meant to handle cases in which there are no statistics, or only a single data point; but in decent-sized samples it work just fine.' This is the intuition that is encapsulated in the catchy but all-too-vague slogan 'Probability is

¹⁹ *Ibid.*

²⁰ See: Hájek, Alan (2010) on finite frequentism in "Interpretations of Probability", *The Stanford Encyclopedia of Philosophy (Spring 2010 Edition)*, Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/spr2010/entries/probability-interpret/>>.

long run relative frequency.’ The reaction is wrong-headed: problems remain even if we let our finite number of trials be as large as we like.²¹

Due to the problem with single case probabilities, many frequentists, like Popper and Peirce, adopt a propensity account of probability. However, the problems with objective probability do not stop by abandoning frequentism. Although it is tricky enough to define what a “propensity” is, it is even more difficult to show that any particular physical entity or event has a definitive propensity, or dispositional quantity at all. This criticism also holds for many other objective probability views that attempt to replace frequentism, including views of objective chance (or “chance theories”) and hybrid frequency-propensity views. I will argue, instead, that an objective Bayesian account is probably the best account that we can use as the formal apparatus to support the pragmatist convergence thesis. That is because an objective Bayesian account will give us two things that these other views cannot give us: (1) a probabilistic model that is epistemic, in that it deals primarily with degrees of belief, without being subjectivist, and (2) an account of belief that is constrained by—not determined by—what we know about the physical world at present. Thus we are not relying on metaphysical commitments to make probability assignments so much as we are employing them as our best guess to constrain our current beliefs. This is one of the three normative requirements for an objective Bayesian account of belief called “calibration:”

According to the version of objective Bayesianism presented in Williamson (2005), one’s beliefs should adhere to three norms:

Probability: The strengths of one’s beliefs should be representable by probabilities. Thus they should be measurable on a scale between 0 and 1, and should be additive.

Calibration: These degrees of belief should fit one’s evidence. For example, degrees of belief should be calibrated with frequency: if all one knows about the truth of a proposition is an appropriate frequency, one should believe the proposition to the extent of that frequency.

Equivocation: One should not believe a proposition more strongly than the evidence demands.

One should equivocate between the basic possibilities as far as the evidence permits.²²

²¹ Hájek, Alan. “Mises Redux”—Redux: Fifteen Arguments against Finite Frequentism in *Erkenntnis*, Vol. 45, No. 2/3, *Probability, Dynamics, and Causality* (Nov., 1996), pp. 209-227.

²² Wheeler, Gregory and Jon Williamson, “Evidential Probability and Objective Bayesian Epistemology,” in *Handbook of the Philosophy of Science, Vol. 7: Philosophy of Statistics*, eds., Prasanta S. Bandyopadhyay and Malcolm Forster, Elsevier: New York (2011), pp. 307-332: The Calibration norm

Here, we can employ frequencies in a way that we could not with frequentism. The entities postulated by this account of objective Bayesianism do not require a metaphysical commitment, but a normative commitment. This model will help us resolve issues with the theory/practice divide by giving us a formal apparatus with which we can say “ought” to constrain belief in practice.

Lastly, we need to consider how we can aggregate and model (e.g., visualize) data in a way that diverse scientific communities can benefit from the successful research of other communities. To this end, I will look to the work of List (2005) and List and Dietrich (2007, 2008) on judgment aggregation and group rationality, and Purchase’s (2008) work on theories of information visualization. I will suggest that a fully developed veritistic social epistemology with working theories of judgment aggregation and information visualization will give us a rich, network-level, or what Goldman (2010) refers to as “systems-oriented,” social epistemology.²³

1.2 The Motivation

One of the new terrains that Alvin Goldman (2010) isolates for social epistemology is the “epistemology of mass collaboration.”²⁴ Or, what we might call “community-based aggregated data.” For the information scientists, this is nothing new. The success of open-source and freeware tools, like WordPress and Wikipedia, has allowed for the creation of new technological spaces that not only use community-based data aggregation, but also analyze the reasons behind its success, e.g., Purchase (2008), Fallis (2008, 2009).²⁵ Simply put, the amount of accessible information that open-source and other internet-based tools created has opened up a new terrain for analytic epistemology. Additionally,

says that the strengths of one’s beliefs should be appropriately constrained by one’s evidence ϵ . (By evidence we just mean everything taken for granted in the current operating context— observations, theory, background knowledge etc.) This norm can be explicated by supposing that there is some set $E \subseteq P$ of probability functions that satisfy constraints imposed by evidence and that one’s degrees of belief should be representable by some $P \epsilon \in E$. Now typically one has two kinds of evidence: quantitative evidence that tells one something about physical probability (frequency, chance etc.), and qualitative evidence that tells one something about how one’s beliefs should be structured.

²³ See: Goldman (2010) for the originating discussion of “systems-oriented social epistemology.” For Goldman (2010), “systems-oriented social epistemology” is a branch—a more expansive branch—of veritistic social epistemology that focuses on the epistemic properties of “epistemic systems.” And an *epistemic system* “designates a social system that houses social practices, procedures, institutions and/or patterns of interpersonal influence that affect the epistemic outcomes of its members.”

²⁴ See: Goldman (2010), “Systems-level Social Epistemology,” in T. Gendler and J. Hawthorne, eds., *Oxford Studies in Epistemology*, vol. 3, 2010, p. 1.

²⁵ For example, see: Fallis (2008), “Toward an Epistemology of Wikipedia, *Journal of the American Society for Information Science and Technology*, 59, 10, (2008): 1662-74.

the ability to hack proprietary online spaces and access large amounts of data has generated alternative scientific communities of equal interest. The “biopunk” community, if it can be called such a thing, works in the name of the democratization of genetic engineering and biotechnology. Biopunks take their particular ethos from the hacking community in computer technology, particularly those who argue in favor of open-source versus proprietary software.²⁶ However, the trouble with democratizing genetic engineering, obvious to anyone who trades in science-fiction scenarios, is that our desire for knowledge may outrun the social and ethical “checks” on knowledge-production usually put in place by traditional academic and governmental institutions. Associated Press reporter, Marcus Wohlsen (2011), looks into the biopunk community in *DIY Scientists Hack the Software of Life*, and notes, “The key question in the democratization of genetic engineering is whether putting the tools and techniques of biotech into the hands of more people will tip that balance...Will more people operating with less supervision unleash biotech for better or worse?”²⁷ In an ideal world there should be a place for all scientific communities to engage in the truth-making process, and to do so without the often-limiting sanctions of institutional science—and without making “truth” and “objectivity” the only practical goals of inquiry. This would include a space for biopunks and institutional science, conventional lab culture, and community-based aggregated data.

I am a big fan of formal approaches, especially those that will allow for a diversity of scientific practices while still feeding into the knowledge-production process. This holds particularly for non-institutional science and community-based data aggregation. I argue that the convergence thesis that is offered by the classical pragmatists (particularly, C.S. Peirce), abridged to do practical work for scientific

²⁶ See: Marcus Wohlsen's *Biopunk: DIY Scientists Hack the Software of Life*, Penguin Group (New York), 2011. In this book, Associated Press reporter, Wohlsen, investigates the world of “do-it-yourself” scientists and “biohackers.” He offers the obvious “end of times” scenario usually thought of when non-institutional science is anywhere sanctioned—it goes something like this: “The disruptive power of DNA changes the terms of the open-source argument. In computer software, knowing the source code allows a hacker to make an app that serves good or destructive ends. The terminology comes straight from biology: Malicious code is a “virus.” A “contaminated” computer is “infected.” Compromised computer security can create havoc. A hacked defense system in theory could send missiles skyward. But the mad scientist worst-case scenario conjured by the idea of biohacking stirs more primal fears. Tomorrow's Dr. Frankenstein would not be building a human-sized monster to stalk the villagers in plain sight. Today's most promising technologies for reading and writing DNA stir worries that he or she would be in the kitchen synthesizing a microscopic superbug no one could see even after it was too late. And maybe the blueprint for that germ would start with the genetic code for a flu virus available to anyone with an Internet connection.”

²⁷ *Ibid*

communities via Bayesian decision theory, will do just that. It will provide a formal apparatus for philosophers and information scientists to analyze large-scale data by providing a model for belief-updating with an aim toward showing how the probability of any, p , will increase toward 1 (certainty) or decrease toward 0 (ignorance) over time. I argue that this method will also broaden the borders of what can be called “analytic epistemology” quite significantly, and for the better. My motivation is the desire to encourage debate between working scientists, social scientists and philosophers of science about the status (and nature) of scientific inquiry, and to make discussions less Balkanized in terms of who is capable of most substantively contributing to our communal knowledge. As Longino (1993) points out: “A community must not only treat its acknowledged members as equally capable of providing persuasive and decisive reasons and must not only be open to the expression of multiple points of view; it must also take active steps to ensure that alternative points of view are developed enough to be sources of criticism.”²⁸ What Longino (1993) has in mind is probably closer to an “inclusion hypothesis,” i.e., having more female and minority scientists involved in institutional science provides differing and often necessarily critical points of view about a hypothesis, h , that wouldn’t have otherwise arisen. However, there is also a broader point here. It is not enough to have the support of institutional science, we also need to extend our epistemology to a diversity of scientific projects and practices—even those that fall outside of what we might normally call “science” in the first place.

The second motivation for this project comes from a deep dissatisfaction with traditional analytic epistemology and its focus on the individual agent as the ultimate unit of epistemic analysis. As many feminist epistemologists have argued, there is something isolationist and unappealing about continuing the project of analytic epistemology as a continuum of Descartes’ Meditations. Epistemic agents are situated agents, and knowledge-production—as the paradigmatic case of lab science illustrates—is unlikely to be an individual affair. Any account of analytic epistemology should not only recognize this problem, but it should make a positive contribution to changing the way that we do epistemology in general. This is just good common sense.

1.3 The Payoff

²⁸ Longino, Helen (1993) in *A Mind of One’s Own: Feminist Essays on Reason and Objectivity*, eds. Louise M. Antony and Charlotte Witt, Westview Press: San Francisco, pp. 267-268.

Generally speaking, the ultimate upshot of this project will be to broaden our notion of what counts as traditional analytic epistemology. Too often, what happens to those who challenge the status quo notions of “truth” and “objectivity” by pointing to the fact that knowledge and knowledge-production are “situated” are accused of not doing “real” epistemology. Even Goldman (1999) has had to defend his account of analytic social epistemology from those who think that it’s not truly epistemology at all (See: Alston (2005)). The accusation typically goes like this: There are projects that are central to our notion of traditional epistemology—and these projects are what defines and gives scope to epistemology—thus projects that fall outside of this are not really epistemology at all. Alston, for one, suggests that Goldman (1999) is really doing something more like sociology of science than epistemology. Of course, this view of what counts as epistemology proper comes perilously close to being question-begging. If we define epistemology as exactly what epistemology already is then we are not really saying anything at all about counts as epistemology other than what we can already point to. As I see it, one big payoff for broadening what counts as epistemology is getting rid of these types of narrow, circular characterizations of epistemology. An additional payoff is that by broadening what counts as epistemology, we give credence to views often promoted by feminist epistemologists, feminist philosophers of science, and those who argue for non-traditional, consensus-based views of epistemology, e.g., Rorty (1979). These views and others like them will be sanctioned by way of the broadened scope of the discipline itself. Now, it’s very easy to counter that this is not really a necessary step to promoting diversity within academic epistemology, i.e., epistemology done within the academy, because many non-traditional epistemological views are already available to us. However, I would argue that there is a substantial difference between the alternative views being available and alternative views contributing to the same discussion about knowledge-production and truth-making. Having a voice on the side of the discussion does not always allow that voice to be considered as a valuable critique in the larger discussion. And, as I (and Longino (1993)) already mentioned, diversity in epistemological viewpoints, and mandating alternative views, and alternative criticism, is an essential component to knowledge-production.

1.4 The Methodology

A good portion of this project will use anti-subjectivism and a deep commitment to methodological naturalism as its philosophical foundation. By “anti-subjectivism” I mean that the goal for any ideal belief

state intended for inquiry is one where our individual preferences and prejudices have been constrained as much as possible by objective evidence. The formal apparatus for epistemic constraint is, as mentioned, an objective Bayesian model attached to a revisionistic metaphysics. This permits the assumption that methodological naturalism is the correct jumping off point for epistemology. That is, we can take our (metaphysical) concepts from the entities postulated by successful inquiry in science. However, we should also allow that such entities could be (and perhaps will be) revised by continuing scientific inquiry. Our metaphysical commitments are, at best, our best guesses and only serve (for my purposes here) to constrain our beliefs. I take it that an alignment with “good” methods of science—as far as that will currently get us—is simply an alignment with good common sense. However, I am not dogmatic about what should (or will) count as “good” methods or “good” science.

I'll admit that the line that I am taking is a bit tricky because the justification for the naturalist's alignment with social epistemology rests primarily on the way that “good” science is conducted. However, to define “good” science as essentially social, communal, or consensus-based, comes dangerously close to begging the question against social epistemology. On this account, it's best to justify this claim with an empirical analysis of the behaviors of different scientific communities and leave the normative claims about “good” science aside. This means, of course, that we can only be weakly justified in suggesting what “good” science is; we can point to what is the case with working science and to suggest (mildly) that it at least has been successful.

1.5 Project Outline

The second chapter of this project is dedicated to a brief argument for a broader notion of science as it relates to methodological naturalism as well as a nod to the argument for the necessity of methodological naturalism's commitment to social epistemology. Methodological naturalism cannot be successful from within the confines of traditional, individualistic epistemology. To respect the thesis that methodological naturalism aligns itself with the best methods in science requires us to affirm what most scientists and few analytic epistemologists are ready to concede: knowledge is never produced in a social, cultural, or political vacuum. Knowledge is produced in a world where the epistemic subject is already contextually situated and, throughout all his or her developmental stages is immersed in a world rich with social content. I will not give a positive argument for the abandonment of individualistic

epistemology, but I will briefly suggest that an honest naturalist must leave behind the Cartesian legacy of the possibility of individualist knowledge. In the third chapter, I will analyze contemporary arguments for social epistemology. I will briefly discuss the presumed historical origins of the project of social epistemology and discuss what bearing these origins have on my project here. I will start the contemporary analysis of social epistemology with a discussion of what can be called, “radical social epistemology,” where the primary goal of the account is to dismantle not only traditional notions of epistemology but epistemology as a worthwhile philosophical pursuit full stop. I will offer a good argument against the project of the radical social epistemologist, and suggest ways in which the project may be saved. The most compelling way to preserve the radical project is by abridging it into a consensus-based social epistemology. However, as I intend to show, other consensus-based projects, supported by Rorty (1980), Nelson (1993), and Longino (1993) are equally problematic. I hope to show that the lack of truth-direction hurts the projects of Longino (1993) and Nelson (1993), and that the alternative account of truth and objectivity given by Longino cannot avoid falling into similar traps as radical social epistemology. I will ultimately suggest that the best truth-directed account of social epistemology comes from Goldman (1999, 2010).

The fourth chapter will be devoted to breaking down Goldman’s (1999) veritistic social epistemology. I will outline the motivations for his account of social epistemology and explain how he extrapolates from a Bayesian account of testimony to a Bayesian account of social epistemology that includes science. Specifically, I will tackle how Goldman (1999) weds his this account of testimony to subjective Bayesian decision theory, and what problems arise from this wedding. I will outline the criticism of his use of Bayesianism by Kitcher (2002), Fallis (2002), and Talbott (2002), and briefly suggest an alternative approach to his veritistic social epistemology, by way of pragmatist analysis, that do not generate the same problems with scientific knowledge.

In the fifth chapter, I will explain what is entailed in a pragmatist account of epistemology, and what a “pragmatist framework” for social epistemology should look like. First, I will discuss some attempts to understand pragmatist epistemology as a “consensus-based” social epistemology and explain (again) why these accounts fail on two levels. A consensus-based social epistemology cannot be suitably truth-directed and that this is problematic because (1) without truth-direction, consensus-based social

epistemology will fall into a radical social epistemology, and radical social epistemology is untenable, and (2) a consensus-based epistemology does not respect the pragmatist commitment to truth and objectivity in inquiry. Second, I will discuss the criticism that pragmatism is fundamentally anti-realist and the debates that the classical pragmatists themselves had concerning their respective metaphysical commitments. I will ultimately claim that although there is some disagreement between the metaphysical commitments of the classical pragmatists, they are all generally committed to an induction based, revisable metaphysics and epistemology.

The pragmatist account of epistemology, however, rests primarily with a fully developed “method of science” tied to the pragmatist convergence thesis. The convergence thesis, familiar to many subjective Bayesians, claims that scientific inquiry (via investigation of any given scientific proposition) carried far enough will eventually lead to probability assignment 1 or 0 for that proposition. What Peirce and Dewey refer to as the “method of science” is the formal mechanism that will eventually generate such probability assignments. The pragmatists assumed that it would be something like a frequency or propensity account of probability, but this can be successfully changed to another objective probability model that still respects the pragmatist convergence thesis. I will also show that those who claim that the pragmatists’ convergence thesis is too idealized to work (See: Hookway 2004) are generally underestimating the crucial nature of a formalized “method of science” in pragmatist epistemology—a method that can suitably constrain the inquiry process. The idealization problem arises because of a need for better epistemic constraints on the method. As Fuhmann (2006) notes: “If X and Y could be brought to exercise the same choices—if, in other words, they would eventually adopt the same selection function...then given sufficient exposition of evidence, their states of belief will indeed converge in the limit.”²⁹ Here the solution seems obvious. Epistemic constraints on inquiry should serve as a normative constraint on believing—that is (or will be) the “selection function” that each inquirer should adopt. Thus we can say—hopefully—that if a given community of inquirers respects the epistemic constraints on inquiry—which, in my view, will be non-frequentists objective probabilities—then their selection function should be identical. I will claim that this is how we can save the pragmatist convergence thesis.

²⁹ Fuhmann (2006).

Lastly, in chapter five, I will discuss the initial account of the “method of science” as formulated by Peirce and then later by Dewey. I will discuss, as mentioned above, the problems with a frequentist account of probability and, to some extent, the problems with Peirce’s later adoption of a propensity account of probability. I will suggest that the “method of science” actually cannot work using either of these accounts and that this is problematic for (1) the convergence thesis and (2) for generating workable epistemic constraints on inquiry in practice. Recall that the question that we want to ask concerns epistemic constraints on inquiry: What ought to constrain our beliefs? Or, perhaps, more specifically – how strongly ought we believe in proposition, p ? The frequentist will say that a frequency-distribution is what should constrain belief, as well as determine exactly how strongly you should believe in p . (A propensity theorist will make a very similar claim here). The subjectivist will claim that such determinations are up to the individual agent, and that we cannot construct non-arbitrary constraints on prior probabilities. Less arbitrary epistemic constraint on belief can take the form of betting arguments, e.g., synchronic Dutch-book arguments. Unlike the frequentist and/or propensity account of probability, subjective Bayesianism is a more compelling account of probabilistic epistemology and has a stronger following in the formal epistemology literature. However, much of the criticism of subjective Bayesianism is that its inability to objectively constrain priors leads to too many unintuitive posterior probabilities. As Joyce (2010) notes:

There is no getting around the fact that Bayesianism is a garbage-in-garbage out enterprise: if one applies the apparatus using a prior that is accurate and well justified, the conclusions derived will be accurate and well justified as well; if one applies the apparatus using a prior that is inaccurate or unjustified, the conclusions will also be inaccurate or unjustified. This is the heart of frequentist misgivings. From the perspective of frequentist statisticians, Bayesian methods carry a massive uncollateralized risk or error.³⁰

However, there are those who argue that the subjectivist position is well worth the unintuitive (and, in some cases, wholly irrational) consequences. This is primarily because of the idea that, given enough time, there will be a “washing out” of unintuitive priors. This involves a type of convergence thesis for the

³⁰ Joyce, James M. “The Development of Subjective Bayesianism” in Dov Gabbay, Stephan Hartmann and John Woods, eds., *Handbook of the History of Logic. Volume 10: Inductive Logic*, Elsevier 2010: 415-476.

subjective Bayesian. The criticism of this runs similar to the criticism against the pragmatist convergence thesis. Namely we do not have an infinite span of time to make sure that our belief-forming method is the correct one and/or that our beliefs can be adequately verified over time. Thus, in the meantime, the subjectivist ends up sanctioning a good amount of irrationality in belief-formation. This is a critical part of the “problem of priors” and this may be reason enough to favor objective probability theories over subjective theories like subjective Bayesianism. As Joyce (2010) notes:

Frequentist statisticians see Bayesians as rogues...They are skeptical of Bayesian methods because they doubt that prior probabilities can be made epistemologically respectable, and feel that their introduction threatens to undermine the accuracy and objectivity of our inductive reasoning.³¹

An objective Bayesian account is one that satisfies the Kolmogorov probability axioms:

1. (Non-negativity): $P(A) \geq 0$, for all $A \in \mathcal{F}$.
2. (Normalization): $P(\Omega) = 1$.
3. (Finite additivity): $P(A \cup B) = P(A) + P(B)$ for all $A, B \in \mathcal{F}$ such that $A \cap B = \emptyset$.³²

And constrains prior probabilities via some objective, or evidential, feature of the world, e.g., objective chances, frequencies, propensities, and/or generally available evidence, etc. Many objective Bayesians (See: Williamson (2005)) also argue that a truly objective account must adhere to three additional normative requirements: (1) probability, i.e., that one’s beliefs should be represented by probabilities, (2) calibration, i.e., that those probabilities should be in accord with the available evidence, and (3) equivocation, i.e., if calibration is not possible due to issues with available evidence, the probabilities should equivocate between “basic outcomes.”³³

The account of convergence that I want to defend here is one that utilizes objective Bayesianism as its formal apparatus (or what the pragmatists’ would call the “method of science”) to constrain beliefs in a broader, wholly social, knowledge generating process. The problem that arises with the pragmatist conception of truth as convergence in its original form does so because of the pragmatist tendency to

³¹ *Ibid.*

³² Hájek, Alan. "Interpretations of Probability", *The Stanford Encyclopedia of Philosophy (Spring 2010 Edition)*, Edward N. Zalta, ed. <<http://plato.stanford.edu/archives/spr2010/entries/probability-interpret/>>.

³³ Williamson, Jon. (2010). *In Defence of Objective Bayesianism*. Oxford University Press: New York.

establish a theory/practice divide. I will suggest along with Pihlström (2004) that this has happened, in part, due to James and Dewey's adoption of a Peircean view of inquiry in conjunction with decidedly more pronounced views about human behavioral psychology. I will suggest, again along with Pihlström (2004), that the primary differences here in the views of truth are not in kind but of degree of emphasis. I will argue that although the pragmatists focused on the method of science ("pure science" or "frequency-distribution") to arrive at truth in the abstract, they did not explicitly show how we will arrive at truth in the case of the concrete or everyday "particulars." Here, I will revisit the claim that the root of this discrepancy actually lies in a metaphysical debate about universals (abstractions or "generals") and particulars. In particular, I will focus on Peirce's fierce refusal to accept nominalism even as an instrument to make our behavioral practices more conducive to explanation—as James and Dewey used them—and how this debate may account for the differences in the pragmatists' conception of everyday inquiry. I will argue that what is called for here is better account of how we can arrive at truth in everyday inquiry, in a way that (1) does not invoke the unnecessary metaphysical debate about universals and particulars, and (2) puts epistemic constraints on everyday inquiry in a way similar to what is called for in the classical pragmatist "method of science."³⁴ Thus we can have an objective model to constrain belief while still engaged in a wide range of inquiry for theory-building.

The final issue that I will discuss in chapter six has to do with how such wide-ranging types of inquiry, as well as far-reaching communities of inquirers, can make practical use of the results of inquiry in different scientific communities to potentially achieve belief convergence in a practicum environment. I will outline some popular theories in science studies particularly as they pertain to mapping the institutional activity of science by way of scientific citation data. I will discuss some methods of data aggregation and data visualization that make use of scientific citation data in order to show emerging and waning trends in scientific research. I will conclude this chapter with what I believe to be information science's most beneficial contribution to social epistemology – providing ways to map, measure, and visualize the various activities of scientific communities. I will also suggest citation studies as a first step

³⁴ For this account, objective likelihoods are constrained by what we know at *present* about the probability (frequency—or "chanciness") of any given hypothesis. This account is "operational" in that all likelihoods will inevitably change (even if the probability of a hypothesis—say, the sun will rise tomorrow, is .99, it will still converge toward 1) as inquiry continues.

for contemporary analytic epistemologists in terms of how they can successfully tackle questions of knowledge-production and analysis in a world of large-scale data production.

In this very brief chapter, I will argue for that the possibility of a “network-level” epistemology that focuses on analyzing large-scale data, and their resulting “data-networks,” using primarily scientific citation studies, scientific citations themselves, and methods of data visualization. Of all of the chapters offered here, this one will perhaps be the most theoretical but will also employ the most recent tools in information and data studies. I will propose a model for what I think a network-level epistemology could look like, given that it is possible to meaningfully analyze (re: for epistemological purposes) large amounts of scholarly data in the form of citations. I will present Goldman’s (2010) “Systems-Oriented Social Epistemology,” which, in large part, also inspires the approach that I will offer here. Goldman (2010) looks at systems-oriented social epistemology as a form of “epistemological consequentialism” where successful outcomes, re: truth-production, are looked at via the social epistemological systems that produce them. This account, although interesting, cannot tell us what we really want to know about how various systems produce knowledge. Goldman’s account (2010), similar to a number of consequentialist accounts, is told in hindsight and without adequate explanation of how the social method and/or social system produced the desired result, which, in this case, is truth. This is particularly the case when dealing with information as large and diffuse as social information. (There is also the problem that unlike consequentialism in, say, ethics, the epistemologist has a set of values that he or she wants the method to maximize that may not have a corollary in terms of prior expectation. That is, we may have some prior concept of a “good” that we wish to maximize, but the same may or may not be said concerning a prior concept of “truth” that we wish to maximize.) How, for instance, does any system isolate data conducive to forming knowledge-claims from “junk in the system”? I take these issues up, along with some others, in this final chapter. I will argue that no matter which network account proves the most problem-free, any reasonable account will need to tackle issues concerning the relation of the data to the institution(s) that produces it. That is, what is the political status of the institutions that produce the units of data that contribute to large-scale social knowledge? Does this status matter? As Coady (2012) claims: do we need to consider if our social epistemology, once rightly conceived, requires an ethics?

Chapter II

On Naturalism

2.1 Methodological Naturalism and the Scope of “Science”

A good place to start thinking about methodological naturalism is to first ask the question of what exactly separates methodological approaches to naturalism from more ontologically robust approaches. I follow Price (2004) in thinking that too often “naturalism” or “philosophical naturalism” is a view asserted without much concern for the more serious details of its adoption. This includes specification of (a) its epistemological relation to the claims made by science and scientific practice, and/or specification of (b) its ontological relation to the claims made by science and scientific practice. As mentioned in Chapter 1, I do not have much interest in ontological approaches. I see no reason for any naturalistic account to over-commit itself to only those entities as described by science. For one, the history of scientific practice has shown us many times how faulty such a logic can be, e.g., consider the case of phrenology or of eugenics. The naturalism that I prefer is of the “methodological” stripe. That is, philosophy ought to align itself with the best truth-producing *methods* or practices of science, rather than with a wholesale adoption of its ontological objects. We might say, as Kuhn (1962) does, that methods (i.e., puzzle-solving) are our best measure of success even if the relevant scientific *objects* are constantly up for scientific, social, and political debate. Although I would not want to say, as Kuhn (1977) does, that with our loosening grip on the metaphysical objects of science we equally loosen our grip on a metaphysically robust account of *truth*. I simply think that any concerns with metaphysics ought to come second to concerns about veritistically adequate scientific methods.¹ I will talk more about this in later chapters. Suffice to say that this is a parting of the ways with many naturalists (and contemporary pragmatists) that the rest of my project will aim to defend. For the moment, however, something has to be said about *how* we can break apart our naturalistic concerns with methods from those of metaphysics. So I start by introducing a

¹ In fact, for Kuhn (1977) the notion of approaching “truth,” or of science reaching some idealized point, was thought to be absurd. There is no notion of “truth” in this sense in Kuhn’s (1977) views at all. In later interviews, Kuhn repeatedly clarified that the philosophical preoccupation with “truth” did not really play a role in how science is either understood or conducted. See: Kuhn (2000), eds. Conant and Haugeland. I will argue in later chapters that we can (1) separate our epistemology from our ontology, as the pragmatists did, and (2) because of (1) we can still focus our interests in methodological naturalism on epistemological issues.

distinction made by Huw Price (2011) between what he calls *subject* naturalism and how we typically understand “naturalism” proper, or *object* naturalism.²

According to Price (2011), object naturalism presupposes the potential mapping of scientific concepts on to the world much like putting a predetermined set of stickers into corresponding shapes in a children’s sticker book.³ Price (2011) rejects this view in favor of what he calls a semantic *subject* naturalism. One of the problems with object naturalism is that it gives us no way of meaningfully talking about concepts that are not currently covered by scientific explanation, i.e., it only allows us to talk about the concepts that are somewhere represented in the sticker book. Now he asks us imagine that we are looking at the world on one side and “all the statements we take to be true of the world” on the other side; what can we say makes each statement on the one side *true* about the world on the other side, the fact that there is a corresponding shape?⁴ In both cases, what happens when there are limitations with respect to the corresponding shapes? I.e., What happens when we have more stickers, or more statements, than we have “truthmakers” and what does this mean for naturalism? Price (2011) writes:

If the matching model is to be incorporated into a scientific perspective, the perspective itself seems to dictate the shape of the available facts and truthmakers. Roughly, the available shapes are the kinds of outlines recognised by natural science. Why does this turn out to be a severe constraint, at least *prima facie*? Because there seem to be many true statements that don’t line up neatly with any facts of the kind uncovered by natural science. Indeed, the problem cases are not just the classic misfits, such as the (apparent?) truths of aesthetics, morality, and other normative matters, or those of consciousness. Arguably, at least, they include matters much closer to a scientist’s heart, such as probability, causation, possibility and necessity, and conditional facts of various kinds; and even, hovering above all, the heavenly truths of mathematics itself. Thus there is a striking mismatch between the rich world of ordinary discourse and the sparse world apparently described by science. A great deal of work in modern philosophy amounts to attempts to deal with some aspect or other of this mismatch.⁵

² See: Price (2011).

³ Price (2011), p. 3.

⁴ *Ibid.*

⁵ Price (2011), pp. 4.

This representational “Naturalism” (or “Big-N” naturalism) aims to be the paradigmatic example as employed by contemporary epistemologists and philosophers of science, and Price (2011) claims that it is *this* version of naturalism that is riddled with problems and hopeless confusions. The first part of which results from the matching problem (or “proto-theory”) described above that tries to align the sentences we think true *of* the world *with* the world, and the second of which results from the idea that science provides (or will provide) all the true facts to be known about the world. Even though these facts, or “truthmakers,” are underdetermined. This is basic naturalist dogma, Price (2011) claims, and it is this dogma (or, as he calls, “mantra”) in combination with the proto-theory that purportedly creates most of the issues in contemporary naturalism.⁶ But what happens if the proto-theory turns out to be inaccurate? It is here that he aims to give an alternative account of naturalistic linguistic behavior which conflicts, and thus undermines, the proto-theory account. Hence the distinction between two types of naturalism: object naturalism and subject naturalism. In his view, *object* naturalism is “Big-N” naturalism which is committed to the proto-theory (or what he calls “Big-R,” *Representationalism*, in what looks to be the Rortyan “mirror” sense) and *subject* naturalism is a philosophically prior view that “begins with the realisation that we humans (our thought and talk included) are surely part of the natural world,” such that we do not need to start with the metaphor of mapping (ourselves, our thoughts and talk) onto the physical world as described by science.⁷

Of course, we might push back on Price’s (2008) characterization by asking about what examples there are of concepts outside of “the world-as-studied-by-science”? Or, rather, what does he have in mind in terms of these concepts? This is because Price’s (2008) worry seems to arise from a fear that “Naturalism” requires that science dictate our concept-reality correspondence and, more importantly, that that concept-reality correspondence be a static relationship. If that is the case, we can easily see what generates his concern that science may not cover all of the concepts we wish to meaningfully use in order to get along in the world. But there is no reason for a “Naturalist” to commit to that. One useful counterexample might be Kitcher’s (2012) argument for giving a new sense to the correspondence theory of truth. Kitcher (2012) inspired by James, argues that we can give a new understanding to the pragmatist view of correspondence so that the “static” relationship between concept and reality allows for

⁶ Price (2011), p. 5.

⁷ Price (2011), p. 5.

both cultural (e.g., the attitudes of scientific institutions and lab culture) and epistemological progress.⁸ Kitcher (2012) claims, and quite correctly, that James's theory of truth (what one might call a "psychologized" version of Peirce's) does argue for an "agreement" between objects and reality. Kitcher (2012) takes this to mean that James is not rejecting the "dictionary idea of truth as correspondence (or agreement)" and, in this sense, his point seems unproblematic. In looking to give sense to the correspondence relationship, Kitcher (2012) aims to marry a "modest" correspondence theory to a Tarskian framework. He claims that this way "we can give James's approach to truth a clearer formulation than *Pragmatism* achieves."⁹

In sketching this new formulation, where our sentences pick out, or represent, "chunks" of reality, Kitcher (2012) also aims to maintain the Jamesian idea that "truth happens to an idea." He accounts for this "truth-making" element by arguing that our sentences are still subject to "world-adjusting success," or "success that accrues to schemes representing part of reality—graphs, equations, maps, diagrams, as well as descriptive statements—in virtue of the systematic kinds of interventions they allow."¹⁰ This allows for the Peircean idea that beliefs get better, or more accurate, over time. Kitcher (2012) suggests, utilizing James's view, that it is our capacities and interests that divide up reality into its respective parts and that there are better and worse ways of doing so—or better and worse ways of "truth-making."¹¹

Kitcher's (2012) understanding of metaphysics, like Price's (2008, 2011), is inspired by classical pragmatism. Although, interestingly enough, their most prominent disagreement occurs at the juncture where Price (2011) makes his main objection to contemporary naturalism (or "Naturalism"). Perhaps giving a new sense to the "correspondence" relationship would ease some of Price's (2008, 2011) worries, but I'm inclined to think not. The question that I have concerns, at least in part, this description of a "static" relationship between our sentences and reality, or "chunks" of reality. Although Kitcher (2012) does try to preserve the pragmatist idea of truth *happening to* or truth *progressing toward* a point with the idea of "world-adjusting success," his view also reiterates a metaphysical relationship between "knowers" and the world that presupposes a certain metaphysics has already obtained. And it becomes our job as

⁸ See: Kitcher (2012), *Preludes to Pragmatism*.

⁹ Kitcher (2012), p. 131.

¹⁰ Kitcher (2012), p. 135. I leave out the discussion of "cultural success" here, not because it's not an important distinction, but for reasons of clarity and brevity on my part.

¹¹ Kitcher (2012), p. 137.

knowers to get in the right (or “best”) relationship to that reality. However, James argues that it is philosophically suspect to assume that reality is even “out there” to correspond to *outside of what can be verified by our experiences*—or verified epistemologically. For the classical pragmatists, the work of epistemology is both practically and theoretically *prior* to any claims about determinate truth and/or claims about *being*. The debate over correspondence that the Jamesian “pragmatist” has with the “intellectualist” in *Pragmatism* concerns just this issue. The pragmatist does insist on correspondence, but there is no correspondence in the abstract. All correspondence pertains to concrete particulars. My belief, *p*, corresponds with that object, *y*, *which I am experiencing*.¹² In this sense, reality is constructed, yes, and solipsism about reality cannot be avoided (as James duly notes), but what is verified by our epistemological work also *verifies* that part of reality about which our beliefs pertain.

For philosophers with metaphysical concerns, the idea of epistemology being prior to metaphysics may appear a bit too much like constructivism or anti-realism.¹³ To “right” the situation and to avoid claims of anti-realism it is tempting to reverse the picture, as Kitcher (2012) does, and make abstract (in the Jamesian sense) the ontological objects in our correspondence theory. Agree or not, for James, this is only to be in methodological error. He writes:

Pragmatist truth contains the whole of intellectualist truth and a hundred other things in addition. Intellectualist truth is then only pragmatist truth in posse. That on innumerable occasions men do substitute truth in posse or verifiability, for verification or truth in act, is a fact to which no one attributes more importance than the pragmatist: he emphasizes the practical utility of such a habit. But he does not on that account consider truth in posse, —truth not alive enough ever to

¹² James (1909) writes: “The pragmatist, being himself a man, and imagining in general no contrary lines of truer belief than ours about the ‘reality’ which he has laid at the base of his epistemological discussion, is willing to treat our satisfactions as possibly really true guides to it, not as guides true solely for US. It would seem here to be the duty of his critics to show with some explicitness why, being our subjective feelings, these satisfactions cannot yield ‘objective’ truth. The beliefs which they accompany ‘posit’ the assumed reality, ‘correspond’ and ‘agree’ with it, and ‘fit’ it in perfectly definite and assignable ways...If our critics have any definite idea of a truth more objectively grounded than the kind we propose, why do they not show it more articulately? As they stand, they remind one of Hegel’s man who wanted ‘fruit,’ but rejected cherries, pears, and grapes, because they were not fruit in the abstract. We offer them the full quart-pot, and they cry for the empty quart-capacity” (192-193).

¹³ As James (1909) writes, “‘When you say the idea is true’—does that mean true for *you*, the critic, or true for the believer whom you are describing? The critic’s trouble over this seems to come from his taking the word ‘true’ irrelatively, whereas the pragmatist always means ‘true for him who experiences the workings.’ ‘But is the object *really* true or not?’—the critic then seems to ask,—as if the pragmatist were bound to throw in a whole ontology on top of his epistemology and tell us what realities indubitably exist. ‘One world at a time,’ would seem to be the right reply here” (177).

have been asserted or questioned or contradicted, to be the metaphysically prior thing, to which truths in act are tributary and subsidiary. When intellectualists do this, pragmatism charges them with inverting the real relation. Truth in posse MEANS only truths in act; and he insists that these latter take precedence in the order of logic as well as in that of being.¹⁴

I take James to mean that potential truths, truth *in posse*, cannot be metaphysically prior to truth *in esse*. This assumption constitutes the inversion of metaphysics and epistemology in terms which one takes logical precedence. This inversion may also lead us to make erroneous claims about the role of metaphysics in a pragmatist “correspondence” view of truth. That said—we may still have good reason to reject James’s argument about truth and ontology here. If we take his claim about the inversion of logical priorities seriously, it may preclude our making a coherent argument about pragmatist “correspondence” or “agreement” altogether. But this is a larger discussion than what I want to discuss here. I simply raise it as a point of interest, and perhaps, as a point of concern.

Brandom (2011) also offers some interesting insights into Price’s (2011) arguments about naturalism, and, in particular, insights into the issue concerning the scope of science as relevant to subject naturalists (or naturalists of a similar sort). “Big-N” naturalism is of primary concern to Price (2004, 2011) because it fails to account for the fact that our linguistic practices seem to outstrip, or *outrage*, the world (“facts”) described by contemporary science. Brandom (2011) relates this argument to his own view where, with regards to linguistic practices, our foremost task is—to borrow from his own phrase—to make said practices “explicit.” As Brandom (2011) writes: “I think of analytic philosophy as having at its center a concern with semantic relations between what I will call ‘vocabularies.’ Its characteristic form of question is whether and in what way one can make sense of the meanings expressed by *one* kind of locution in terms of the meanings expressed by *another* kind....”¹⁵ Most semantic views, he claims, are steeped in a kind of logicism (“*semantic logicism*”) where logical vocabulary is given priority in clarifying the semantic relations between a given *target* and *base* vocabulary.¹⁶ However, what Brandom (2011) takes from the influence of classical pragmatism on

¹⁴ James (1909), pp. 205-206.

¹⁵ Brandom (2011), p. 158-159.

¹⁶ See: Brandom (2011), p. 159. Here he offers Russell and Whitehead’s *Principia Mathematica* as well as Frege’s *Grundgesetze* as paradigmatic examples of “semantic logicism,” (*see footnote*), and broadly specifies two philosophical programs as examples of using a logical vocabulary to clarify the semantic

contemporary philosophy is the shift in emphasis (in terms of semantic analysis) from *meaning* to that of *use*. But what does this actually entail? It first means we start with the new assumption that “the only explanation there could be for how a given *meaning* gets associated with a vocabulary is to be found in the *use* of that vocabulary.”¹⁷ That is, we should aim for a kind of “semantic pragmatism” where the fundamental relation of concern is the one between our “practices” or “practical abilities” and use. And, second, we speculate as to what relations need to hold between vocabularies and our “practices” and/or “abilities.” Brandom (2011) argues that there are two basic relationships of sufficiency that must hold between our “practices-or-abilities” and our vocabulary. The first is the relationship of “PV-sufficiency,” which is relationship that holds between our practices-or-abilities and our vocabulary such that we can say that someone can be said to be *using said vocabulary*: “It obtains when engaging in a specified set of practices or exercising a specified set of abilities is sufficient for someone to count as *deploying a specified vocabulary*.”¹⁸ The second relationship of sufficiency, “VP-sufficiency,” is the relationship between the vocabulary and a set of practices-or-abilities such that we can say that *that vocabulary* can be used to say something.¹⁹ And although relationships of sufficiency can be explained in terms of each other, i.e., where the most basic relation will be a “*pragmatically mediated semantic relation* between

relations between a “target” and a “base” vocabulary: empiricism and naturalism. He writes: “If we ask *which* were the vocabulary-kinds whose semantic relations it was thought to be important to investigate during this period, at least two *core programs* of classical analytic philosophy show up: *empiricism* and *naturalism*. These venerable modern philosophical traditions in epistemology and ontology, respectively, were transformed in the twentieth century, first by being transposed into a *semantic* key, and second by the application of the newly available *logical* vocabulary to the self-consciously semantic programs they then became. As *base* vocabularies, different species of *empiricism* appealed to phenomenal vocabulary, expressing how things appear, or to secondary-quality vocabulary, or less demandingly, to observational vocabulary. Typical *target* vocabularies include objective vocabulary formulating claims about how things actually are (as opposed to how they merely appear); primary-quality vocabulary, theoretical vocabulary; and modal, normative, and semantic vocabularies. The generic challenge is to show how what is expressed by the use of such target vocabularies can be reconstructed from what is expressed by the base vocabulary, when it is elaborated by the use of logical vocabulary” (159).

¹⁷ Brandom (2011), p. 166.

¹⁸ *Ibid.*

¹⁹ This view does not presuppose that we can adopt a system of universal or “*ur-*” practices to analyze. All of the practices or practical abilities under consideration should be relativized to a vocabulary. Brandom writes: “Talk of practices-or-abilities has a definite sense only insofar as it is relativized to the vocabulary in which those practices-or-abilities are specified. And that means that besides PV-sufficiency [Note: That is, “practice-vocabulary sufficiency”], we should consider a second basic meaning-use relation: ‘vocabulary-practice sufficiency,’ or just ‘VP-sufficiency,’ is the relation that holds between a vocabulary and a set of practices-or-abilities when that vocabulary is sufficient to *specify* those practices-or-abilities. VP-sufficient vocabularies that *specify* PV-sufficient practices let one *say* what it is one must *do* to count as engaging in those practices or exercising those abilities, and so to deploy a vocabulary to *say* something. PV-sufficiency and VP-sufficiency are two basic *meaning-use* relations (MURs)” (166-167).

vocabularies,” such relations are ultimately *not* relationships of “definability, translatability, reducibility, and supervenience.”²⁰ For Brandom (2011), this is the “pragmatic” part of how we should understand meaning comes from his view of classical pragmatism (and particularly of Dewey and James) as being primarily *instrumental* about the analysis of meaning. He argues:

[The classical pragmatists] manifest their endorsement of methodological pragmatism by taking it that the point of our talk about what we mean or believe is to be found in the light it sheds on what we *do*, on our habits, our practice of inquiry, of solving problems and pursuing goals. They manifest their endorsement of semantic pragmatism by taking it that all there is that can be appealed to in explaining the meaning of our utterances and the contents of our beliefs is the role those utterances and beliefs play in our habits and practices.²¹

This understanding of pragmatism, which I explicitly reject in Chapter 5, clearly corresponds to Brandom’s (2011) particular variety of semantic pragmatism. Although, similar to Price (2011), his view takes most of the metaphysical (i.e., representational) weight out of meaning, he still extends the “metaphysical scope” of his project beyond what Price (2011) wants to argue for. This is particularly in terms of how we ought to understanding the relationship of “representation.” Price argues for what he calls *internal* representations, or “i-representations.” He writes, “It is open to us to take the view that at least by the time we get to language, there is no useful external [representational] notion, *of a semantic kind*—in other words, no useful, general, notion of relations that words and sentences bear to the external world, that we might usefully identify with truth and reference.”²² Price (2011) argues that we can distinguish two “nodes” of representationalist talk. On the one hand, there is the idea that representations “co-vary” or are distinctly (and sometimes uniquely) related to something in the external world. This is the idea behind his sticker metaphor mentioned above. Another “node” of representationalist talk—the one

²⁰ To this, Brandom (2011) writes: “Being a pragmatic metavocabulary is the simplest species of the genus I want to introduce here. It is a *pragmatically mediated semantic relation* between vocabularies. It is pragmatically mediated by the practices-or-abilities that are *specified* by one of the vocabularies (which *say* what counts as *doing* that) and that *deploy* or are the *use* of the other vocabulary (what one says *by* doing that). The semantic relation that is established thereby between the two vocabularies is of a distinctive sort, quite different from, for instance, *definability, translatability, reducibility, and supervenience*” (*my italics*) (167-168).

²¹ Brandom (2011) p. 71.

²² Price (2011), *Naturalism Without Mirrors*, “Moving the Mirror Aside,” p. 21.

that Price (2011) prefers—“gives priority to the *internal* cognitive role of representation.”²³ He argues: “A token counts as a representation, in this sense, in virtue of its position, or role, in some sort of cognitive or inferential architecture—in virtue of its links, within a network, to other items of the same general kind.”²⁴ This line of argument is enthusiastically deflationist in that it asserts, as the semantic deflationist does, that the concept of representation as a language-world relationship is illusory. As Price (2011) notes:

On this view, the impression that there are such external relations will be regarded as a kind of trick of language—a misunderstanding of the nature of disquotational platitudes. But we can think of this without rejecting the internal notion: without thinking that there is no interesting sense in which mental and linguistic representation are to be characterized and identified in terms of their roles in network of various kinds.²⁵

Brandom (2011) argues that Price’s (2011) characterization of internal representations is interesting but lacking in substantive notion of *representation*. He suggests that Price (2011) uses i-representations as a “placeholder” for a more developed concept: “My reasons for saying that will emerge if we ask what makes the notion of I-representation a notion of a kind or sense of ‘representation.’”²⁶ Internal representations are such because they do not need to stand in any veridical, or “mapping or tracking relations to something outside the system,” they are fundamentally about networked relations within our own cognitive system.²⁷ But, as Brandom (2011) asks, why do we need to treat them as representations at all? At the very least, internal representations needs to begin with an explanation of what we track on to when we track what is expressed by our semantic practices. The representational relationship cannot be just the networks that exist amongst those representations (once we figure out what exactly they are). As Brandom (2011) argues:

Price likes the idea—at the core of my own thought—that a decisive line is crossed when we become entitled to think of the relations they stand in to one another as *inferential* relations. Indeed, I think we then become entitled to think of them (for the first time) as expressing *propositional* contents. For me, such contents are just what can play the role of premises and

²³ Price (2011), p. 20.

²⁴ *Ibid.*

²⁵ Price (2011), p. 21.

²⁶ Brandom (2011), p. 211.

²⁷ *Ibid.*

conclusions of inferences—what can both serve as and stand in need of *reasons*. But what results from that view is at least to begin with a notion of *I-expression*, not *I-representation*. For what does expressing propositional contents in *this* sense have to do with *representation*? Here it looks as though Price is seeking to procure by terminological fiat what can legitimately be secured only by honest toil.²⁸

Brandom's (2011) account of representation, unlike Price's (2011), attempts representationalism by matching (or "mapping") a discursive representational vocabulary to a pragmatic metavocabulary. It begins with Price's (2011) concern but gives a more concrete analysis to what we might mean by describing representations as internal by relating the contents of our propositions (via a representational *vocabulary*) with those conditions under which we can say that our use of language represents a successful example—or successful *use*—of "knowing *how*" to employ a given discursive representational vocabulary. I pause here to note that it is still unclear how either Price (2011) or Brandom (2011) understand the relationship internal representationalism has to science—even if that only entails a more detailed understanding of our cognitive apparatus. Given that, it makes it equally unclear how we should understand internal representations and/or the semantic practices that invoke them in the context argued for above. How does this suggest anything beyond a very narrow naturalism? I.e., how does this account for anything more than a naturalism that only covers the mental? Should we understand this as a "naturalist" theory at all except perhaps incidentally?

There is a problem with tying naturalism to a deflationist theory of truth – of which Brandom's (2011) semantic theory is a variety – without first trying to give meaning to the "science" part of naturalism. Although I'm sure that it is possible to incidentally or even accidentally advocate for naturalism, Brandom's (2011) view buries the explanation of how we "represent" the external world into an internally constructed, non-veridical system that is only tracked externally by our reliably produced semantic practices. And to add insult to injury, Brandom (2011) labels this as a kind of "pragmatism." But pragmatists do not unhinge the relationship between knowers and the "world-to-be-known" – particularly not in a way that demands a tracking relationship between our external and internal cognitive practices. But my grievance here is not with Brandom's (2011) questionable understanding of pragmatism, it is with

²⁸ Brandom (2011), p. 213-214.

how both he and Price (2011) seem to correctly diagnose a problem with understanding naturalism and then go on to commit seemingly similar errors. In their case, the problem of representing the external world simply becomes the problem of having some system or method of internally “tracking” our external linguistic practices. In any case, neither approach allows us to align (or even substantively *fail* to align) naturalism with a recognizable characterization of modern science. So as far as the “naturalism debates” are concerned, they both leave us hopelessly stymied. However, the distinction made by Price (2011) between object naturalism and subject naturalism is a very useful one – if only perhaps to show that subject naturalism needs to be anchored in some understanding of science (via method or content) to actually *be* a naturalist view.

2.2 Methodological Naturalism and The Pragmatist Conception of Science

We ought to co-opt “subject naturalism” and try to root it where I think it actually belongs – in a more robust understanding of the epistemological practices we invoke by using representational discourse. These epistemological practices may include linguistic practice, but the concept is by no means exhausted by linguistic practice. “Epistemological practices,” will include any practice that utilizes the scientific method, however broadly conceived. “Methodological naturalism” entails this commitment and it goes hand-in-hand with a broad conception of “science” itself. What it borrows from Price’s (2011) subject naturalism is that our regular discourse need not be so reductionist as to posit that the objects of our discourse be representational in a 1-to-1 relationship – or representational at all. Where it veers away from Price’s (2011) subject naturalism is in its denial of needing to track our discourse to internal representations and in its denial in insisting that linguistic practice is how we go about truth tracking. In this way, my interpretation is more in line with the pragmatist spirit of avoiding distinctions that necessitate a mind-world dualisms, although I’m sure Price (2011) would not like my sticking representations back into the external world even in this sense. Still, the appeal of Price’s (1997, 2011) subject naturalism is in its limited role for science and in its refusal to claim that science provides us a 1-to-1 description of the ontological objects invoked in regular discourse.²⁹ It aligns philosophical naturalism with science in a

²⁹ Price (1997) argues: “Object naturalism gives science not just center stage but the whole stage, taking scientific knowledge to be the only knowledge there is (at least in some sense). Subject naturalism suggests that science might properly take a more modest view of its own importance. It imagines a scientific discovery that science is not all there is—that science is just one thing among many that we do with ‘representational’ discourse.”

limited way that allows for the objects of science to be amended and/or updated as science progresses. It avoids the hard line of object naturalism and the ontological burden it places on what we might call our epistemologically justifiable linguistic practice. But a limited representational discourse, an abridged subject naturalism, can be thought of as doing something other than representing the objects of science in a strong sense. Rather, it can be seen as invoking the objects of science as the content of epistemological inquiry without having to commit to the existence of said objects. Even if we prioritize practice, we have to invoke objects merely by the default necessity of assertions being *about* some entity, *x*. Price's (1997, 2011) view won't really allow this latter claim and it is this assumption that I propose we reject. Instead, we ought to move the framework of discussion from the objects of science to the *methods* of science. And, even though our "science talk" will always invoke objects, we should make a distinction between the metaphysical objects we invoke by using a representational discourse and the successful epistemological practices that we implicitly or explicitly endorse while using making assertions. Accurately representing the objects of science via our linguistic practices should not be the foundation of a reasonable naturalism. A reasonable naturalism ought to entail that the objects we invoke in regular discourse can be at least potentially put to scrutiny via the most successful methods of scientific practice, not that they actually be confirmed as existing by science before we can use them in "representational-talk."

So what would the epistemological practices that we invoke by using a representational discourse actually be? For instance, if we are arguing about quantum loop gravity or stem-cell research, are we not invoking scientific ontology and not scientific practice? How we answer this seems to warrant a change of philosophical mindset as opposed to a substantive change in our world-view. And it is here that I see an appeal to pragmatism and pragmatic concerns as necessary. Consider this: if we were to adopt object naturalism then the object represented in the assertion, "Those are stem cells," would be the metaphysical object referenced in the statement itself. The stem cells. This case would seem easy enough, even if I were incorrect in identifying the object as stem cells. But consider an assertions such as, "Dark matter must be otherworldly." In this case, it is difficult, if not impossible, to map the referenced object to an existing scientific object. As Price (2011) would say, our "sticker book" method will fail. However, it will not fail for the same reasons that it might in the first case. In the second case, the "object"

is the result of a mathematical formula and an empirical conjecture. As such, there may not be an existent object at all. This is one variety of what Price (2011) refers to as the *material conception* of the “placement problem.” The problem arises for object naturalists as the result of assuming that existent objects create the semantic relations that dictate our linguistic behaviors. It assumes that if there is a *term X* there should equally be an *object X*.³⁰ But what if we were not concerned with the existence of dark matter in this case? Instead, what if we were concerned with the epistemological practices that justify the existence of dark matter instead? This is the pragmatic switch, it is the turn from questions about metaphysics to questions about epistemology. And it requires us to ask the follow-up question: how are we to understand the scientific method in relation to our epistemic practices? Here I appeal to an argument made by Rosenthal (2009) about pragmatic naturalism, where she asks:

Though the pragmatic focus on scientific method is a focus on pure method in the sense that it is a focus on scientific method as opposed to content, the method itself has far-ranging philosophical implications. What, then, does classical American pragmatism find when it examines scientific methodology by focusing on the lived experience of scientists rather than on the objectivities they put forth as their findings or the type of content that tends to occupy their interest; on the history of modern science rather than on its assertions; and on the formation of scientific meanings rather than on a formalized deductive model?³¹

Specifically Rosenthal (2009) seems to ask what the scientific method finds for pragmatic naturalists, or naturalists who’ve made the “pragmatic switch” described above, if not scientific “objects”? If the end result is that our epistemological practices merely represent the causal events and/or process that scientists go through to *confirm* objects, then “epistemological practices” is just Price’s (2011) “Big-N” naturalism via an alternate description. But the view is subtler than that, as we might make use of the pragmatist conception of science and truth to fill out our new naturalism.

³⁰ I have not delved into Price’s (2011) discussion of the “Priority Thesis” or why object naturalism and its resulting problems have to be validated first (and perhaps not at all) by subject naturalism. I think that his argument only works because of his initial characterization of object naturalism’s inability to capture human behaviors - and particularly human linguistic practices - which subject naturalism does capture. It is not entirely clear to me why object naturalism, and especially the “science” part of object naturalism must be characterized in this way. For that reason, most of this discussion has remained purposely absent from my use of Price’s (2011) distinction.

³¹ Rosenthal, Sandra. 2009. “Pragmatic Natures” in *The Future of Naturalism*, eds., John R. Shook and Paul Kurtz. New York: Humanity Books (Prometheus Books), pp. 77-96.

Consider Peirce (1877) – For Peirce, if we “fix” belief in the best way, we do so via the “method of science.” The other methods of inquiry - tenacity, authority, aprioricity – eventually lead us to a certain unease with our beliefs. It is the method of science that leads us from the “irritation of doubt” to a sustainable belief state. Thus we should model our general epistemic practices with the “method of science” which will ultimately align our beliefs with actual facts. However, these “actual facts” unlike the “objects of science” mentioned above, may or may not be sufficiently empirically verified. Allowing the method of science to dictate our epistemic practices will also allow our beliefs to converge toward a destined end or “external permanency.” As Peirce (1878) writes:

[The followers of science] may at first obtain different results, but, as each perfects his method and his processes, the results are found to move steadily together toward a destined centre. So with all scientific research. Different minds may set out with the most antagonistic views, but the progress of investigation carries them by a force outside of themselves to one and the same conclusion.³²

Peirce (1877) claims that this “external permanency,” or reality, is existent but only arrived at by a community of inquirers via the method of science. Thus our beliefs are constrained by, or fixed by, a reality that is still in the process of being determined. So reality may not fix belief in the short-run, as with adequately representing the objects of science for “Big-N” naturalists, but it will fix belief in the long run. If so, we can now say that our argument for thinking about subject naturalism in terms of epistemological practices (1) doesn’t beg the question against Price’s (2011) argument against representationalism by positing the same objects in alternate terms, and (2) doesn’t reduce the concept of “epistemological practices” as the foundation of an abridged subject naturalism to only linguistic practice, behavior, and analysis.

2.3 Social Epistemology and a More Honest Naturalism

This may be a more honest naturalism in that it does not make any explicit ontological assumptions about what already exists as confirmed by science or not. However, the ground that is gained by aligning naturalism with pragmatism and the “method of science,” is not freely gained. The price is a commitment to the pragmatist conception of truth and inquiry. And this may entail more than is

³² Talisse and Aikin (2011), p. 63.

initially apparent here. The method of science and the process of inquiry for the pragmatists rely heavily on the social dimensions of knowledge-production, or “communities of inquirers.” Truth seeking, then, cannot be an individualistic endeavor. And if that is the case, naturalists and naturalized epistemologists, cannot truly adhere to the foundational tenets of individualistic epistemology. Instead, all epistemology that sufficiently naturalized must be *social* epistemology.

In the next chapter, I give an overview of the most compelling arguments for social epistemology and suggest ways in which social epistemologically can respect its foundations in naturalism and be philosophically viable. I argue briefly in Chapter 3, and more thoroughly in Chapters 4 and 5, that the most viable form of social epistemology is veritistic social epistemology. Goldman (1999) is the originator of this veritistic view. However, as I hope to show, Goldman’s (1999) social epistemology, taken as an extrapolation from a Bayesian account of testimony, is terribly problematic. So, in Chapter 5, I offer a different version of social epistemology, a pragmatist-inspired social epistemology that has its roots in the version of pragmatic naturalism that I’ve suggested here.

Chapter III

Arguments for Social Epistemology

3.1 Introduction

More often than not it seems that social epistemology is either (1) defined as a question about how the “social” fits into traditional projects in epistemology¹ or (2) as a “correction” to the central projects of traditional epistemology.² In the former case, social epistemology becomes a project of figuring out what is meant by calling epistemology “social.” And, more specifically, it is an attempt to determine what would constitute a meaningful social epistemology in contrast with what we might consider traditional, individualistic epistemology. In the second case, the “corrective” project, we can divide the project into two distinct corrective (sub-) approaches. The first approach involves a complete revision of the goals of traditional epistemology. Among other things, this may include abandoning the epistemic goals of truth, objectivity, and normative criteria for rationality. This approach is what Goldman (2010) calls “revisionist” epistemology, with Richard Rorty (1979) as one of its most prominent defenders. The second corrective approach involves abridging the content or subject matter of traditional epistemology to include epistemological problems concerning groups, aggregates, institutions, and collective doxastic agents. The goals of this approach to social epistemology are still aligned with the goals of traditional epistemology; and, in particular, with traditional epistemology’s emphasis on truth, objectivity, and strict criteria for rationality.

Schmitt (1994) writes in favor of the first interpretation of social epistemology and that’s where we’ll start the discussion of the question: “What constitutes a *social* epistemology?” Schmitt (1994) claims: “Social epistemology is the conceptual and normative study of the relevance of social relations, roles, interests, and institutions to knowledge.”³ It “differs from the sociology of knowledge, which is an empirical study of the contingent social conditions or causes of knowledge or of what passes for knowledge in a society...”⁴ Additionally, he claims “social epistemology centers on the question whether

¹ See Schmitt (1994).

² See Rorty (1979), Kitcher (1993, 1994), Kornblith (1994), Goldman (1992, 1999, 2010), Longino (1990, 1993, 1994), Nelson (1993), and Solomon (1994, 2001).

³ Schmitt (1994), p. 1.

⁴ *Ibid.*

knowledge is to be understood individualistically or socially,” but that it is yet unclear where we might find an “informative” characterization of what it means for knowledge to be individual *or* social.⁵ He suggests that one way in which epistemology might be meaningfully social is by attributing beliefs to groups and showing how we might engage in the process of justification for group beliefs. However, he also claims that “those who go so far as to say that all knowledge—even the knowledge of individuals—is ultimately group or communal knowledge” have gone too far.⁶

Of course, this characterization of social epistemology is not unproblematic. Other accounts of social epistemology claim that social factors or social influences on individual cognition are crucial not just for traditional epistemology but for all epistemological projects in general (See: Kitcher (1994)). This would hint, in small part, at the claim concerning social epistemology’s being either individualistic or social as a bit inaccurate. Instead what we want to know is what role social factors or social influences play in our belief-making process; we are not interested in justifying a claim about whether or not they play a role at all. But why is that the case? Again, if one is skeptical about methodological naturalism then this claim about social epistemology will seem a bit less obvious. However, as Kornblith (1994) notes: “Knowledge is a socially mediated phenomenon. That social factors play a role in the production of belief is sometimes transparent, as in cases of deference to expert opinion. But the influence of social factors extends far beyond such cases, and pervades our entire corpus of beliefs.”⁷

Perhaps part of the reason for Schmitt’s (1994) claim that all knowledge cannot be group or communal knowledge is due in part to his hypothesis about finding the origins of social knowledge in the philosophy of testimony. Schmitt (1994) suggests that we can see the elevation of testimony as a “way of knowing” as a historical consequence of the growth of early-modern science “with its break from Aristotelian and common sense pictures of the physical world” and its “remapping of the distinction between knowledge and opinion.”⁸ These two notions, inspired by “new science,” required that, along with reason, perception had to play a role in knowledge production as well. That is, reason could not tell us everything that we sought to know about the world. For this we needed perception, thus empiricism. Schmitt (1994) attributes the birth of social epistemology to the rise in importance of testimonial beliefs in

⁵ *Ibid.*

⁶ Schmitt (1994), p. 257.

⁷ Kornblith (1994), p. 97.

⁸ Schmitt (1994), p. 2.

Hume's treatment of perceptual beliefs. For Hume, testimonial beliefs can be weakly justified on the basis of non-testimonial, perceptual beliefs. Testimonial beliefs, like much else, serve as evidence that will be either worthy or not worthy of consideration.⁹ So, although testimonial beliefs will have genuine epistemic status, it is at best a secondary status. According to Schmitt (1994), it was Thomas Reid who through his treatment of testimonial beliefs made social relations a primary source of knowledge rather than a secondary relation. Schmitt (1994) writes:

Reid (1969, 1975) also assimilated testimonial belief to perceptual belief, but he assimilated it to belief whose justification derives from a primary source rather than to a belief deriving from perception. In offering this alternative to Hume, Reid not only broke with the traditional assignment of a secondary status to beliefs indebted to social relations. He also made the status of social factors a topic of controversy. In doing so he founded social epistemology as a significant, self-conscious philosophical subject.

Schmitt (1994) argues that social epistemology as an area of inquiry has grown primarily out of an epistemology of testimony, but it is not a domain committed to the idea that the "social" is the correct sphere for all of our epistemological questions. Instead, it may be that there exist areas of the "social" for which we can attribute beliefs, e.g., groups and group beliefs, but that individualistic projects in epistemology are still the requisite norm. Specifically, this reinforces the traditional epistemological notion of the individual agent as the appropriate unit of epistemic analysis.

It's easy to see how Schmitt's (1994) view might be the result of the "origin story" that he chooses to tell about social epistemology. In this story, the basic unit of epistemic analysis (via testimony) is the individual agent and that agent's beliefs can only be supported or not supported via the input of an outside epistemic agent. This is because the beliefs themselves are still justified by the individual's believing p is warranted according to his or her own lights. C. A. J. Coady's (1973) arguments about proper justifications of testimony also support this line of argument.¹⁰ In "Testimony and Observation," Coady writes, "Hume is, indeed, one of the few philosophers I have read who had offered anything like a

⁹ See: Hume (1748, 2000) on miracles.

¹⁰ See: Schmitt (1994) where he mentions Coady's influence in terms of the history of social epistemology.

sustained account of testimony and if any view has a claim to the title of 'received view' it is his."¹¹

However, Coady (1973) finds the Humean account of testimony to be problematic because its justification rests up on a vicious circularity. In fact, it is the same vicious circularity that underpins Humean induction in general. As in the cases where we justify inductive hypotheses via their conformity to supposed (observable) expectations, testimony is also justified as such. Testimony is justified as observable phenomena of mankind; it does not have to be justified by any particular agent himself but only by *someone* at some past point, as with other, more general causal claims. It is important to note that it is not the act of testimony itself that is ever justified for Hume. Testimony, as Schmitt (1994) writes, has only a secondary epistemic status (Note: I think that it might be arguable whether testimony itself has any epistemic status at all. Rather, it seems that the epistemic significance of testimony is not in the "report" at all, but in the observation relationship that the "reporter" recounts). However, this emphasis seems a bit problematic. As Coady writes:

The idea of taking seriously someone else's observations, someone else's experience, already requires us to take their testimony (in this case, reports of what they observe) equally seriously. It is ludicrous to talk of their observations being the major part of our justification in taking their reports seriously when we have to take their reports seriously in order to know what their observations are.¹²

In any case, we can also say that Hume's account of testimony is problematic for far less complicated (and metaphysically-tinged) reasons. Like general Humean induction, our inferences about testimony hinge together because of what we infer about the uniformity of laws of nature (i.e., the "Uniformity of Nature" hypothesis) along with our (mankind's) general observations of testimony in the past. Specifically, this means that the foundation of our inferences rests upon a prior belief that (1) there are laws of nature that hold universally, and (2) these laws are what underwrite the inferences that we make—these laws of nature justify the conformity that we witness (e.g., all Fs are Gs). However, the hypothesis itself cannot be justified except by induction, which generates the vicious circularity that Coady (1973) also sees in Hume's account of testimony.

¹¹ Coady, C. A. J., "Testimony and Observation," in *American Philosophical Quarterly* 10 (1973), pp. 149-155.

¹² Coady, C. A. J. "Testimony and Observation" in *American Philosophical Quarterly*, Vol. 10, No. 2, (Apr., 1973), pp. 150.

Thus, as Coady (1973) claims, if we take testimony as a general phenomenon and subject it to scrutiny, we will find that it is also subject to the above-mentioned vicious circularity as well. This is because Hume views testimony as a common or universally justified phenomenon. We may, as Coady (1973) mentions, consider evaluating testimony as a particular phenomenon that can be justified by particular agents, but the problem with epistemic justification just crops up that much quicker. This is because it is equally unclear how any particular agent would justify individual testimony reports with a timeliness and consistency that matches how often we rely on testimony in general:

[I]t seems absurd to suggest that, individually, we have done anything like the amount of fieldwork that [this view] requires. As mentioned earlier, most of us have never seen a baby born nor have we examined the circulation of the blood nor the actual geography of the world nor any fair sample of the laws of the land nor have we made the observations that lie behind our knowledge that the lights in the sky are heavenly bodies immensely distant nor a vast number of other observations that [the view] would seem to require.¹³

What Schmitt (1994) takes from Hume's account and from Coady's analysis of testimony is that the social aspects of what might be called "traditional" epistemology enter into our questions about knowledge much earlier in the history of social epistemology than we might have previously thought. In some ways, Schmitt's (1994) account dovetails nicely with Goldman's (1999) view that it is testimony, rightly conceived, that will give us a robust account of social epistemology. However, their positions differ in that Schmitt (1994) lacks Goldman's commitment to the primacy of the "social," and his willingness to extrapolate social epistemological norms from a philosophy of testimony.¹⁴ Goldman's (1992, 1999, 2010) approach, like many other contemporary social epistemologists, rest on the assumption that social epistemology is a "corrective" project.¹⁵ However, it is a corrective project of the non-revisionist sort. It does not aim to supplant traditional epistemic criteria—truth, objectivity—with other norms of inquiry. Instead, it insists that traditional epistemology cuts its boundaries too narrowly. There are knowledge-

¹³ Coady, C. A. J., "Testimony and Observation," in *American Philosophical Quarterly* 10 (1973), pp. 151.

¹⁴ This is *not* to say that Schmitt (1994) does not give an account of how we can *argue* for the primacy of the social. This is only to say that he does not take it as given that all of our epistemological questions are, so to speak, "tainted" by it. See: Schmitt (1994): "The Justification of Group Belief."

¹⁵ Goldman (1992) claims: "Concentration on the individual to the exclusion of the social is inappropriate." Instead, he argues for a view called "social epistemics," where the aim of epistemological questions is not only normative and prescriptive, but dependent on all of the social, cultural, and political factors that influence the beliefs of individuals and communities.

producing “bodies” (e.g., juries, scientific laboratories, political bodies) that are left out of how we conceive of traditional epistemology. Thus what social epistemology *is* is something like “epistemology without the ego;” or, epistemology without an absolute reduction to the individual knower. This view is dominant among contemporary naturalists and naturalized epistemologists (see those listed above) because it is also aligned with the practice of scientific inquiry.¹⁶ Similar to Hume’s elevation of the epistemic status of perception, contemporary naturalists find that a crucial piece of evidence in science rests in the context in which science itself evolves—in the sphere of practicum. Thus the practical, or social arena of knowledge-production, ought to be another piece of evidence in our attempts to resolve epistemological problems.

3.2 Radical Social Epistemology

First, however, something needs to be said about social epistemology as a revisionist epistemology, or a radical attempt to undermine the main tenets of traditional (re: individualistic) epistemology. I will follow Kitcher (1994) in calling this collection of views, “radical social epistemology.” In this section, I will give an overview of the three-part characterization of radical social epistemology offered by Kitcher (1994). He argues that radical social epistemology as a set of philosophical arguments was born out of the denial of the claim that agreement, or consensus, in terms of a given hypothesis is distinct from what makes that particular hypothesis true.¹⁷ “Traditionalists,” Kitcher (1994) claims, “presuppose that the notion of truth is epistemically independent, that we are not to reduce the notion of truth in terms of what people know, or believe, or what the members of a society accept.”¹⁸ However, this is exactly what radical social epistemologists aim to do. The three-part breakdown can be described as the following:

The first campaigns for relativism by appealing to the fundamental tenets of the Strong Programme in the sociology of knowledge (most notably, the Symmetry Principle). The second invokes a venerable anti-realist attack on the correspondence theory of truth. The third rests on a thesis about the underdetermination of our claims about reality by our encounters with reality.¹⁹

¹⁶ See also: Putnam (1981), Longino (1990), and Kitcher (1992).

¹⁷ Kitcher (1994), p. 119.

¹⁸ Kitcher (1994), p. 119-120.

¹⁹ Kitcher (1994), p. 120.

According to Kitcher (1994): “The Strong Programme in the sociology of knowledge is based upon the attractive idea that *all* beliefs should submit to causal explanations, and that this causal explanation will involve social causes (See: Bloor (1991))...All suppose that ‘the same type of causes’ must be invoked to explain both true and false beliefs [*Note: This is the Symmetry Principle*].”²⁰ The radical nature of the Strong Programme depends heavily on how we interpret the Symmetry Principle. As Kitcher notes:

[I]f we are out to explain *X*’s belief that *p*, we shall surely do so by identifying *X*’s cognitive capacities, *X*’s interactions with reality, and *X*’s social background. Even in cases of perception, *X*’s socialization will be relevant, if only to help us understand why *X* forms a belief under the conditions that obtain and why *X*’s belief employs the categories that it does. Hence, *at a very general level*, the same types of causes will be invoked to understand any belief, irrespective of its truth value.²¹

What makes the Symmetry Principle problematic is how broadly or how narrowly we interpret the “type” of cause that the principle invokes. As Kitcher (1994) notes, if we interpret cause more narrowly while also making “traditional assumptions about truth” [i.e., non-relativism, the existence of a “truth-maker,” or a reliable truth-making processes, etc.] we will run into very counterintuitive results. Consider the following examples: “(a) perceiving macroscopic objects in good light; (b) forming conclusions about probabilities through the use of careful sampling...(c) forming beliefs by listing the first thirty-eight propositions that come into one’s head...(d) ingesting large quantities of alcohol...and forming beliefs about the number objects of various kinds that are present...”²². Now if we assume that the Symmetry Principle is true, then the same above-mentioned processes (a-d) can be used equally (i.e., symmetrically) to explain both true and false beliefs—whereas what “we should find [is] that some processes turn up more frequently in the explanation of true beliefs while others occur more frequently in the explanation of false beliefs.”²³

So what can save the Symmetry Principle from such unintuitive consequences? Kitcher (1994) claims that the absurdity surrounding the Symmetry Principle can be dispelled if a community already has in place a coherent notion of truth that identifies truth with something like “community-wide belief” or

²⁰ *Ibid.*

²¹ *Ibid.*

²² Kitcher (1994), p.121.

²³ *Ibid.*

community-wide acceptance. He suggests: "While we live in a culture that has institutionalized processes like (a) as truth promoting, others might just as well (and perhaps even do) take processes like (d) as truth promoting."²⁴

Furthermore, he claims: "[I]t is possible to hold that different types of processes can equally be made part of a community's standards for truth, so that, while the symmetry may be broken locally, the breach of symmetry can and must be explained in terms of the particular choices that the community has made."²⁵

We might now ask what type of social epistemology such a position about truth would generate? This type of reading of the Symmetry Principle could easily support a very radical social epistemology where truth is not only "non-traditional" and widely variant between communities, but also subject to absurdly quick changes in truth-conditions. This is particularly the case if there is no clear constraint on *what* kind of influence can alter a community's standards for truth, e.g., political violence, bribery, etc., or on how a community should go about determining such standards in the first place.

However, Kitcher (1994) claims that it is not the Symmetry Principle alone that pushes us toward very radical versions of social epistemology. The Symmetry Principle must first be given a much stronger reading where we individuate "type" more narrowly; and, as we've already seen, this is only problematic for those who adopt traditional approaches to truth. For those to want to abandon traditional approaches to truth, the relativism that the Symmetry Principle inspires is actually not a product of the principle itself, but of the prior belief that truth should already be viewed as non-traditional and highly relativized, e.g., reduced to something like community acceptance. Kitcher (1994) writes:

Appeal to the Symmetry Principle alone should not force us to adopt a version of social epistemology more radical than that discussed in section II...For there is no compelling motivation for adopting the Symmetry Principle in the strong sense in which types of processes are individuated narrowly. Defenders of nonepistemic approaches to truth will view such adoption as leading quickly to absurdity. However, for those who already believe that traditional, nonepistemic, notions of truth must be abandoned, the absurd consequences are artifacts of a misbegotten approach to truth, and the Symmetry Principle can be given a far stronger interpretation. The popular belief that acceptance of the Symmetry Principle thus leads to

²⁴ Kitcher (1994), p. 121.

²⁵ *Ibid.*

relativism seems to me to be quite mistaken. Only in the context of independent arguments against nonepistemic notions of truth does the principle obtain that kind of force. At best, the first line of argument can only reinforce conclusions that have already been reached on different grounds.²⁶

But what are we to make of the characterization of radical social epistemology that Kitcher (1994) offers in section II? We have already said how the view of radical social epistemology offered by Kitcher (1994) rejects traditional notions of truth as “epistemically independent,” or not reducible to common belief, consensus, social acceptance, etc., and this is due in large part to the three arguments already mentioned—(a) relativism as supported by the Strong Programme in the sociology of knowledge, (b) an anti-realist argument against the correspondence theory of truth, and (c) “a thesis about the underdetermination of our claims about reality by our encounters with reality” [note: I think that what Kitcher (1994) has in mind here is an “underdetermination of our claims about reality” via some empiricist thesis about reality, not via our direct interactions with reality. I will, however, come back to this claim later].²⁷ We can also say that, although he does ultimately reject radical readings of social epistemology, Kitcher (1994) does not reject the idea that we can successfully argue for a veridical social epistemology in general. In fact, Kitcher (1992, 1994) argues against the idea that epistemology can be committed to a thoroughgoing individualism full stop—or, rather, committed to what he calls the “reductionist program.”

Kitcher (1994) describes the “reductionist program” in traditional epistemology as “a set of propositions—the individualistic basis—that we can know without reliance on others.”²⁸ “Reliance,” however, means something rather weak in this interpretation. As Kitcher (1994) notes, “we are supposed to be able to assess the reliability of potential sources by checking their deliverances against propositions in the basis. Once a source’s reliability has been evaluated in this way, simple inductive inferences can lead us to employ that source in instances in which individualistic checking is impossible.”²⁹ The primary argument against the reductionist program is discussed in some detail in the Introduction to this project. However, I do want to reiterate that the main critique of the reductionist program can be boiled down to the following claim: our reliance on others in the reductionist view grossly understates the role that other

²⁶ Kitcher (1994), pp. 121-122.

²⁷ Kitcher (1994), p. 120.

²⁸ Kitcher (1994), p. 112.

²⁹ *Ibid.*

agents play in our development as knowers. First, this is because their reliability as knowers is crucial to justifying the claims that it is not possible for us to justify ourselves. This involves our recognizing their contribution as sources of knowledge as well as their contribution in the production of our own knowledge. Second, the reductionist view requires that we downplay just how often we are in the position of relying on others as sources of justification for the claims that we cannot ourselves justify, e.g., as in the case of “reductionist” views of testimony. Ultimately, as Kitcher (1994) seems to suggest, the amount to which others contribute to our knowledge—especially over the course of typical human development—would make the reductionist program seem wildly optimistic about individualistic claims to knowledge. That is unless we can reasonably assume some blank, assumption-free foundation for inquiry that does not rely upon any outside knowers as resources.³⁰ However, it is unlikely that this claim can be justified if the initial reductionist hypothesis cannot. In any case, the stronger rebuttal, although not entirely necessary to make the point here, might be better expressed as the following: “[T]here may be no propositions that we can know without being epistemically dependent on others.”³¹ Thus all epistemology must be at least minimally social.³² Although the question now becomes: what does it mean to be minimally social?

The idea that all epistemology is (at least minimally) social epistemology does not originate with Kitcher (1994) as I have attempted to show by the brief discussion of the early-modern philosophy of testimony. However, there have been more exhaustive articulations of this idea, one of which can be found in Fuller (1988, 2002).³³ However, the extent to which epistemology is social—even minimally so—

³⁰ Kitcher (1994), p. 112, “Unless we hold, as Descartes did, that there is some presuppositionless point from which we can begin inquiry, we must abandon the individualistic reduction as a failure.”

³¹ *Ibid.*

³² Kitcher (1994) claims: “Social epistemology begins at the point of rejecting the individualistic reduction...[T]he exact point at which epistemology becomes social is in the appreciation of the possibility that whether or not a subject is justified (or whether or not a belief-forming process counts as reliable in the pertinent sense) turns on the properties of other people or of the group to which the subject belongs...We can thus present the elements of a *minimal social epistemology* as follows: (1) Individuals are the primary subjects of knowledge. To ascribe knowledge to a community is to make an assertion about the epistemic states of members of the community. (2) *X* knows that *p* if and only if (a) *X* believes that *p* and (b) *p* and (c) *X*'s belief that *p* was formed by a reliable process. (3) The reliability of the process that produces *X*'s belief that *p* depends on the properties and actions of agents other than *X*,” (113).

³³ See: Fuller (1988, 2002): Fuller often refers to himself as one of the “founders” of contemporary social epistemology, a claim that I have no intention of trying to confirm or deny here. He is also considered, for a variety of reasons, to have a quite controversial take on epistemology in general. He argues for a different “origin story” for social epistemology rooted in the work of Auguste Comte and John Stuart Mill, and sketches little to no distinction between social epistemology and the sociology of knowledge. Although, I do not intend to dwell on this particular issue, I will say that the uniqueness of his work in the

requires that we eventually take another look at what role “truth” is to play in our epistemological view. This is because one of the primary consequences of ridding epistemology of both its foundationalist and reductionist impulses is also ridding it of traditional notions of truth and truth-making. Moreover, as many social epistemological projects often begin with, or are motivated by, claims about the status of truth; they are equally often motivated by more radical agendas in terms of truth and truth-making.

For instance, to generate a minimal social epistemology, Kitcher (1994) has to argue against the radical version of social epistemology; and, in particular, the radical version’s critique of traditional, nonepistemic truth. We have already seen how appealing to the Strong Programme in the sociology of knowledge, and the Symmetry Principle in particular, does not get us any closer to rejecting radical social epistemology. The Symmetry Principle alone does not justify the assumption of relativized (or consensus-based) views of truth; it only supports what the radical social epistemologist already believes. It also seems that radical social epistemology has the implausibility of the reductionist thesis on its side. If we cannot begin with any real Cartesian presumptions, we may have to start looking for the justification of our knowledge from the “outside-in” as opposed to the “inside-out.”³⁴

3.2.1 Against Radical Social Epistemology

The second prong of Kitcher’s (1994) argument comes in the form of an anti-realist argument against the correspondence theory of truth; and serves as the primary argument against radical versions of social epistemology. This is because those who defend radical social epistemology, and non-traditional notions of truth, specifically, seek to undermine the correspondence theory of truth as one of the most dominant interpretations of “truth-making.” One of the most serious critiques of the correspondence theory, and the one that Kitcher (1994) takes issue with here, is the following idea: a “correspondence relationship,” and thus the correspondence theory of truth, presupposes that we would be able to see (or verify) both sides of the alleged correspondence. That is, we would need a way to see

field of contemporary social epistemology cannot be understated; and, as such, it warrants some detailed explication here.

³⁴ Fuller (2002): The distinction between “inside-out” and “outside-in” justifications for knowledge comes from an illustration sketched by Fuller (1988, 2002) concerning two common “problems of knowledge.” The first problem asks: “How...do I determine whether other possible things exist, and, if they exist, how can I know them, given that they seem quite different from my own mind?” While the second problem asks: “What...accounts for [the] differences in our access to our common reality, and what enables us to ignore them in everyday life, as we suppose that our own access is the one shared by all (right-minded) people?” (xii).

both the conceptual scheme and the parts of reality that match up, so to speak, with the concepts that we employ in the scheme. If we do not have access to the reality of which the conceptual elements are supposed to correspond, then we have no reason to believe that the “correspondence” relationship holds at all, i.e., the correspondence relationship is false. As Kitcher (1994) writes:

Allegedly, to establish or to scrutinize the correspondence between thought/language and reality would require the attainment of some perspective from which both sides of the dichotomy could be viewed and the connections between them identified. Since there is no such out-of-theory position, no sense can be given to the correspondence or failure to correspond, and the notion of correspondence is senseless/incoherent/useless.³⁵

Kitcher (1994) rejects this argument as a failed “Inaccessibility of Reality Argument” (IRA). He dismisses the IRA argument for demanding an epistemic position from the realist (or the “real realist”) that is unmotivated by any realist concerns for symmetrical epistemic access: “There is no reason to claim that acceptance of the correspondence theory of truth requires some transcendent perspective, or that it presupposes some privileged position for the epistemic subject in which reality directly manifests itself.”³⁶ Although the radical social epistemologist might argue that claiming symmetrical epistemic access is unimportant is, at least, as equally unjustified as claiming the importance of symmetrical epistemic access. The line of argument that is offered by the realist in the case that Kitcher (1994) sketches is a scenario where what we are concerned about is how we can assess certain types of correspondence relationships from the outside. Consider the following example: “Imagine that you are observing the

³⁵ Kitcher (1994), p. 122.

³⁶ Kitcher (1994), p. 124: He also seems to suggest that the IRA argument also presupposes a naïve form of symmetry between the knower/observer and the world that we want to know/observe. One alternative to this is to suggest that knowers can stand in different—and importantly different—relationships to reality and thus can provide insight into the “limiting conditions,” so to speak, of knowledge-producing situations and dispositions. As Kitcher (1994) notes, “Standpoint Theory,” a subset of feminist epistemology and/or feminist postmodernist studies, provides some very good examples of this (See: Harding (1987, 1993), Haraway (1988), Longino (1990)). Although he does not claim that an appreciation for “standpoints” requires that our concerns about truth and truth-conditions for knowledge should be downgraded to concerns about standpoints full stop. Kitcher (1994) writes: “Some approaches in contemporary epistemology proceed far too swiftly from appreciation of the socio-historical situatedness of the knower to dismissal of the independence of what is known. Feminist epistemology offers important insights in its recognition that each of us occupies a standpoint, and that standpoints make epistemic differences. But, in light of my response to the IRA, I propose that the way to extend this insight is not to dismiss the ideal of objectivity, nor to reject the correspondence theory of truth, which supplies its most obvious underpinnings, but rather to probe systematically the ways in which different standpoints make available more or less epistemically apt dispositions, more or less reliable ways of generating true beliefs” (124).

behavior of another person and that you know not only what that person desires and intends but also how she represents the objects with which she interacts.”³⁷ Now how might you explain that person’s successful interactions with their environment? You might, for instance, “appeal to the accuracy of her representations,” meaning that, “she gets what she wants, to the extent that she does, because she represents the objects whose properties she controls, modifies, or compensates for in her own actions in ways that correspond to their actual dispositions.”³⁸ In this case, a big part of the story is that the correspondence relationship is not being assessed via symmetrical access (as the IRA argument assumes), but via the successful interactions of a subject, and a subject’s behaviors, with their respective environment. As Kitcher (1994) notes, “given the approach I have adopted, there is no reason to claim that acceptance of the correspondence theory of truth requires some transcendent perspective, or that it presupposes some privileged position for the epistemic subject in which reality directly manifests itself.”³⁹ As suggested in the scenario sketched above, we may make assumptions about reality by way of an agent’s interactions with their respective environment(s). Kitcher (1994) argues that we can have “perspectival” relativism, i.e., subjects standing in different reporting relationships to their respective environment(s), while still being realists about our ontological commitments. That is, we can have competing schemes of representation, relative to different cultures and different scientific communities, without abandoning realism.

However, Kitcher (1994) is not entirely convincing as to why this is the case. Part of the problem here comes from the anti-realist claim that because we have no way of deciding between differing conceptual schemes (or “schemes of representation”), we equally have no way of resolving issues presented by varying schemes appealing to differing ontological objects. Therefore, behavioral success is the only marker that our beliefs have been formed in the right way—i.e., it is the only marker of epistemic success. Proponents of this argument often rely here on the Duhem-Quine thesis to justify the central claim of the argument.⁴⁰ That is, proponents of IRA rely on the thesis to support the claim that our

³⁷ Kitcher (1994), p. 123.

³⁸ Kitcher (1994): He notes: “Reflection on such cases should bring home to us the importance of explaining the behavior of *others* by recognizing their representations, the correspondence of elements of those representations to objects that are independent of the individuals under study, and the connection between accurate (true) representations and success in dealing with those objects” (123).

³⁹ Kitcher (1994), p. 124.

⁴⁰ Kitcher (1994), p. 126.

inability to make a conclusive determination between competing conceptual schemes further reinforces the anti-realist claim that “reality” itself is under-determined.⁴¹ Furthermore, they argue that problems with observer biases, or rather, correcting for observer biases amongst varying conceptual schemes, cannot be resolved by appealing to broader notions of representational “success.”⁴²

However, Quine (1975) clarifies that there is a substantive difference between “under-determination” and “holism,” and that the Duhem-Quine thesis is really an argument for the latter.⁴³ Under-determination is the by-product of confirmation holism—which, at least in its strongest form, claims that all knowledge, as propositional knowledge, is under-confirmed or under-determined by the empirical data used to support any particular proposition. In any case, both under-determination and confirmation holism are claims about theory not claims about reality.⁴⁴ In fact, Quine’s (1977) thesis about ontological

⁴¹ See: Gillies, Donald (1998). “The Duhem Thesis and the Quine Thesis,” in Curd and Cover eds. *Philosophy of Science: The Central Issues* (New York: Norton), pp. 302-319.

⁴² Kitcher (1994) attempts to resolve some of the difficulties presented by dealing with rival “schemes of representation” by appealing to a broader notion of behavioral success via an agent’s given scheme of representation. That is, one might argue that different cultures not only utilize different schemes of representation, but that those schemes are not at all empirically equivalent, p. 127. Kitcher (1994) argues: “Now the issue of the potential relativity of success, to which we have been led, may itself look threatening. To understand the extent of possible trouble, let us consider the example of those non-Western societies that hold what we regard as radically false beliefs about the properties of certain taboo animals. It is quite possible that we should find, after detailed analysis, that these false beliefs play an important part in the way in which members of the society cope with reality. Suppose, specifically, that the false beliefs are invoked to defuse various types of social friction, and that, in consequence, these beliefs are partially responsible for the social order that members of the culture value. Thus we have an example of two societies in which incompatible claims are made about certain kinds of animals, both of which facilitate successful practices...But the important point is that what we view as the faulty beliefs about animals are not implicated in any way in the generation of successful interventions (whether predictive or manipulative) with respect to those animals. Hence it would appear possible to achieve a broader set of representations that would incorporate the Western biological views and the non-Western social understanding in a system that would preserve both sets of successes. Because there would be no internal inconsistency, there would be no challenge to the link between success and accuracy” (127-128).

⁴³ Quine (1975) on the Duhem-Quine thesis: “This doctrine of empirical under-determination is not to be confused with holism. It is holism that has rightly been called the Duhem thesis, and also, rather generously, the Duhem-Quine thesis. It says that scientific statements are not separately vulnerable to adverse observations, because it is only jointly as a theory that they imply their observable consequences. Any one of the statements can be adhered to in the face to adverse observations, by revising others of the statements. This holism thesis lends credence to the under-determination theses. If in the face of adverse observations we are free always to choose among various adequate modifications of our theory, then presumably all possible observations are insufficient to determine theory uniquely” (313).

⁴⁴ See: Chalmers, David J. (2009) “Ontological Anti-Realism,” in *Metametaphysics: New Essays on the Foundations of Ontology*, eds. David J. Chalmers, David Manley, Ryan Wasserman (New York: Oxford University Press). He writes: “Ontological realism is often traced to Quine (1948), who held that we can determine what exists by seeing which entities are endorsed by our best scientific theory of the world. In recent years, the practice of ontology has often presupposed an ever-stronger ontological realism, and

relativity is particularly influential in terms of the realism/anti-realism debates because he argues that all of our metaphysical claims are embedded in our best working theories. And, more importantly, that relying on the conceptual scheme to verify object-claims is a surefire way to end up making “Platonic objects” out of the conceptual scheme itself. (This debate is fully illustrated in Quine’s argument against Carnap’s use of conceptual frameworks. See: Carnap’s (1950) “Empiricism, Semantics, and Ontology”).⁴⁵ Kitcher (1994) rightly argues that the use of the Duhem-Quine thesis does not really apply to the argument(s) against IRA and, more specifically, IRA-supported radical social epistemology.⁴⁶ One problem is that the Duhem-Quine thesis seems to lack sufficient constraints on what will count as a viable, and/or “serious,” hypothesis.⁴⁷ As it stands, there are an infinite number of rival hypotheses for any hypothesis, *H*, under scientific consideration; however, as Kitcher (1993) notes: “Our prior practice recognizes certain kinds of processes as occurring in nature and not others...by abandoning that practice and adopting something different we could change the assessment of the serious possibilities.”⁴⁸ That is, our background practices (and, I would imagine, something like prior probabilities as they accord with an accepted background scientific method) lend credence to the hypotheses that we currently want to consider. Background practice will then suggest that some current hypotheses that do not accord with background practice will be eliminated from consideration. (Notice a similar circularity to the one that appears in

strong versions of ontological realism have received explicit statements by Fine (this volume), Sider (2001; this volume), van Inwagen (1998; this volume), and others. Ontological anti-realism is often traced to Carnap (1950), who held that there are many different ontological frameworks, holding that different sorts of entities exist, and that while some frameworks may be more useful than others for some purposes, there is no fact of the matter as to which framework is correct. In recent years, versions of ontological anti-realism have been developed by Putnam (1987), Sidelle (2001), Yablo (this volume), and others” (77-78).

⁴⁵ Carnap, Rudolf. (1950). “Empiricism, Semantics, and Ontology” in *Meaning and Necessity: A Study in Semantics and Modal Logic*. (Chicago: University of Chicago Press), pp. 205-221.

⁴⁶ Kitcher (1994) does hint that there is a substantive problem with using the Duhem-Quine thesis as a partial defense of IRA, but does not go into detail about *why* this is the case: “In the most obvious versions, Duhem and Quine are concerned with the possibility that incompatible sets of statements might prove observationally equivalent, either in the sense of yielding the same set of observational consequences or in the sense of accommodating the same set of stimulations of nerve endings...The argument from the Duhem-Quine thesis needs considerable development if it is to show that there are alternatives to commonsense ideas about the nutritive value of different substances *that would abandon or reverse our most basic ideas on this subject*, and that would be equally successful in keeping us alive. When we consider commonly cited examples of cultural variation, it is quite clear that the situations studied by Duhem and Quine are very different from those relevant to the assessment of the connection between practical success and representational accuracy” (126-127).

⁴⁷ See: Kitcher (1993). *The Advancement of Science: Science without Legend, Objectivity without Illusions*. (New York: Oxford University Press), pp. 247-248.

⁴⁸ Kitcher (1993), p. 248.

Hume's justification for the universality of nature is also at work here). Kitcher (1993) writes: "This point about scientific practice is readily accommodated in some accounts of science. Thus, Kuhn's conception of 'normal science' allows for determinate resolution of issues because of the constraining role of background practice (the "paradigm")."⁴⁹

So what remains of the IRA critique if conceptual scheme-based anti-realism and infinite rival hypotheses are not supported by appeal to Duhem-Quine? Perhaps the most damaging critique is the one that remains—there is, as of yet, no clear way to correct for observer bias in (a) creating conceptual schemes, and/or (b) validating the objects of conceptual schemes. What the radical social epistemologist may see that the traditional epistemology may not is that there are no benign observers and thus there are no benign objects of observation. So even if ontology isn't the primary issue, methodology still is. Specifically, the radical social epistemologist wants to examine the normative question about who is establishing the terms of debate in social epistemology. Kitcher (1994), wisely enough, leaves the issue for last. He notes:

I save until the end what has increasingly come, it seems to me, to be a major source of motivation—if not of argument—for a radical version of social epistemology founded in relativism. For those whose voices have traditionally been silenced, or ignored, an epistemology that seeks objective standards may appear inevitably oppressive, so that resistance to it is based more on concern than on the construction of detailed arguments...Mindful of the harm that has been done by treating some standards as objective, some conclusions as established, and some positions as superior, radical critics of traditional epistemology propose that we rethink our reigning metaphors. They envisage different ends for inquiry—not the control of nature grounded in the apprehension of truth, but the enhancement of human life through the sympathetic exploration of rival viewpoints and the development of attitudes of care and concern.⁵⁰

In Section 3.2.2, and for the remainder of this chapter, I will consider the question, "Can radical social epistemology be saved?" And, if so, how? I will begin by starting the discussion where Kitcher (1994) leaves off—there is a tension between the idea of treating the epistemic standards of truth and objectivity as the right ones, and leaving the question of "*who* decides?" unanswered. This is because the radical

⁴⁹ Kitcher (1993), p. 248, *footnote*.

⁵⁰ Kitcher (1994), pp. 128-129.

social epistemologist asks a question whose answer the traditional epistemologist often takes for granted, i.e., that truth and objectivity are the primary goals for any epistemology irrespective of the context of the knower. Although we can have primacy of truth and objectivity *and* have pluralism in terms of epistemic positions (or contexts), we first need to address this challenge put forth by radical social epistemology (RSE).⁵¹ We need to examine what motivates this tension in the first place.

3.2.2 Can RSE Be Saved?

Code (1993) argues that the absence of subjectivity in traditional epistemology comes from the view that there is something like a universal knower who is “capable of achieving a ‘view from nowhere’⁵² that allows them, through the autonomous exercise of their reason, to transcend particularity and contingency.”⁵³ She refers to this position as a “positivist-empiricist” position because it is “generated and enforced by appeals to such paradigms”; that is, paradigms that presuppose universally necessary and sufficient conditions for knowledge.⁵⁴ Code (1993) writes:

The ideals presuppose a universal, homogeneous, and essential ‘human nature’ that allows knowers to be substitutable for one another...Such epistemologies implicitly assert that if one cannot see ‘from nowhere’...if one cannot take up an epistemological position that mirrors the ‘original position’ of ‘the moral point of view’—then one cannot *know* anything at all. If one cannot transcend subjectivity and the particularities of its ‘location,’ then there is no knowledge worth analyzing.⁵⁵

Code (1993) further argues: “despite the disclaimers, hidden subjectivities produce these epistemologies and sustain their hegemony in a curiously circular process.”⁵⁶ That is, those who argue from what Code (1993) calls the “*S*-knows-that-*p*” positivist-empiricist position, assume that their position is context-free and/or context-independent and *isn’t* subjectively determined at all. Rather they assume that they have arrived at universally necessary and sufficient conditions for knowledge full stop. The presumption that

⁵¹ As Kitcher (1994) notes: “The challenge for the more conservative versions of social epistemology is to respond to genuine concerns about the oppressive force of standards without abandoning the benefits that the search for such standards makes possible” p. 129.

⁵² Code, Lorraine. (1993), *footnote*: “I allude here to the title of Thomas Nagel’s book, *A View From Nowhere*, (Oxford: Oxford University Press, 1986).

⁵³ Code (1993), p. 16.

⁵⁴ *Ibid.*

⁵⁵ *Ibid.*

⁵⁶ Code (1993), p. 19.

Code (1993) takes up here is that this model of knowledge is not only taken for granted—as opposed to being justified—but it assumes that knowledge, to be called “real” knowledge, must “transcend the particularities of experience to achieve objective purity and value neutrality.”⁵⁷ Because of this assumption, subjectivity is never addressed as a properly epistemic issue.

Code’s (1993) claim encapsulates one major challenge put forth by radical social epistemology (RSE) in the following sense: The radical social epistemologist wants to argue that “truth” and “objectivity” are the by-products of a subjectively determined epistemological position, i.e., positivist-empiricist position, and that the adoption of this epistemological position is merely taken as the right one by traditional epistemologists. So when Code (1993) asks the question “*Who says?*” we are forced to address this assumption made by traditional analytic (and naturalistic) epistemology. Code (1993) “denies the possibility of the disinterested and dislocated view from nowhere” and claims that most “knowledge-producing activities” are embedded in both a political context *and* an ethical framework—and this requires the epistemic subject to be responsible to their *epistemic communities*, not solely to empirical evidence.⁵⁸ Her argument rests on the assumption that “it is impossible to sustain the presumption of gender-neutrality that is central to standard epistemologies: the presumption that gender has nothing to do with knowledge, that the mind has no sex, that reason is alike in all men, and man ‘embraces’ woman.”⁵⁹

Instead Code (1993) suggests a “constructivist epistemology” or “constructivist reorientation” of epistemology that “requires epistemologists to take subjective factors, factors that pertain to the circumstances of the subject, *S*, centrally into account in evaluative and justificatory procedures.”⁶⁰ Code (1993) writes of such a project:

Because differing social positions generate variable constructions of reality and afford different perspectives on the world, the revisionary stages of this project will consist of case-by-case analyses of the knowledge produced in specific social positions. These analyses derive from a recognition that knowers are always *somewhere*—and at once limited and enabled by the specificities of their locations...my argument in this essay points to the conclusion that necessary

⁵⁷ Code (1993) pp. 18-19.

⁵⁸ Code (1993), p. 20.

⁵⁹ *Ibid.*

⁶⁰ Code (1993), pp. 20-21.

and sufficient conditions for establishing empirical knowledge claims cannot be found, at least where experimentally significant knowledge is at issue...Yet the relativism that my argument generates is sufficiently nuanced and sophisticated to escape the scorn—and the anxiety—that ‘relativism, after all’ usually occasions.

Code’s (1993) view aims to provide “a hybrid breed of relativism” where empirical (scientific) realism and epistemological relativism are not at odds. That is, we can have the radical elements of radical social epistemology without having to commit ourselves to anti-realism, which initially sounds like a very good thing! However, the argument for how we are to achieve this prescription is a lot less compelling than prescription itself. Code (1993) asks us to “take subjectivity into account” in way that does not equate our epistemology with radical relativism, and that accommodating “subjectivity” will not, in the end, commit us to an unwelcome variety of relativism. However, the “relativisms” at issue here are epistemological *and* metaphysical and it is unclear what makes the epistemological relativism that she offers of a tenable variety. I will come back to this issue in a moment. First, it is important to note what Code’s (1993) argument does get right.

The upshot of Code’s (1993) argument is that it does not push epistemological relativism into the same philosophical corner as most (RSE) arguments. What makes (RSE) unpalatable is not that it claims that subjectivity and the particularities of experience do not play a role in knowledge-production, but that it claims they are an *absolute hindrance* to truth and objectivity. The “radical” qualifier of radical social epistemology gets wedded tightly to a questionable metaphysics (see: Kitcher (1993, 1994)) rather than to a reimagining of epistemology—that is, it gets wedded to (very) strong anti-realist claims. This is (very) troubling for a naturalist, and I think it is for this reason that Code (1993) gives the following argument for at least a “minimally realist” metaphysics:

People need to be able to explain the world and to explain their circumstances as part of it; hence they need to assume its ‘reality’ in some minimal sense. The fact of the world’s intractability to intervention and wishful thinking is the strongest evidence of its independence from human knowers...People cannot hope to transform their circumstances and hence to realize

emancipatory goals if their explanations cannot at once account for the intractable dimensions of the world and engage appropriately with its patently malleable features.⁶¹

The epistemological subjectivity that she has in mind must have something to fall back on that is itself not subjective and/or relative. Otherwise, the want to “take subjectivity into account,” along with considerations of epistemic diversity and plurality, would lack substantive meaning in a world where our metaphysics is also subjective and/or relativistic. (Although, one might argue that this is the least of our problems with this view.) Those who argue for the form of (RSE) that denies the correspondence theory of truth and, in essence, objective truth-making, advocate for epistemic subjectivity because reality gives us no verifiable ontological objects to “correspond” to. Code (1993) suggests that the value of subjectivity isn’t at issue because our access to “reality” is flawed – but because the epistemological positions that we privilege in order to access reality are unjustifiable *as privileged* epistemic positions. That is, our methods of determining which epistemic position(s) are most appropriate in order to accurately talk about reality are flawed. Although I do not intend to take up the argument concerning which approach gives us the most accurate metaphysics, an account that does not require even a tentative commitment to anti-realism seems preferable—and Code’s (1993) argument for subjectivity has the least ontological baggage in that respect.⁶²

However, as mentioned, the epistemological relativism that Code (1993) argues for is not without other difficulties. Even if it were possible to entertain an infinite number of subjective positions in our epistemological analyses, we don’t have a systematic way to weed through these epistemological positions. This is crucial omission on Code’s (1991, 1993) part, as she argues against radical relativism *and* radical subjectivism. This means that, in her view, some epistemic positions will have to be more or less accurate, or will have to correspond more or less with objective facts (or what she calls “objective

⁶¹ *Ibid.*

⁶² See Code (1991), where she gives an extended argument for the claim(s) about realism, relativism, and subjectivity offered in Code (1993). Code (1991) makes the following claim about epistemic subjectivity and realism: “Facts may change and evolve in processes of interpretation and critique; hence ‘reality’ is indeed open to social structuring. Social practices, attitudes, institutions are far from constant, yet neither are they mere ephemera of a researcher’s imagination. They are there, present for analysis. Facts may mean different things to different people, affect some people profoundly and others not at all: hence they are both subjective and objective” (45-46).

social realities).”⁶³ To rectify this issue, Code (1991) seems to appeal to an epistemic ideal view, or “ecological model” of epistemology, as opposed to an epistemic method for sorting through the possibly infinite number of subjective views. She argues:

In arguing for a remapping of the epistemic terrain, I am taking issue with the implicit belief that epistemologists need only to understand propositional, observationally derived knowledge, and all the rest will follow. I am taking issue also with the concomitant claim that epistemologists need only to understand how such knowledge claims would be made and justified by autonomous, self-reliant reasoners, and they will understand all the rest. Beliefs of this sort are, in the end, politically oppressive in that they rest on exclusionary assumptions about the nature of cognitive agency and mask the experiences one might expect a theory of knowledge to explain...An ecological model builds on the mutual relations of organisms with one another and the relations between organisms and their environment...ecological thinking analyzes the implications, for organisms, of living in certain kinds of environments, and the possibilities, for those organisms, of developing strategies to create and sustain environments conducive to a mutual empowerment that is exploitative neither of the habitat nor of other inhabitants.⁶⁴

Code (1991) claims that the ecological model can be implemented by starting with “a model of personal relationships derived from friendship” where the notion of friendship can be modified or extended to “construct a regulative ideal for communal living.”⁶⁵ This model for communal living offers the most potential for any particular subjective view not only to be heard by the community, but to be realized by the individual him or herself.

One might have thought, that the aim of proposing an ideal model is to gain a more accurate understanding of “objective social realities” (e.g., our public and private epistemological spaces, such as our beliefs, opinions, popular attitudes, political spaces, governing institutions, etc.) via the inclusion of different, and often omitted, epistemological perspectives (e.g., women, individuals with disabilities, ethnic

⁶³ Code (1991) argues: “A contention that there are *no* objective social realities would obliterate the purpose of feminist political projects...She can, at most, construct a temporally, culturally, and geographically located analysis, according to the best evidence available. But such an acknowledgment does not deny that there are right and wrong descriptions of social realities or that evidence needs to be counted, objectively, as evidence” (45).

⁶⁴ Code (1991), p. 269-270.

⁶⁵ Code (1991), p. 278.

minorities, etc.). In this case, taking subjectivity into account matters because the traditional epistemological framework, with its “view from nowhere” stance, is too often a privileged position that is adopted without justification and without questioning the “whom” of the perspective-less knower. The inclusion of subjectivity and the contextual knowledge of a “situated knower” enrich our conception of objective reality, as opposed to narrowing it by tacitly enforcing an epistemic ideal. When Code (1991) claims that reality is objective but has a social dimension, we can take her as claiming, “Yes, there are ‘facts of the matter,’” but they are only accessible via situated knowers – and feigning otherwise does a disservice to knowledge and knowers alike. The problem is that is not what Code (1991) is arguing for in terms of the ecological model. The problem is that she assumes that an epistemology that “takes subjectivity into account” will support the “ecological model” that she prefers—and, so much so, that she specifies the unit of “friendship” as the necessary relation to bring this model about.⁶⁶ But adopting an ecological model involves prioritizing a set of values, e.g., friendship, mutual empowerment, a denouncement of oppression, etc., that our epistemology has yet to determine. That does not mean that those values may be incorrect determinations of our communal values, it only means that – according to her own problematization of objectivity – epistemic values ought to be neutral until a given community’s epistemological project is, so to speak, complete.

Code’s (1991, 1993) argument for subjectivity, like many others, ends up not respecting the process for knowledge-production that it actually advocates for. If we are knowers in the way that she argues, in that *what* we know, *how* we know it, and *to what extent* we are believed and trusted by others when making assertions is determined in large part by our association with other knowers, then knowledge, including knowledge of values, cannot be determined until the knowers within our given community are appreciated and included as such. This is the point of the subjectivity that she argues in the first place. Furthermore, the only way that Code (1991) can successfully argue for the priority of one model over another is if she denies that there are any objective facts, or any “objective social realities,” for which we require a variety of situated knowers—a collective “we”—to weigh-in upon in order to determine. And she explicitly denies that. As Code (1993) notes, quoting Fraser (1989), “The very goal of achieving

⁶⁶ Code (1991) writes: “It is clearly of paramount importance for community standards of responsible cognitive practice to prevail if members of epistemic communities are to be able to know well and to construct environments where they can live in enabling rather than oppressive circumstances” (270).

'as much intersubjective agreement as possible,' of extending 'the reference of 'us' as far as we can,' with the belief that tolerance will do the job when conflicts arise, is unlikely to convince members of groups who have never felt solidarity with the representers of the self-image of society."⁶⁷ The goal is not to construct the ideal epistemic model, but to offer a better epistemic method in order for all knowers to generate the model (i.e., the epistemological values) themselves. This way we can get the "situated knower," at all social and economic levels, at all levels of inclusion, involved in a new social epistemological project without falling back on a model that arises from epistemic privilege and the "view from nowhere." In many ways, Code (1991, 1993) commits the same mistake as those who claim that the "S-knows-that-p" model best represents appropriate epistemological values. To respect her own epistemological account, Code (1991) must allow for the unwanted scenario: an inclusionary epistemology, with situated knowers, which still generates severely individualistic, unequal, and oppressive spaces.⁶⁸ The larger implication here is that arguments that begin with the assumption that epistemic communities ought to determine epistemic values will have to live with the consequences of whatever those values turn out to be.

To get this analysis off the ground we do have to make some assumptions about the status of the epistemological approaches that are already culturally dominant (e.g., the positivist-empiricist position), but to assume that they are incorrect or misguided just begs the question against community-derived epistemic values. The upshot of a view that argues for subjectivity, and ultimately an *inclusionary epistemology*, is also one of its biggest liabilities. If we (re: Code (1991)) are wrong to articulate an epistemic ideal model for which all subjective positions should aim, then our epistemic values are at the mercy of the community itself, and this approach will not have "objectivity" as a constraint. This is what led Code (1991) to offer an epistemological model that she assumed was workable. However, and perhaps unfortunately, to take subjectivity seriously is to bar any presumptions about epistemological

⁶⁷ Code (1993), p. 24. It is important to note that I am not going to take up this particular case against Rorty (1979) here. This is because I think that this is far from the best argument against the type of anti-objectivist, anti-foundationalist epistemology that Rorty (1979) argues in favor of in *Philosophy and the Mirror of Nature* and elsewhere. Furthermore, I intend to follow-up on the arguments against consensus-based epistemology in Section 2.3, primarily with the intention to show why both Code's (1993) and Rorty's (1979) arguments are flawed.

⁶⁸ To be fair, Code (1991) does acknowledge the problem of "academic women, from their privileged place" speaking for other women, and other disadvantaged knowers (288). However, recognizing a problem is not the same thing as providing a resolution for it.

values – and this particularly disconcerting for epistemologists who are wedded to even a minimal conception of rationality. This is not a problem unique to Code (1991, 1993) but a problem inherent to both radical and consensus-based social epistemology in general. Consensus-based social epistemology theoretically succeeds by appealing to a multiplicity of epistemic values; but fails because it has no methodological apparatus to correct for observer biases and prejudices. In 3.3.3-3.3.5, I will address this critique of in further detail. Ultimately, as I hope to show in Chapters 4 and 5, the problem of the observer can be corrected by appealing to truth-conditions for all situated knowers in a way that still respects, but does not *endorse*, some key features of the consensus-based approach. Like Code (1991), I think there is a way to have our epistemological subjectivity and enjoy it without relinquishing too much of what underpins our most basic claims about reality.

3.3 Consensus-Based Social Epistemology

It is important to start by saying something about what a “consensus-based” epistemology looks like. I suggest the criteria sketched below as broad requirements for what should count as a consensus-based epistemology, although there may be some narrower requirements, as suggested by Goldman (1999).⁶⁹ To start, I suggest that consensus-based social epistemology (CSE) does not assume (a) either an arbitrary *or* a non-arbitrary selection of the individuals “who count,” epistemically-speaking, within a given community of inquirers, (b) or that any or all individuals will—or should—enter into a community of inquirers on equal footing, e.g., social, political, economic, educational, etc., or (c) or that “truth,” in any form, is an objective standard by which we can make non-arbitrary value distinctions between the consensuses reached in one community and another. For reasons that I will lay out in this section, as well as in Chapter 5, the third requirement, (c), will be the most difficult requirement for (CSE) arguments to negotiate. This requirement (c) is also what makes (CSE) arguments generally untenable, but perhaps still epistemically useful (I’ll come back to this later).

3.3.1 Background

For the most part, what motivates (CSE) arguments is very similar to what motivates (RSE) arguments: an assumption that an individualistic, foundationalist framework for knowledge claims is false.

⁶⁹ Goldman’s (1999) discussion of “rival conceptions of social epistemology” features an extended characterization of what he sees as “consensus” social epistemology, or “consensus consequentialism.” I will revisit Goldman’s claims about “consensus consequentialism” in Chapter 3.

What makes (CSE) arguments different from (RSE) arguments is that (RSE) usually advocates for “social constructivism” in the strong, anti-realist sense.⁷⁰ Consensus-based social epistemology does not generally wed a subjectivist epistemology to a radically relativistic metaphysics. This difference is what makes (CSE) arguments more palatable, philosophically speaking, but less so, methodologically speaking. The paradigmatic argument for (CSE) is probably the one offered by Rorty (1979), and so his grievances against traditional epistemology are a good place to start looking at what has, and did, motivate much of the subsequent (CSE) arguments, particularly amongst, neo-pragmatists, feminist social epistemologists, and feminist philosophers of science. For his part, Rorty (1979) claims: “The eventual demarcation of philosophy from science was made possible by the notion that philosophy’s core was ‘theory of knowledge,’ a theory distinct from the sciences because it was their *foundation*.”⁷¹ Rorty (1979) argues that Kant institutionalized the idea of “theory of knowledge” as the foundational, underlying structure of all inquiry, focusing particularly on the question, “How is our knowledge possible?”⁷² However, if it was Kant who gave us the task of philosophic inquiry, i.e., theory of knowledge, it was Descartes and Locke who gave us the conundrum of “mind” which in turn gave epistemology its content matter. Here Rorty (1979) writes: “Descartes’s invention of the mind—his coalescence of beliefs and sensations into Lockean ideas—gave philosophy new ground to stand on.”⁷³ This “new ground” was an area of philosophical inquiry that “seemed ‘prior’ to the subjects on which the ancient philosophers had had opinions...it provided a field within which certainty, as opposed to mere opinion, was possible.”⁷⁴

Rorty (1979) claims that it was Descartes who essentially changed the language of skepticism from its Pyrrhonian roots to a new “veil-of-ideas” skepticism that divided the external world from our knowledge of it. Rorty (1979) claims, “Traditional skepticism had been troubled principally by the ‘problem of the criterion’—the problem of validating procedures of inquiry while avoiding either circularity or dogmatism.”⁷⁵ Descartes’ “clear and distinct ideas” attempted to resolve this problem by establishing a criterion, but instead established one of the most problematic dualisms of modern philosophy. He made an “inner space” out of the “mind,” and a more mysterious “outer space” out of the external world. Rorty

⁷⁰ See: Goldman (1999) and Kitcher (1993).

⁷¹ Rorty (1979), p. 132.

⁷² *Ibid.*

⁷³ Rorty (1979), p. 136.

⁷⁴ *Ibid.*

⁷⁵ Rorty (1979), p. 139.

(1979) claims that “the idea of a ‘theory of knowledge’ grew up around this latter problem—the problem of knowing whether our inner representations [of this outer space] were accurate (*my notes*).”⁷⁶ However, he does not attribute the beginning of “traditional epistemology” to Descartes and Kant alone. Essential to this picture is Locke’s mistaken account of knowledge, one that vacillates problematically between “knowledge *of*” and “knowledge *that*,” or rather, between the conditions that make knowledge possible, physiologically, or mechanistically, speaking, and the conditions that justify a particular propositional statement.⁷⁷ Rorty (1979) writes:

Locke did not think of ‘knowledge that’ as the primary form of knowledge. He thought, as Aristotle had, of ‘knowledge of’ as prior to ‘knowledge that,’ and thus of knowledge as relation between persons and objects rather than persons and propositions. Given that picture, the notion of an examination of our ‘faculty of understanding’ makes sense, as does the notion that it is fitted to deal with some sorts of objects and not with others. It makes even more sense if one is convinced that this faculty is something like a wax tablet upon which objects make *impressions*, and if one thinks of ‘having an impression’ as in itself a *knowing* rather than a causal antecedent of knowing.

Rorty (1979) suggests, following Green (1968), that Kant attempted to make this fundamental confusion in Locke’s “theory of knowledge” right by necessarily wedding our perception of objects to “acts of mind.” That is: “Kant’s discovery was supposed to have been that there are no ‘qualified things’—no objects—prior to ‘the constitutive action of the mind.’”⁷⁸

Without digressing too completely into explication of Kantian epistemology, suffice to say that Kant began with the familiar Cartesian assumption that “inner space” and “outer space” were separated by a real explanatory gulf. Kant’s solution was to show that outer space, the external world, cannot be appropriately differentiated, as to be perceptible by the mind, unless it is done so *by the mind itself*. The

⁷⁶ Rorty (1979), p. 139-140.

⁷⁷ Here, Rorty (1979) appeals to an argument made by T. H. Green (1968) in *Hume and Locke*: “The fundamental confusion, on which all empirical psychology rests, between two essentially distinct questions—one metaphysical, What is the simplest element of knowledge? the other physiological, What are the conditions in the individual human organism in virtue of which it becomes a vehicle of knowledge?” (19). Green (1968) illustrates here the fundamental confusion of Locke’s epistemological account—the wedding of a mechanistic account of *how knowledge is physiologically possible* with an epistemic account of *what justifies a knowledge claim*.

⁷⁸ Rorty (1979), p. 147.

mind is fixed with *a priori* concepts, “categories,” which organize undifferentiated experience for us, and give experience its regularity and consistency. In Kant’s view, the mind is not a passive receptor of impressions but an active constituent in the synthesis of categories and intuitions (representations) into pure knowledge or empirical knowledge. So in one sense, Kant has corrected Locke’s mistake by not leaving the explanation of how knowledge is possible detached from how we justify knowledge claims; but, in another sense, Kant makes the same mistake as Locke. Kant’s theory of knowledge still makes “knowledge” an activity of mind that is distinct from outer experience, so much so that Kant postulates that we cannot even grasp experience, “outer appearances,” without some activity, or fundamental capability, of the mind, i.e., as is the case with our ability to represent objects in space and time.

We can see Kant as repeating the same error that Locke does in this respect, as well as introducing the new problem of setting up the mind as necessary to “constituting” or organizing the outer appearances of nature (or experience).⁷⁹ Rorty (1979) claims that Kant was right in thinking that the Lockean causal project was flawed, but wrong in thinking that the solution was to make knowledge transcendental. The solution, as he puts it above, taken from Sellars (1956), is to move from looking for the “foundations” of knowledge to justification of belief—and to focus the process of justifying our beliefs around rational discourse, and giving and revising reasons. He thinks of Sellars (1963) and Quine (1968) as wiping away our foundationalist tendencies in favor of what he calls “Epistemological Behaviorism,” or the view that empiricism tied to socially generated methods of verification are the only methods by which we can (or even ought) to justify belief.

3.3.2 Rortyan Pragmatism

The unwillingness to divorce philosophy (or what Rorty (1979) calls “philosophy-as-epistemology) from the search for an underlying “structure” is the legacy of Descartes,’ Locke’s, and Kant’s attempts to

⁷⁹ To this effect, Rorty (1979) claims: “Kant did not, however, free us from Locke’s confusion between justification and causal explanation, the basic confusion contained in the idea of a ‘theory of knowledge.’ For the notion that our freedom depends on an idealistic epistemology—that to see ourselves as ‘rising above mechanism’ we have to go transcendental and claim to have ‘constituted’ atoms and the void ourselves—is just Locke’s mistake all over again. It is to assume that the logical space of giving reasons—of justifying our utterances and our other actions—needs to stand in some special relationship to the logical space of causal explanation so as to insure either an accord between the two (Locke) or the inability of the one to interfere with the other (Kant). Kant was right in thinking accord was senseless and interference impossible, but wrong in thinking that establishing the latter point required the notion of the ‘constitution’ of nature by the knowing subject” (161).

“ground” our knowledge claims.⁸⁰ It is this particular thesis—the “Mirror of Nature” thesis—that drives the belief that the “way to have accurate representations is to find, within the Mirror, a special privileged class of representations so compelling that their accuracy cannot be doubted. [*And to assume*] these privileged foundations will be the foundations of knowledge (*my notes*).”⁸¹ But what happens to epistemology if the “Mirror of Nature” assumption is false?⁸² Rorty (1979) argues that the idea of “privileged representations” has been well refuted in modern philosophy by the work of Sellars (1963) and Quine (1968), and that this has made it possible to reimagine what, if anything, is left for “epistemology” to do.⁸³ If we untie ourselves from the idea of “philosophy-as-epistemology,” then what will take its place? The uncomfortable answer is: nothing. The hope of Rorty’s (1979) thesis is that we will move from epistemology to *hermeneutics*, and that “the cultural space left by the demise of epistemology will not be filled—that our culture should

⁸⁰ Here Rorty (1979) claims “traditional epistemology” rooted in Kantian epistemology is misguided as both an *extension* of his view or as a *repudiation* of his view. He writes: “the difference between the ‘mainstream’ Anglo-Saxon tradition and the ‘mainstream’ German tradition in twentieth-century philosophy is the expression of two opposed stances toward Kant. The tradition which goes back to Russell dismissed Kant’s problem about synthetic a priori truths as a misunderstanding of the nature of mathematics, and thus viewed epistemology as essentially a matter of updating Locke. In the course of this updating, epistemology was separated off from psychology by being viewed as a study of the evidential relations between basic and nonbasic propositions, and these relations were viewed as a matter of ‘logic’ rather than of empirical fact. In the German tradition, on the other hand, the defense of freedom and spirituality through the notion of ‘constitution’ was retained as the distinctive mission of the philosopher. Logical empiricism and, later, analytic philosophy were dismissed by most German (and many French) philosophers as not ‘transcendental,’ and therefore neither methodologically sound nor properly edifying. Even those with the gravest doubts about most Kantian doctrines never doubted that something like his ‘transcendental turn’ was essential. On the Anglo-Saxon side, the so-called linguistic turn was thought to do the job of demarcating philosophy from science, while freeing one of any vestiges of, or temptation to, ‘idealism’ (which was thought the besetting sin of philosophy on the Continent)” (161-162).

⁸¹ Rorty (1979), p. 163.

⁸² Or, as Rorty (1979) writes: “If knowledge is not a matter of accuracy of representations, in any but the most trivial and unproblematic sense, then we need no inner mirror, and there is thus no mystery concerning the relation of that mirror to our grosser parts” (126-127).

⁸³ Rorty (1979) also notes the attempts of 19th century philosophers, particularly Russell and Husserl, to reassert the need for foundationalism in philosophy. This turn-around, on the heels of the birth of “naturalized” philosophy, succeeded in bringing talk of “foundationalism” and “apodictic truth” back into philosophy. He writes: “Husserl saw philosophy as trapped between ‘naturalism’ and ‘historicism,’ neither of which offered the sort of ‘apodictic truths’ which Kant had assured philosophers were their birthright. Russell joined Husserl in denouncing the psychologism which had infected the philosophy of mathematics, and announced that logic was the essence of philosophy. Driven by the need to find something to be apodictic about, Russell discovered ‘logical form’ and Husserl discovered ‘essences,’ the ‘purely formal’ aspects of the world which remained when the nonformal had been ‘bracketed.’ The discovery of these privileged representations began once again a quest for seriousness, purity, and rigor, a quest which lasted for some forty years. But, in the end, heretical followers of Husserl (Sartre and Heidegger) and heretical followers of Russell (Sellars and Quine) raised the same sorts of questions about the possibility of apodictic truth which Hegel had raised about Kant” (166-167).

become one in which the demand for constraint and confrontation is no longer felt.”⁸⁴ The assumption of epistemology, which hermeneutics explicitly denies, is that conversation will yield commensurability. Where “commensurable” means something like “a set of rules which will tell us how rational agreement can be reached...These rules tell us how to construct an ideal situation, in which all residual disagreements will be seen as ‘noncognitive’ or merely verbal...capable of being resolved by doing something further.”⁸⁵

The fear of hermeneutics, according to Rorty (1979) is that if there is no commensurability then there can be no rationality. Conversations will take place, and perhaps agreements will be reached, but there will be nothing to ensure that said conversations, and said agreements, will not be irrational. To this, Rorty (1979) responds by reinterpreting the traditional epistemological use of “rationality” and replaces it with the following definition: “For hermeneutics, to be rational is to be willing to refrain from epistemology—from thinking that there is a special set of terms in which all contributions to the conversation should be put....”⁸⁶ It is helpful that Rorty (1979) thinks that the traditional notion of rationality, one that arises from thinking of epistemology as creating the proper epistemic path to accurate representations of reality, has been successfully refuted by the behavioral accounts of belief (“epistemological behaviorism”) offered by Sellars (1963) and Quine (1968). These antifoundationalist accounts make up an empirical position that gets at the heart of Rortyan Pragmatism as it clears the way for him to ask *and* affirmatively answer the question: “Can we treat the study of ‘the nature of human knowledge’ just as the study of certain ways in which human beings interact...?”⁸⁷ This answer can be specifically attributed to (1) Quine’s (1968) argument for indeterminacy in the translation of the natives’ reference of “*gavagai*,” and the pervasive inscrutability of reference,⁸⁸ and (2), to the “‘pragmatic view of truth’ as opposed to the views that lead to ‘ontological’ explanations of the relations between minds and

⁸⁴ Rorty (1979), p. 315.

⁸⁵ Rorty (1979), p. 316.

⁸⁶ Rorty (1979), p. 318.

⁸⁷ Rorty (1979), p. 175.

⁸⁸ As Quine (1968) famously notes, the inscrutability of reference begins “at home.” We have “no private language,” and no way to guarantee even our own referential objects (in language) outside of our background language. Rorty (1979) considers this argument, along with Sellars’ (1963) discussion of the social justification of assertions and denial of traditional Cartesian view of privileged access. For a good discussion of this, see: Loux and Solomon (1974), “Quine on the Inscrutability and Relativity of Reference,” in *Notre Dame Journal of Formal Logic*, Vol. XV, No. 1, January 1974, pp. 16-24, and Lehrer and Stern (2000), “The ‘Denouement’ of ‘Empiricism and the Philosophy of Mind,’” in *History of Philosophy Quarterly*, Vol. 17, No. 2, April 2000, pp. 201-216.

meanings, minds and immediate data of awareness, universals and particulars, thought and language...and so on.”⁸⁹ These arguments comprise two very important prongs of Rorty’s (1979) pragmatism.

It seems obvious what Rorty (1979) sees in Quine’s (1968) arguments for the necessity of a background language in which we conduct our inquiry, and by way of which we can sufficiently justify belief statements. The Quinean view makes social spaces, and social justification, a necessity. There is no vantage point from which we can “understand reality as it ‘really’ is,” or employ the “Mirror of Nature,” because our empirical statements are ontologically supported only by our background language (and, in some cases, by the “principle of charity”). Our background language makes referring to those entities shockingly indeterminate, in both the case of our pinpointing the references of others, as well as pinpointing references for ourselves. This may make it seem as if there is no referential “fact of the matter,” as Quine (1968) calls it, or no way of determining what is more or less accurate. But assuming this claim is true is only to revisit the assumption that there is some correspondence relationship between our statements about events and objects in the world and the events and objects themselves. This is just to invoke the “mirror.” Quine (1968) argues that to think that we can talk and/or refer to objects without some system in which to talk and/or refer to those objects is nonsense, “the inscrutability of reference is not the inscrutability of fact” because *there is no fact*.⁹⁰ As he writes:

Reference *is* nonsense except relative to a coordinate system...It is meaningless to ask whether, in general, our terms ‘rabbit’, ‘rabbit part’, ‘number’, etc., really refer to rabbits, rabbit parts, numbers, etc., rather than to some ingeniously permuted denotations. It is meaningless to ask this absolutely; we can meaningfully ask it only relative to some background language...Querying reference in any more absolute way would be like asking absolute position, or absolute velocity, rather than position or velocity relative to a given frame of reference. Also it is very much like

⁸⁹ Rorty (1979), p. 175. It is arguably the “pragmatic view of truth” which Rorty (1979) attributes primarily to James and Dewey that foreshadows the move from epistemology to hermeneutics. If Sellars (1956) and Quine (1968) cleared the way for empiricism, behaviorism, and antifoundationalism in epistemology, then the pragmatic view of truth (as Rorty (1979) interprets it) cleared the way for conversation, a culture of civility, and conformity to societal norms to replace traditional epistemology with its focus on rationality and commensurability. See Rorty (1979), p. 315-394.

⁹⁰ Quine (1968), p. 199.

asking whether our neighbor may not systematically see everything upside down, or in complementary color, forever undetectably.⁹¹

Rorty (1979) thinks of Quine's (1968) argument (as well as Dewey's and Wittgenstein's arguments against private language) as providing the groundwork for a new path in epistemological thinking that is inherently social, and it is from Quine's (1968) work that Rorty (1991, 1995) sees the arrival of Davidson's (1984) views on truth and language.⁹² But it is the "pragmatic" tendencies toward truth that Rorty (1979, 1991) sees in Davidson's work, or the pragmatist view of truth, as Rorty (1979, 1991) understands it, that inspires both his own pragmatism and his appreciation of Davidson's view of truth and language (1973, 1984).⁹³

Rorty's (1991) view mainly focuses on what he sees as Dewey's and Davidson's (1973, 1984) common deflationist view of truth, in conjunction with a repudiation of a correspondence theory of truth and any substantive metaphysics that may go along with it. As we already know from Rorty's (1979) appreciation of Quinean epistemological behaviorism, the emphasis of Rorty's pragmatism is on language and, eventually, on conversation. Rorty (1991) uses an abridged Jamesian view of truth to illustrate the primary theses for pragmatists:

- (1) 'True' has no explanatory uses.
- (2) We understand all there is to know about the relation of beliefs to the world when we understand their causal relations with the world; our knowledge of how to apply terms such as 'about' and 'true of' is fallout from a 'naturalistic' account of linguistic behavior.
- (3) There are no relations of 'being made true' which hold between beliefs and the world.
- (4) There is no point to debates between realism and anti-realism, for such debates presuppose the empty and misleading idea of beliefs 'being made true.'⁹⁴

⁹¹ Quine (1968), pp. 200-201.

⁹² Rorty (1991), "Pragmatism, Davidson and truth" in *Objectivity, Relativism, and Truth, Philosophical Papers, Vol. 1*, Cambridge: Cambridge University Press, p. 126. See also: Rorty (1979), p. 165-212, and 257-311.

⁹³ Rorty (1991) notes: "Davidson, however, has explicitly denied that his break with the empiricist tradition makes him a pragmatist. He thinks of pragmatism as an identification of truth with assertibility, or with assertibility under ideal conditions. If such an identification is essential to pragmatism, then indeed Davidson is as anti-pragmatist as he is anti-empiricist" (126).

⁹⁴ Rorty (1991), p. 128. He notes of thesis option (2), "This thesis does not, of course, entail that you can define intentional terms in non-intentional terms, nor that a semantical metalanguage can somehow be 'reduced' to Behaviorese." Additionally, he notes of thesis option (4), "Jamesian pragmatists heartily

Rorty (1991) claims that this is actually *not* a “theory of truth,” as the classical pragmatists, aside from Peirce, did not offer one in the traditional sense.⁹⁵ To be a pragmatist, as Rorty (1991) understands it, is to deny the “traditional problematic about truth” and start with the assumption that “‘true’ has no explanatory uses.”⁹⁶ Instead, we can use “true” to denote a certain quality in linguistic behavior, or to denote a certain quality of a given utterance. Consider the given descriptions offered by Rorty (1991). He claims that the term, “true,” can have one of the following uses:

- (a) An endorsing use;
- (b) A cautionary use, in such remarks as ‘Your belief that S is perfectly justified, but perhaps not true’—reminding ourselves that justification is relative to, and no better than, the beliefs cited as grounds for S, and that such justification is no guarantee that things will go well if we take S as a ‘rule for action’ (Peirce’s definition of belief);
- (c) A disquotational use: to say metalinguistic things of the form ‘S’ is true iff ____’.⁹⁷

The notion of truth that Rorty (1995) employs is primary deflationist in that it emphasizes “justificatory practices” instead of a notion of “truth” that appeals to the “mirror.” Rorty denies that there is a substantive difference between being “justified in believing p ,” and “accepting p as true.” Additionally, he denies “that there is anything substantive to be said about truth...once we have finished talking about justification.”⁹⁸ We may invoke what he calls a “cautionary” use of “true,” such that what we really mean is that p is extremely justified (or highly probable) given the inquiry conducted by our epistemic community thus far. (We may also invoke the disquotational use of “true,” or even others if necessary for our aims).⁹⁹ But there is no weightier, more “realistic” notion of truth available to us. It is better to think of arriving at

agree with Dummett’s claim that lots and lots of traditional ‘problems of philosophy’...are best seen as issues between realists and anti-realists over whether there are ‘matters of fact’ in, e.g., physics, ethics, or logic. But whereas Dummett sees himself as having rehabilitated these fine old problems by semanticizing them, the pragmatist sees him as having conveniently bagged them for disposal” (128, *footnote*).

⁹⁵ I will argue in Chapter 4 that this assumption by Rorty (1979, 1991) is actually a mistake, and that the classical pragmatist not only had a robust “theory of truth,” but that James and Dewey both aligned themselves with the Peircean view of truth. Their differences, should they be relevant, in fact stem from James’ focus on religion and psychology and Peirce’s dislike of nominalism. It was Dewey who defended their views as both consistent *and* related from pragmatist critics, e.g., Russell.

⁹⁶ Rorty (1991), p. 127.

⁹⁷ Rorty (1991), pp. 127-128.

⁹⁸ Rorty (1995), p. 149.

⁹⁹ Rorty (1995), p. 150.

being “justified in believing p ” via a conversational process that is subject to conversational norms, and ridding ourselves of “truth-talk” entirely. Rorty (1979) suggests that this idea comes from the classical pragmatist view that truth is whatever our fellow inquirers will “let [us] get away with saying.” That is, whatever we can “get away with saying” according to the current norms of conversation and justification that happen to be in place at the time.

The last prong of Rortyan pragmatism comes from Kuhn’s (1970, 1977) arguments concerning objectivity in science. Rorty (1979) uses Kuhn’s (1977) distinction between “normal” and “revolutionary” science to mirror the distinction between “traditional” epistemology and hermeneutics. Instead of looking at science as the paradigmatic example of knowledge, Rorty (1979) argues that Kuhn (1970) takes science to be closer to “ordinary conversation.”¹⁰⁰ This is because Kuhn (1962, 1977) denies that there is commensurability “between groups of scientists who have different paradigms of a successful explanation;” and thus makes it imperative that scientists engage in hermeneutics at some point.¹⁰¹ The working scientist cannot appeal to any definitive, truth-directed notion of theory-choice in order to achieve commensurability, “the criteria of choice between theories,” or what turns out to be the *justification* for theory-choice, is determined in large part by the values shared by a community of scientists (See: Kuhn (1977) where he specifically isolates and explains some of these shared values). For Rorty (1979), the ultimate takeaway from a Kuhnian analysis of science is that it further suggests that the “Mirror” or “correspondence to reality” view of knowledge is flawed, as Kuhn (1977) refuses to give meaning to the terms “truth” and/or “objectivity” as though they referred to some entity or point at which science will arrive. He warns philosophers of science with the tendency to regard scientific “truth” with the end point of inquiry that science *in practice* simply does not require, or necessitate, such a demand. “Normal science” will proceed—not until it arrives at “truth”—but until a “revolutionary” period begins. If this view of science is accurate, then the traditional epistemological hope for science as the guarantor of truth and knowledge, in a naturalistic but still heavily metaphysical sense, will not pan out. We will be forced to take up other scientific values as our markers for successful science. Given this, it is easy to see why Rorty (1979) supports this particular view of science.

3.3.3 Feminist Philosophy of Science and Feminist Epistemologies

¹⁰⁰ Rorty (1979), p. 322. Original reference: Kuhn (1962).

¹⁰¹ *Ibid.*

As with Rorty (1979), there is a sub-section of feminist philosophy of science and feminist epistemology also inspired by Kuhn's (1962, 1977) work on science and objectivity. These views generally characterize consensus-based social epistemology (CSE) due to their emphasis on "situated knowledge" and "standpoints," and general skepticism toward traditional epistemic notions of scientific truth and objectivity. For this reason, this section will conclude with a representative overview and discussion of these views. I focus on Solomon's (2001) discussion of Longino's "standpoint epistemology of science" as the lens for discussion, but suffice to say that there is an unavoidable arbitrariness to the selection of work that I'm considering here. This is only because I don't have the space to address this area of philosophic thought as thoroughly as I would like. (That will have to be a project for another time, and I hope that any glaring omissions in this section will be forgiven as such.) The primary goal here is to illustrate that these philosophical arguments, arguments that employ "standpoints," like all (CSE) arguments, suffer from the same problems inherent to radical social epistemology, (RSE). In previous sections (See: 3.2.1 and 3.2.2), I suggested that Kitcher (1994) was right to address the tension between treating truth and objectivity as the *only* rightful epistemic ideals without asking the question, "Who decides?" I offered Code's (1993) argument in favor of abandoning the positivist-empiricist position—the dislocated "view from nowhere"—and incorporating subjectivity into our epistemological accounts. I also offered what I believe to be a reasonable critique of Code's (1991, 1993) arguments for subjectivity on the grounds that her approach generally begs the question against the positivist-empiricist view by offering another ideal epistemological model, the "ecological model." This particular problem of subjective views advocating ideal models is not unique to Code (1991, 1993). The impetus to offer ideal models comes from the belief that in order for the community to determine epistemic values without non-community generated epistemic constraints is, at least in part, to surrender to epistemic irrationality. However, attempts have been made in consensus-based epistemology, (CSE), to wed the subjective approach preferred by (RSE) views with epistemic constraints that may prevent radical subjectivity and/or irrationality. These views fall under the umbrella of "consensus-based social epistemology," with feminist philosophers of science and feminist epistemologists as its most prominent defenders.

3.3.4 Solomon (2001) and Standpoint Epistemologies of Science

In *Social Empiricism*, Solomon (2001) offers Longino's (1990) epistemology of science as one of the best and most "normatively sophisticated of feminist standpoint epistemology." She argues that other feminist standpoint epistemological accounts (e.g., Harding, Haraway, Keller, Nelson) do not "go far enough" in their analyses of biases and prejudices, or "decision vectors" as she (2001) calls them, in terms of theory choice and hypothesis construction.¹⁰² Still, however, she claims that standpoint epistemologies of science do offer one of the most significant challenges to traditional "epistemology of science," (Note: What I would still want to call "philosophy of science" proper) in that they do not presuppose the "view from nowhere" thesis—or, rather, that there is some "neutral perspective from which to identify and eliminate 'bias.'"¹⁰³ Solomon (2001) makes the following claim:

Few epistemologies of science devote attention to the identification of decision vectors (or 'bias') or to realistic ways for improving scientific decisions (or 'addressing bias'). A major exception is standpoint epistemologies of science. It is for this reason that I view them as the deepest challenge to traditional (traditional naturalistic or analytic) epistemology of science.¹⁰⁴

Solomon (2001) contrasts "social empiricism" as a "standpoint approach" from both "feminist standpoint epistemology" and Antony's (1993, 1995) "feminist empiricism."¹⁰⁵ She characterizes feminist empiricism as a philosophical rejection of the standpoint approach, primarily on the grounds that it leads to a paradoxical understanding of utilizing "standpoints" in the first place: "Antony coined the term 'bias paradox' to characterize the conundrum facing feminist epistemologists: if the goal is truth, and all positions are biased, and bias reduces the chances of getting truth, why regard feminist epistemologists

¹⁰² Solomon (2001) offers a compelling analysis of how empirical and non-empirical decision vectors factor in to a "social empiricist" view of scientific practice. It's probably fair to say that her account of empirical and non-empirical decision vectors is not uncontentious, and this does impact the claim here that some feminist standpoint epistemologists have not carried the idea of biases, social context, and social situatedness far enough in their own accounts. I mention this only to say that it is an important consideration in terms her understanding of the normative significance of Longino's (1990, 1994, 2002) work. However, for my purposes here, I will leave this assessment unchallenged.

¹⁰³ Solomon (2001), p. 141.

¹⁰⁴ Solomon, Miriam. *Social Empiricism*. Cambridge: Massachusetts. MIT Press (2001), p. 141.

¹⁰⁵ It may be important to note here that Solomon (2001) does not see "social empiricism" as a version of *feminist* standpoint epistemology, but as a version of standpoint epistemology in general. It is not terribly clear what's making the substantive difference here if not to say that there is something about *feminist standpoint epistemology* that—like the above footnote mentions—is problematic due to how feminist standpoint epistemologists traditionally understand and incorporate "decision vectors." However, the contrast between "social empiricism" as such and Louise Antony's version of "feminist empiricism" is much more philosophically substantive. See Solomon (2001), pp. 138-151.

as any more credible than the gender-biased scientists they are criticizing?”¹⁰⁶ Solomon (2001) objects to Antony’s characterization on several grounds, however, I think the most important of her objections is that it is still not clear that “bias” is an “epistemic flaw:” Instead, she claims through *Social Empiricism* that “decision vectors of various kinds can often contribute to scientific success and truth.”¹⁰⁷ This is one of the reasons that she picks up on Longino’s (1990) work as an example of where feminist standpoint epistemology (and feminist philosophy of science, more generally) goes right. Solomon (2001) writes:

Longino recognizes that traditional canons of scientific method...are not sufficient for doing science. Scientific work is underdetermined by those canons. According to Longino, ‘values,’ which come from ideologies, fill the gap. Different values lead different scientists to develop and pursue different theories. The result can be a healthy pluralism...She argues that ideological factors benefit science when the scientific community is structured so as to be, in her sense, ‘objective.’ This requires equality of intellectual authority, public forums for criticism, responsiveness to criticism and shared standards including the standard of empirical adequacy. Longino expects that ‘individual biases’ such as cognitive bias, error, motivational bias, etc., will be eliminated through social criticism and that deeper values, or ‘ideological bias,’ while not eliminated...will be made manifest and treated more democratically.¹⁰⁸

What is nice about Longino’s (1990) view (see also Longino (2002)), and what Solomon (2001) picks up on, is that it offers a way to understand and incorporate bias without the “forced idealization” problem that results from Code’s (1993) view. (See also Anderson (2004)).¹⁰⁹ That is, in Longino’s (1990) view, we can have “objectivity” in some sense, if not wholesale. In her view, the objectivity that is at play is a *social* objectivity that results from individuals working together in a variety of social processes, e.g., public forums, public debates, etc., in conjunction with a set of adopted institutional norms, e.g., assumed intellectual equality, shared standards for empirical adequacy, etc., to achieve scientific objectivity. Longino’s (1990, 2002) view assumes that we can, by way of shared social processes and norms, wipe away individual biases and prejudices (not ideological prejudices) and “transform the subjective into the

¹⁰⁶ Solomon (2001), p. 142.

¹⁰⁷ *Ibid.*

¹⁰⁸ Solomon (2001), p. 143.

¹⁰⁹ Anderson, Elizabeth. “Uses of Value Judgments in Science: A General Argument, with Lessons from a Case Study of Feminist Research on Divorce” in *Hypatia*, Vol. 19, No. 1, Feminist Science Studies (Winter, 2004), pp. 1-24.

objective” by way of “critical discursive interactions.”¹¹⁰ That is, in Bayesian-speak, she offers a way for us to minimize prior probabilities. However, as Solomon (2001) rightly mentions, Longino’s (1990, 2002) view may be more “intuitively plausible” than naturalistically justifiable. She does not appeal to cases where we can see these types of social processes and norms at work in achieving (something like) scientific objectivity. Additionally, as Solomon (2001) notes, it is not even clear that such adoptions will result in objectivity as opposed to deeper individual entrenchment in one’s own prior views. She writes: “There is already reason for doubt...criticism frequently results in entrenchment of prior positions (due to pride, confirmation bias, etc.)...And, for example, no one has a realistic model for a scientific community in which inequalities in intellectual authority do not play a large role.”¹¹¹

Solomon’s (2001) critique hits at the heart of the problem for this account of scientific objectivity—it is too idealized. And, for a naturalistic epistemology, this is no small problem. Additionally, and perhaps more pessimistically, Longino’s (1990, 2002) view also does not deal with the issue of epistemic *taint*, unintentional or otherwise. Scientific objectivity in her view is at the mercy of the members of the scientific community to adopt and/or adhere to the rules of the social practices and social norms that have been agreed upon. A cynic might even take a moment to consider that Longino’s (1990, 2002) account mirrors many of the positive assumptions that we already have about the way in which scientific practice ought to conduct itself, i.e., open-minded, responsive to criticism, not reliant on intellectual authority, etc., and that this has not been terribly successful so far in weeding out individual or ideological biases.

3.3.5 Against Consensus-Based Social Epistemology (CSE)

So where does all of this leave consensus-based social epistemology? Of which feminist standpoint epistemology and feminist philosophy of science are often a part? In some of Kuhn’s follow-up work to the *Structure of Scientific Revolutions*, he argues that the philosopher’s demand for truth and objectivity in scientific practice is a misguided notion if it is not understood as part of the scientific paradigm itself. There is no weightier or more metaphysically or epistemologically *accurate* version of

¹¹⁰ Longino (2002) argues: “Critical discursive interactions are social processes of knowledge production. They determine what gets to remain in the public pool of information that counts as knowledge. Thus, a normative account of knowledge must rest on norms governing such interactions. Criticism must be epistemologically effective—by helping a community avoid falsehood and by helping to bring its accepted content into alignment with its cognitive goals and its cognitive standards. Effective critical interactions *transform the subjective into the objective*, not by canonizing one subjectivity over others, but by assuring that what is ratified as knowledge has survived criticism from multiple points of view” (129) [*my italics*].

¹¹¹ Solomon (2001), p. 144.

“truth” and/or “objectivity” to be had. Although we may hope that there is some *thing* that science is progressing *to*, Kuhn (1977) vehemently argues that there is not. In this case, consensus within a scientific paradigm, along with efficiency, creativity, utility, etc., will be our closest approximation to what we might call “truth.” There is, of course, more to say about the Kuhnian impact on consensus-based social epistemology. However, for the moment, I will focus on this feature—the role of “truth.” For (CSE) views, “truth,” in any form, cannot be an objective standard by which we can make non-arbitrary value distinctions between the consensuses reached in one community versus those reached in another. This is a lasting remnant of Kuhn’s (1962, 1977) work on the practice of science, and, in light of this, it makes sense as to why Rorty (1979) and others have understood both science and epistemology as a world without capital “T” truth. “Truth” is what we can get inside of our scientific paradigm not something towards which all of our scientific inquiries ultimately progress. This is a compelling story. However, it isn’t the only one or even the correct one. I do not plan to argue against (CSE) as much as I hope to merely “problematize” this view of truth in science (and in inquiry generally) and offer a reasonable alternative.

Consensus-based social epistemology, while susceptible to what Solomon (2001) would call “naturalistic investigation” or naturalistic justification, fails in terms of epistemological justification. (Note: There is cause to separate naturalistic from epistemological justification here, where “naturalistic” in Solomon’s (2001) sense is closer to something like “broadly empirical,” or “determined by cases.” However, this is not always an appropriate separation.) I *am* sensitive to Kuhn’s (1977) views about “truth” in science. I am also sensitive to (CSE) views that argue for consensus in conjunction with agreed upon social practices and norms, e.g., Longino (1990, 2002). But the underlying claim—that such consensus will generate the closest approximation to what we mean when we talk about “truth”—is hopelessly unsatisfactory. Goldman (1992) argues that this type of “relativism” in social epistemology cannot avoid irrationality—that is, if rationality is a worthwhile epistemic goal—because consensus is too easily tainted. He writes: “Consensus *per se* is not a reliable sign of rationality. It depends on how consensus is reached. All sorts of methods can yield consensus: brainwashing, the threat of the rack or burning at the stake, totalitarian control of the sources of information. Consensus reached by these

means does not guarantee rationality.”¹¹² However, even if we are not worried about “rationality” as such, these views still cannot guarantee that any of our consensus-based processes will have basic utility. If consensus is swamped with biases, prejudices, ideologies, etc., then we have no way of ensuring its use at all. Furthermore, any principled way of tooling biases, etc., to avoid the problems that are generated via consensus will have the same problems as well, i.e., the use of biases in a process used to correct for bias will itself be biased unless the process itself is idealized. That is why, in many cases, reasonable (CSE) views cannot avoid accusations of idealization. This is because they need a way to (a) orient or direct consensus-building to avoid said biases, prejudices, etc., from wildly tainting the consensus, as well as (b) a principled way to decide between consensuses reached in multiple communities of inquirers. As a result, no (CSE) view can offer an account of (a) or (b) that does not hinge on some type of idealized principle (or set of practices). What they are doing—but perhaps not acknowledging—is trying to fill the vacuum that the absence of a substantive notion of truth leaves. But this is an unnecessary compromise. As I hope to show in Chapters 4 and 5, we have no reason to abandon truth on the grounds that Rorty (1979) and other (CSE) views set out. There are alternative possibilities for available to us, veritistic possibilities, which arise from joining social epistemology, Bayesian decision processes, and the pragmatist conception of truth.

¹¹² Goldman (1992), p. 186-188.

Chapter IV

Veritistic Social Epistemology

4.1 Goldman (1999) and the Social Epistemological Project

It's not easy to parse Goldman's (1992, 1999) social epistemological project, but its importance to social epistemology cannot really be overstated. He is one of the first and earliest advocates of social epistemology from within analytic epistemology. For that reason, there are many points of conflict and confusion that arise from his most prominent work in social epistemology. I cannot hope to offer a veritistic social epistemological account myself without first placing it in some relation to Goldman's work, and, in particular, his (1999) book, *Knowledge In a Social World*. One of the first questions usually asked of the social epistemologist concerns the relationship between traditional analytic epistemology and with its focus on individual doxastic agents and social epistemology with its focus on knowledge communities, and the activities of those communities, groups, institutions, etc. The approach offered in this project tries to circumvent the question of the relationship between traditional and social epistemology by starting with the assumption that individualistic accounts of knowledge are only theoretically possible, and are not really normatively desirable given that methodological naturalism aims to adopt the best knowledge producing practices of science, and those practices rely on the social dimensions of knowledge production. In short, methodological naturalists really ought to hold that all knowledge is social knowledge. An explicit argument for this does not fall within the scope of this project, although it was briefly mentioned in Chapter 2. My hope is that the argument has been successfully embedded within the parameters of my account of social epistemology as a whole. So I will not dwell on this issue. I mention it in order to illustrate, not only a difference in starting points between Goldman's (1999) social epistemology and the one that I will offer here, but also to justify the starting point of this chapter. That is, for Goldman (1999), we do need an account of how he conceives of his social epistemological project in light of the more dominant tradition in analytic epistemology that focuses on the individual.

In a more recent account, Goldman's (2010) "Why Social Epistemology is Real Epistemology?" he attempts to place his own version of social epistemology within the three most dominant traditions in social epistemology, what he calls "revisionism," "preservationism," and "expansionism." In the previous

chapter, two dominant strains of social epistemology were referred to: “radical social epistemology” and “consensus-based social epistemology.” I argued that radical social epistemology was too problematic to be salvageable, and that consensus-based epistemology, in its current form, is untenable. It is important to note this here because Goldman (2010) also attempts to understand social epistemology, like traditional analytic epistemology, as existing within a spectrum of social epistemological concerns. But, as Goldman (2010) correctly notes, “revisionist” social epistemology, or “radical social epistemology,” is not philosophically workable. Goldman (2010) separates out revisionist social epistemology for criticism for primarily two reasons: (1) he attempts to clarify, in light of Alston’s (2005) criticism of social epistemology in general, that there are types of “social epistemology” that do fall outside of the domain of philosophy, (2) he attempts to separate out the revisionist project as an example of (1) in that it seeks to dismantle traditional notions of truth, knowledge, and objectivity in philosophy, and/or that it seeks to answer questions posed in non-philosophical disciplines, e.g., post-modern literature, sociology of science, science studies, etc. The motivations for Goldman’s (2010) analysis are also useful to note here. In a passing comment, Alston (2005) asks whether or not we can think of Goldman’s (1999) work, *Knowledge In a Social World*, as being a philosophical work, or rather, as being a contribution to sociology, or social science studies, etc. Goldman (2010) notes Alston’s (2005) comment and frames his reply by first reposing the question:

“Is social epistemology real epistemology?”... He raises it en passant, while noting that epistemology’s boundaries are controversial, drawn differently by different thinkers. To illustrate his point, he suggests that much of the material in my book on social epistemology, *Knowledge in a Social World* (Goldman, 1999), “would be rejected by many contemporary epistemologists as ‘not real epistemology’” (Alston 2005: 5). The epistemologists in question, says Alston, would relegate much of this so-called social epistemology to sociology, social psychology, or other social sciences, or perhaps to the philosophical foundations thereof.¹

Goldman (2010) concludes that there is good reason for such skepticism about some of the projects attributed to social epistemology. Thus he makes the above-mentioned distinctions between varieties of

¹ Goldman (2010) and also see: Goldman (2009), “*Why Social Epistemology Is Real Epistemology*,” in Adrian Haddock, Alan Millar, and Duncan Pritchard, eds., *Social Epistemology*, New York: Oxford University Press (2010). Here, however, I will refer exclusively to Goldman (2010).

social epistemology and claims that it is only revisionist social epistemology that does not fit the requisite criteria for being “philosophy.” He uses Rorty (1979) as a primary example of the revisionist approach.

However, I argued in the last chapter that, although radical social epistemology (RSE) can look quite similar to consensus-based social epistemology (CSE), because they are both motivated by the assumption that an individualistic, foundationalist frameworks for knowledge are false, what makes (CSE) arguments different from (RSE) arguments is that (RSE) argues for “social constructivism” in the strong, anti-realist sense.² Consensus-based social epistemology does not support such a radically relativistic metaphysics. This difference is what makes (CSE) arguments easier to negotiate, philosophically speaking, as well as easier wed to veritistic approaches, which is crucial for those of us with concerns about truth. Goldman (2010) argues that Rorty (1979) is a “social constructivist” whose “mantra was to ‘keep the conversation going rather than to find objective truth’ (1979: 377)” and thus does fit squarely into the “constructivist” category.³ He writes: “A few philosophers such as Richard Rorty fall into the same category, sometimes waxing emphatic about the bankruptcy of traditional epistemology. Rorty (1979) declared the “death” of epistemology and proposed a vague replacement for it in the form of a ‘conversation of mankind.’”⁴ But this isn’t a charitable reading of Rorty’s (1979) project.

On the one hand, yes, Rorty (1979) does in fact argue for a dismantling of a traditional analytic notion of truth. He also argues that we ought not to bother trying to replace “truth” with something else. Instead, he argues for a position that he calls “epistemology-as-hermeneutics.” Rorty (1979) writes:

Hermeneutics sees the relations between various discourses as those of strands in a possible conversation, a conversation which presupposes no disciplinary matrix which unites the speakers, but where the hope of agreement is never lost so long as the conversation lasts. This hope is not a hope for the discovery of antecedently existing common ground, but simply hope for agreement, or, at least, exciting and fruitful discussion.⁵

“Antecedent common ground” here is a synonym for “truth,” or “knowledge” in something like the Platonic sense. It is a real ontological entity that we can, if our methods are correct, have direct access to. Rorty (1979) argues against this assumption—the assumption that philosophy done correctly can give us a

² See: Goldman (1999) and Kitcher (1993).

³ Goldman (2010).

⁴ *Ibid.*

⁵ Rorty (1979), p. 318.

“mirror of nature”—or a mirror of what reality is really like. He also argues that if we could appreciate how the history of philosophy since Descartes, and through Locke and Kant, has allowed us to fool ourselves about such things, we would see that “epistemology,” and “truth” and “knowledge” with it, are merely fictions of our collective intellectual history. He claims: “The word knowledge would not seem worth fighting over were it not for the Kantian tradition that to be a philosopher is to have a ‘theory of knowledge,’ and the Platonic tradition that action not based on knowledge of the truth of propositions is ‘irrational.’”⁶ Thus Rorty (1979) argues if we can simultaneously acknowledge the history of the problem of truth, and then move on from it, philosophy would be the better for it.

An initial reading of this view can be misleading as it is constructivist in method but is not constructivist in terms of ontological content – and that difference is key. Rorty (1979) accepts that we can have knowledge of the external world and he does not deny the basic entities posited by a broadly empiricist metaphysics. What he denies is the thesis that there is a certain epistemic vantage point that puts the knower in the correctly “mirrored,” or objectively accurate, relationship with the external world and/or to those ontological entities. This does not mean that he thinks – as radical social epistemologists do – that “revisionist” or “radical” epistemology is rooted in a rejection of the correspondence theory of truth. That is, our epistemology is rooted in the problem that we can’t say anything substantive about the external world. Rorty’s (1979) problem is not of that stripe and his metaphysical outlook is far from being anti-realist. This is an important distinction that Goldman (2009) fails to acknowledge. As Kitcher (1994) argues, many of the constructivist views that one finds in epistemology and the sociology of science are really by-products of a rejection of the correspondence theory of truth – but this rejection, and the corresponding adoption of constructivism or “revisionism,” is most likely not a failure of philosophy so much as it is a failure of the imagination. It is a failure to conceive of the correspondence relationship in any other way, and a want to adopt Rorty’s (1979) “mirror of nature” hypothesis as metaphysical, not epistemological, fact. Kitcher (1994) notes: We can have a correspondence theory of truth that does not require some transcendent perspective from which we view reality. Instead, we can have multiple knowers, or communities of knowers, in multiple schemes of representation who are all realists about their ontological commitments, i.e., we could have subjects standing in different reporting relationships to

⁶ Rorty (1979), p. 356.

their respective environment(s), where ontology is determined by successful/unsuccessful interactions with their respective environment(s). There are problems with Kitcher's (1994) analysis here, but the overarching point still stands. Social constructivism is a result of a failure to imagine versions of the correspondence theory of truth that do not result in anti-realism—and Rorty (1979), despite the many confusions in his view, is not guilty of this.

That aside, revisionism doesn't succeed for many of the same reasons that radical social epistemology does not succeed. These reasons have already been covered in the last chapter and Goldman (2010) mostly treads upon the same ground: revisionist (or radical) social epistemologists, if they use these terms at all, have to deflate most traditional epistemic concepts, such as "knowledge," "truth," and "objectivity" to accommodate the social units or bodies under consideration. This makes avoiding irrationality a very difficult task. And although "rationality" can be a questionable normative epistemic requirement, we can also say that the radical subjectivism (including the radical prejudices) that revisionism allows would overburden any epistemological project. These views cannot correct for what might be called "epistemic taint," or the idea that too many undesirable epistemic factors, e.g., non-evidence based prejudices, bending to authority, concern for reputation, "group think," etc., preclude us from engaging in any meaningful, veridical epistemic task. Making truth a priority in our epistemic methodology is the best way of dealing with such epistemic taint. But how might we get truth in an inherently social context?

Goldman (2010) argues for an "expansionist" social epistemology that can take into account: (1) collective doxastic agents, and (2) social systems, but that is still truth-directed, or primarily veritistic. The most substantive defense of this version of social epistemology appears in Goldman's (1999) *Knowledge In a Social World*. An analysis of his argument here may also settle the issue mentioned above – Goldman (1999) seems to posit a two-pronged epistemological theory. On the one hand, he does not see epistemology as wholly individualistic; on the other hand, he does not clearly lay out how individualistic approaches square with his new social epistemological ones. This may not be necessarily philosophically burdensome, unless you subscribe to a methodological naturalism that denies individualistic epistemology in practice. As this is the preferable way of conceiving of naturalism, I will consider Goldman's (1999) social epistemological project in light of that. Goldman (1999) is no doubt a

naturalist – it is just unclear of what stripe his naturalism takes. In any case, I will use this chapter to do three things: (1) to outline Goldman’s (1999) argument, as his veritistic social epistemology is what motivated much of my own position, (2) discuss Goldman’s (1999) social epistemology as it stands in relation to his take on individualistic epistemology, and (3) to analyze areas where his social epistemological approach fails to capture what most social epistemologists take as a crucial feature of the project: adequately accounting for how a variety of knowers, with their own biases and prejudices (or “prior probabilities,” as I will say later), can contribute veridically and veritistically to knowledge production.

4.2 V-Values, Non-Epistemic Interests, and Questions of Interest

For Goldman (1999) there is “veritism” and then there is “veritism” as it is applied to collective doxastic agents. In both cases, veritism has social dimensions, but it is in the latter case that it especially pertains to multiple knowers. But let’s consider the details of veritism before considering how it applies to either individual or social epistemology. First, it’s important to note that Goldman (1999) analyzes veritism in terms of the correspondence theory of truth. He briefly argues that rival theories, including pragmatic (instrumental), epistemic, and deflationist theories are more problematic, hence less philosophically desirable, than the correspondence theory.⁷ I will not reiterate his arguments here, nor will I aim to clarify how his argument conflates “pragmatism” and the “pragmatic theory of truth” with “instrumentalism.” I will give a positive account of the pragmatic theory of truth as non-instrumental in Chapter 5. I mention Goldman’s (1999) preference for the correspondence account of truth only to say that this is how he fills out the “what is truth?” part of veritism. The other part of a veritistic analysis is the selection method. As a selection method, veritism helps knowers select which social practices “best advance the course of knowledge.”⁸ For this to work, there need to be two-orders of selection practices: “Call the set of practices from which veritism wishes to make a selection the target practices. [And] call the practices used to select among the target practices the selection practices.”⁹ However, Goldman (1999) identifies three problems with two-ordered selection:

- (1) Which selection practices should be used to choose among the target practices?
- (2) Is there any guarantee that the selection practices will accurately identify the veritistically best target practices?

⁷ Goldman (1999), pp. 41-68.

⁸ Goldman (1999), p. 79.

⁹ *Ibid.*

(3) Is there any guarantee that the process of selection will lead different veritistic theorists to agree on the choice of target practices?¹⁰

We might say that the first issue here is one of circularity, although Goldman (1999) argues that he's not convinced that social practices can be epistemically circular. This claim will need to be sorted out a bit more, so let's come back to it in a moment. The larger issue here is that we don't have a sense of how selection practices can identify the best veritistic target practices in manner that would make a veritistic analysis of social practices substantively different than non-truth seeking analysis of social practices – Goldman (1999) takes note of this, arguing that the problem with choosing initial selection practices is a problem for rival theories as well. If so, all things being equal, this should not count against veritism as a theory.¹¹ And if we can dismiss the initial selection issue – how exactly should veritism work to select the best practices to advance knowledge? Consider Goldman's (1999) example:

Suppose that a local weather bureau wishes to establish the veritistically best practice for predicting the weather. It has five weather forecasters available and wishes to 'pool' or amalgamate their daily judgments about the day's weather into a single prediction by the bureau, and they wish to maximize the bureau's forecasting accuracy. What social practice of amalgamating the experts' opinions is veritistically best? To make the case more concrete, suppose that two of these forecasters each has a 90 percent success rate in making categorical predictions of rain versus nonrain, and the other three have a 60 percent success rate each. These numbers are assumed to represent their propensities...not simply their past track record. Finally, assume that the forecasters' judgments are reached independently, and that two

¹⁰ Goldman (1999), p. 79-80.

¹¹ Goldman (1999) notes: "There is not difference on this score between veritism and any of its rivals. If the standard of evaluation for social epistemology were, for example, consensus consequentialism rather than veritism, it would still be necessary to utilize some sort of selection procedures to decide which practices are optimal according to the consensus consequentialist criterion. Whichever selection practices are used, they could be veritistically deficient... Clearly, there is no logical guarantee that the veritistic enterprise—that is, the enterprise of *trying* to identify veritistically desirable social practices—will succeed. But the absence of guaranteed success is no count against veritism as compared with its rivals. Like any enterprise, the project of veritism could fail; but unless one is a total defeatist, that is not much of a reason against trying. Similar points apply to the problem of reaching consensus on the veritistically best practices. There is no guarantee that two people who both try to identify the veritistically best practices will agree. If their initial selection practices differ and they staunchly resist change, they might easily select different target practices. This applies equally, however, to veritism's rivals" (80).

possibilities, rain and nonrain, are a priori equally likely. How should their daily predictions be pooled to arrive at the bureau's judgment?¹²

So what are the possibilities?

- (1) The first is unweighted majority rule: the prediction favored by a majority of the experts should be the bureau's prediction (belief). This yields a probability of correctness for the bureau of .877;
- (2) The second possibility is dictatorial rule: let the bureau adopt the prediction of its most competent expert. In the present case of a tie for maximal competence—the two experts with 90 percent accuracy—the bureau chooses one of the two at random and lets his or her judgment constitute the bureau's judgment. This would yield a correctness probability for the bureau of .900;
- (3) The third possibility uses weighted voting: A theorem due to several independent authors (Shapley and Grofman 1984, Nitzan and Paroush 1982) says that a maximally truth-conducive weighting scheme is one that assigns a weight w_i to each expert i that satisfies the following formula: $w_i \propto \log(p_i/(1-p_i))$ where p_i represents the probability that expert i makes a true prediction. In the present example, one weight assignment that would conform with this rule is: .392, .392, .072, .072, .072. That is, the higher voting weight, .392, is given to each of the two forecasters with 90 percent accuracy, and the lower weight, .072, is assigned to each of the three forecasters with 60 percent accuracy...Using a weighted voting scheme with these weights, and letting the weighed vote of the five forecasters determine the bureau's own judgment, the bureau's probability of correctness is .927.¹³

Clearly the best results are produced via the weighted voting method and/or weighted voting theorem. However, as Goldman (1999) asks, how would the members of the weather bureau – or any collection of doxastic agents – come to agree that weighted voting is the best veritistic practice? He suggests three ways in which collective doxastic agents might agree that a particular practice is veritistically optimal, (i) by providing external justification for the veracity of the practice, e.g., a well-supported proof or a theorem (ii) by choosing a selection practice, e.g., majority rule, dictatorship, etc., and using that to justify a chosen

¹² Goldman (1999), p. 81.

¹³ Goldman (1999), p. 81.

target practice, or (iii) by coming to some interpersonal consensus on the best selection practice and using that consensus to justify a chosen target practice. But don't such justifications run the risk of circularity?

Goldman's (1999) argues that it is not clear that social practices can be epistemically circular in the same way that individual practices can. Circularity is only important – or relevant – insofar as it is individuals who engage in initial selection practices. This is because individual selection relies upon individual perception and memory, and the epistemic justification for these functions is difficult to ground epistemically. But this is not a problem for social epistemology in particular; it is a problem for all epistemological projects that do not rely on mere speculation. He argues: "Does anyone seriously propose that we cease all attempts to assess social practices with the help of perception and memory? Does anyone propose that from this day forward we should cease regarding perception as more reliable than idle speculation?"¹⁴ This does not so much as answer the circularity question as proclaim it a question for epistemology in general. And for this reason, a charitable reading of veritism should not consider this issue as reason enough to disregard it in general.

Moving on from the issue of initial selection, we still want to see how veritism or a veritistic analysis works. We have seen how collective doxastic agents might come to some agreement about the best practice to produce the best veritistic outcome – but what does Goldman (1999) say about the outcome itself? How ought we analyze the outcomes? First, he suggests that we adopt a trichotomous approach to understanding the veritistic outcome, or knowledge output. Goldman (1999) considers "knowledge" in what he calls the "weak" sense of "true belief," so; in this case, the object of evaluation will be belief. And here he suggests taking a trichotomous approach similar to one basic division of credences (or "credal attitudes") made by subjective Bayesians: (i) belief as a state of "knowledge," or true belief, (ii) belief as a state of ignorance, or withheld belief, and (iii) belief as a state of error, or false belief. These approaches will generate a "V-value," or veritistic-value, in terms of a probability assignment, and we can evaluate V-values on two levels: beliefs states, like mentioned above, will have a "fundamental" veritistic value or disvalue according to their belief state outputs; social practices have "instrumental" veritistic value according to the ability to "promote or impede the acquisition of fundamental

¹⁴ Goldman (1999), p. 87.

veritistic value.”¹⁵ As the trichotomous approach employs credences or “credal attitudes,” Goldman (1999) aims to align the belief states of knowledge, ignorance, and error with subjective degrees of belief (DB) and/or subjective probabilities. Although, he argues, this is primarily a matter of convenience, as opposed probabilistic accuracy (I will come back to this issue later):

There is no entirely unproblematic method of mutually translating the trichotomous scheme and the DB scheme. For example, should plain ‘belief’ in the trichotomous scheme be translated as ‘DB .50’? Again, this is an imperfect translation. Although no perfect translations are known to me, my V-value assignments will sometimes appear to presuppose certain equivalences. This is more a matter of convenience than a matter of conviction that there are perfect equivalence here.¹⁶

Now Goldman (1999) suggests that we begin the veritistic project with what he calls “questions of interest.” He suggests, “V-value should always be assessed relative to questions of interest.”¹⁷ This is because a holistic view of an individual’s belief states should not be analyzed in terms of questions that he/she does not have an explicit interest in; that is, an individual’s ignorance about a given P should not count against their V-values as a whole, as P might be question that they merely haven’t yet entertained. Thus veritistic analysis should only count for or against an individual for those questions he/she has entertained.

Now that we have some account of initial selection, an account of veritistic states, and a two-pronged account of how to analyze veritistic outcomes in terms of “questions of interest” – can we reasonably apply veritistic analysis to all of our social practices?

Goldman (1999) argues that the collective case, or the case of applying veritistic evaluation to social practices, will look quite similar to the individual case, except, in the collective case, we will need a method to aggregate the V-values of all doxastic agents involved in analyzing the question of interest. He suggests that one simple way of doing this by having individuals complete the veritistic task and then taking the group’s mean V-value. He also suggests that an aggregate analysis can also be achieved by taking the group’s root mean square, or a weighted average. He claims, “These measures all share the property that if some individuals’ V-values rise and no individuals’ V-values decline, the aggregate V-

¹⁵ *Ibid.*

¹⁶ Goldman (1999), p. 88.

¹⁷ Goldman (1999), p. 89.

value rises.”¹⁸ Ultimately he leaves the choice of aggregation method open and suggests that those with particular mathematical bents choose the method that appeals to them. However, as List (2008) argues, some judgment aggregation methods have trouble weeding through the inconsistency and irrationality of individual judgments, so we need to be thoughtful in choosing our methods of aggregation. Initially this does not seem to be a problem for Goldman’s (1999) veritistic view because his model aligns knowers with their credences in true propositions – and, presumably, if an individual utilizing the trichotomous scheme assigns a .50 probability to a proposition for which he/she is withholding belief, there is a method for increasing V-value in relation to that true proposition. I.e., there is a method or practice by which the knower’s V-value will definitively change to 0 or 1, respective to the question of interest. Goldman (1999) writes:

Suppose that question Q_1 begins to interest agent S at t_1 , and S applies a certain practice π to question Q_1 . The practice might consist in a certain perceptual investigation, such as visually scanning the environment, or it might consist in asking a friend for information and drawing a conclusion from her response. One possible result of the practice is to change S’s state of belief (vis-à-vis question Q_1) at t_2 .¹⁹

Let’s pause here to highlight an emerging issue – one that will be discussed in more detail in 4.3. Goldman’s (1999) analysis of V-value in terms of a subjective probability assignment confuses the role that subjective probability aims to play in formal approaches in epistemology. Goldman’s (1999) view suggests that an individual (i) has an important stake in the P under consideration, and (ii) has a role in determining the subjective probabilistic value of said P. This suggests that an individual can analyze and assign their own belief states, at least initially. But how does Goldman (1999) square this with the correspondence theory of truth that he also supports? If the correspondence theory of truth is true, however broadly conceived, an individual cannot come to know a proposition is true or false via veritistic evaluation – at best, an individual can engage in better and worse practices, π , to increase his or her V-value. The knower would stand in a subjective probabilistic relation to a proposition whose truth-value is determined by the correspondence relationship - although it is not clear how an individual would know if a given practice brought them into the correct relationship with the truth. Goldman’s (1999) view makes

¹⁸ Goldman (1999), p. 94.

¹⁹ Goldman (1999), p. 90.

epistemology responsible to metaphysics in way that seems to preclude any meaningful use of subjective probabilities. This view also seems to involve the knower in something of a no-win situation in regard to the truth. This is because it would be possible for S to have a 1.0 degree of belief that P is true and for S's belief state in this case to be called "knowledge," under veritism, while P's actually being true would still actually be knower-independent. This seems fine in the case where our unconditionalized degree of belief in P is subjectively determined and we have a good way to engage in belief-updating, i.e., conditionalization, etc., but not in the case where we have already said that P would be either true or false via correspondence to some entity X. What would work with this view is some account of objective probability to which we would want our credences to match via some normative requirement. I will argue for this in Chapter 5.

4.3 Extrapolating from Testimony: Does Goldman (1999) Give Us a Bayesian Social Epistemology?

One way to make sure that our social practices avoid the problem described above is to suggest a general social practice that is conducive to truth-making in a similar way that many epistemologists think that the correspondence view of truth is conducive to truth-making. This is the closest Goldman (1999) gets to suggesting that veritism should adopt something like an objective probability view to anchor the truth-making process. He starts by suggesting that the social practice most applicable to veritistic analysis is testimony, and because testimony covers such a vast range of communicative practices and behaviors, it could potentially benefit greatly from the truth-orientation veritism provides. Unfortunately, it is as of yet unclear how any practice, π , including testimony, does this and actually contributes to an individual's better or worse V-value. So Goldman (1999) weds Bayesian inference as a practice to testimony in a way that looks at testimonial reports in terms of likelihoods – and he claims that Bayesian inferences can really be used for any type of reporting environment. At first glance, this wedding of the practice of testimony, and/or reporting in general, to Bayesian inference might seem a promising route by which we could make veritism's truth-direction, V-values' subjective assignment, and the correspondence relationship consistent. Under the conditions outlined, below, Goldman (1999) claims that the use of Bayes' Theorem, as a practice for bettering V-values, is "objectively likely" to raise an individual's degree of knowledge. So what are these conditions?:

When a reasoner starts with accurate likelihoods (analogous to true premises), it is objectively probable that Bayesian inference will increase his degree of knowledge (truth possession) of the target proposition. More precisely, after using Bayes' Theorem, the objectively expected degree of truth possession (or V-value) associated with the reasoner's posterior DB [degree of belief] will be greater than the objectively expected degree of truth possession (or V-value) associated with the reasoner's prior DB...If a probabilistic reasoner begins with inaccurate likelihoods, she cannot expect the Bayesian method to improve her V-value...So let us ask what the Bayesian method will do if the reasoner has accurate likelihoods, that is, if her subjective likelihoods match the objective likelihoods. Here is where we locate our mathematical results. If subjective likelihoods match objective likelihoods, use of Bayes' Theorem leads to an objectively expected increase in degree of truth possession (V-value). In other words, under these conditions a Bayesian practice exhibits positive V-value.²⁰

There are a number of problems with Goldman's (1999) treatment of Bayesianism here. But let's just start with (i), "When a reasoner starts with accurate likelihoods...it is objectively probable that Bayesian inference will increase his degree of knowledge." First, this only seems likely if what we mean by "accuracy" is actually something closer to "objectivity." A subjective Bayesian is primarily interested in self-consistency, not external or objective accuracy, so to be "accurate" in this sense will mainly involve assigning "accurate" priors according to "one's own lights" – and for this interpretation of Bayesianism, that is perfectly fine. Another sense in which a reasoner may have "accurate likelihoods," is in that they match (or are constrained by) something objective, such as chances, frequencies, and/or the laws of statistical mechanics. In this case, to say that a reasoner starts with accurate likelihoods is just to comment on how their priors match those determined by the theory that constrains them – it does not tell us much at all. It is unclear why Goldman (1999) thinks that individuals will start with accurate objective likelihoods unless they are either (a) non-subjective and determined by some theory of constructing priors, or (b) based on trivial, mostly chance-based, propositions, such as those pertaining to coin tosses. Given the following claim, it is clear that Goldman (1999) thinks that priors are clearly subjective and determined by the agent (he also argues this in earlier chapters): "If a probabilistic reasoner begins with

²⁰ Goldman (1999), pp. 116-117.

inaccurate likelihoods, she cannot expect the Bayesian method to improve her V-value..."²¹. However: "If subjective likelihoods match objective likelihoods, use of Bayes' Theorem leads to an objectively expected increase in degree of truth possession (V-value)" (my italics).²²

What seems clear is this: Goldman (1999) cannot use the Bayesian model in this way. He does not offer any justification for why a "reasoner's" subjective credences ought to match objectively likelihoods – although Lewis' (1980) Principal Principle is available to him – and he does not explain how a reasoner's credences can be sufficiently constrained by objective factors at the outset of probabilistic inquiry. He seems to want subjectivity but does not want to deal with the "problem of priors." Instead, he claims that he's only concerned with the priors that initially match some objective likelihood. If this method aims to extend veritism to all of our reporting practices and behaviors, then Goldman (1999) has made too small a subset out of those reporting behaviors. Veritism will probably only apply to those few individuals (engaging in a social practice) who are fortunate enough to have accurate priors. Or Goldman (1999) is just more optimistic about how people construct their belief systems than seems wise.

4.3.1 Goldman (1999) and the "Problem of Priors"

Goldman (1999) would like the pay-off of the objective Bayesian view while holding on to the self-assignment of probabilities possible via subjective Bayesianism. It is not clear which end he would ultimately give up, but he cannot reasonably hold on to both features. If we extrapolate from his wanted the trichotomous scheme to subjective probabilities, i.e., the "DB-view," as well as his insistence that subjective probabilities must match objective likelihoods, it's fair to say that Goldman (1999) wants to be subjectivist about probability assignments. If so – and to truly do justice to his view – he will need to say something different about constraining subjective priors such that it's possible that they match objective likelihoods. That is, he will need to say something substantive about the problem of priors.²³ There are at

²¹ *Ibid.*

²² *Ibid.*

²³ Talbott (2008) gives a quick overview of the problem of priors in the *Stanford Encyclopedia of Philosophy*. The most pressing point for subjectivists is as follows: "The problem of the priors identifies an important issue between the Subjective and Objective Bayesians. If the constraints on rational inference are so weak as to permit any or almost any probabilistically coherent prior probabilities, then there would be nothing to make inferences in the sciences any more rational than inferences in astrology or phrenology or in the conspiracy reasoning of a paranoid schizophrenic, because all of them can be reconstructed as inferences from probabilistically coherent prior probabilities." See: Talbott, William (2008). "Bayesian Epistemology," *The Stanford Encyclopedia of Philosophy* (Fall 2013 Edition), Edward N. Zalta (ed.), <<http://plato.stanford.edu/archives/fall2013/entries/epistemology-bayesian/>>.

least two different tactics that Goldman (1999) could here: (1) he could adopt Lewis' (1980) "Principal Principle" or something like it to constrain priors, or (2) he could advocate for a convergence thesis – common amongst many subjectivists – which claims that end-state probabilities (determined over infinite time) converge to either 1 or 0. Thus initial probability functions, no matter how they are formed, will converge toward truth and produce the same end-state probabilities as objective probabilistic methods. This will work as long as an individual's initial credence about end-state probabilities is itself countably infinite, and does not violate the sigma-additivity axiom of probabilism. Of the two options, adopting Lewis' Principal Principle might require the least suspension of disbelief and might be the most reasonable in terms of constraining priors. That is because, more so than the second option, Lewis' interpretation of chance presents us with more of a known quantity regarding future probabilities.

In "A Subjectivist's Guide to Objective Chance," Lewis (1980) argues that subjectivists do not have to take an antagonistic view of the concepts of science, and particularly the concept of objective chance. Subjectivists can talk meaningfully about objective chance and incorporate it into their credence functions as a normative constraint via something like a "matching principle." Lewis (1980) calls this normative requirement the "Principal Principle." He writes:

We subjectivists conceive of probability as the measure of reasonable partial belief. But we need not make war against other conceptions of probability, declaring that where subjective credence leaves off, there nonsense begins. Along with subjective credence we should believe also in objective chance. The practice and the analysis of science require both concepts. Neither can replace the other. Among the propositions that deserve our credence we find, for instance, the proposition that (as a matter of contingent fact about our world) any tritium atom that now exists has a certain chance of decaying within a year. Why should we subjectivists be less able than other folk to make sense of that?²⁴

An easy way to understand the Principal Principle is to consider a simple formulation of our prior probability in the chance of some event's occurring at some specified time and/or some specified time and possible world. For simplicity, we'll only consider events at some specified time, not time-and-

²⁴ Lewis (1980), p. 263.

world.²⁵ So consider the following formulation: $hp(A | \langle \text{chance of } t \text{ at } x \rangle)$. We can read this as a formulation of our prior (or hypothetical) probability of A given the chance of A at t being x . Now, Lewis (1980) notes, consider a coin toss that is going to happen at noon today. Assume that you are not worried about trick coins, or other such nonsense, so you think that the chance of the coin's coming up heads will be .50, just like the chance of the coin's coming up tails. Now consider the proposition: "The coin tossed today at noon today falls heads."²⁶ What should your credence or degree of belief be in that proposition? Given that the coin is a fair coin, and the set up of the coin toss can be reasonably assumed to be a fair setup, your credence should match the objective chance of the coin's coming up heads being .50.

But what happens in the case where you somehow know information about that future coin toss? – Or what if a reliable witness told you what the results would be already?²⁷ Should your credence stay at .50? No, in this case, Lewis (1980) argues, your credence should match something approximating certainty, if not certainty itself. Thus your degree of belief in the proposition about the coin toss should be 1 (or 100%), or at the very least, it should be very, very close to 1. But if you had information about what x will be before time t , you are most likely in possession of "inadmissible evidence."²⁸ So the principle should really be formulated like, $hp(A | \langle \text{chance of } t \text{ at } x \rangle) + E$, where E is "admissible evidence." Lewis (1980) writes: "Admissible propositions are the sort of information whose impact on credence about outcomes comes entirely by way of credence about the chances of those outcomes."²⁹ In terms of the coin toss, "once the chances are given outright, conditionally or unconditionally, evidence bearing on them no longer matters."³⁰ Thus the evidence is inadmissible by default. There are other types of inadmissible evidence – although Lewis (1980) argues that he cannot offer a definition of admissibility – but generally this will include historical evidence after time t . Goldman (1999) may be able to make use of the Principal Principle, if for no other reason than it provides one substantive framework for

²⁵ Lewis (1980) writes: "**The Principal Principle.** Let C be any reasonable initial credence function. Let t be any time. Let x be any real number in the unit interval. Let X be the proposition that the chance, at time t , of A 's holding equals x . Let E be any proposition compatible with X that is admissible at time t . Then $C(A|XE) = x$ " (266).

²⁶ Lewis (1980), p. 264.

²⁷ Lewis (1980), p. 265.

²⁸ Lewis (1980), pp. 272-273.

²⁹ Lewis (1980), p. 272.

³⁰ *Ibid.*

subjectivists to constrain priors. Of course, it is not without its own problems. The primary ones being that Lewis (1980) assumes (i) that the past is no longer “chancy,” or as Meacham (2005) writes, “that a chance distribution assigns only trivial chances (0 or 1) to events in the past;” and that (ii) the “best system” theory of chance will allow us to pick out unique past chances.³¹

4.3.2 Replies: Kitcher (2002) and Fallis (2002)

Fallis (2002) agrees that there is reason to be skeptical about reading too much into the “chanciness” of the universe particularly as it pertains to past events. Although the universe may in fact be chancy, we simply don’t have access of the objective probabilities of most of those non-trivial events. Simply put, there are just not enough known objective probabilities to go around. He writes:

It should be noted that there are certainly some cases where it is feasible to determine the objective probability of a chancy event. For example, we can determine the objective probability that a particular fair coin will land heads when I flip it five minutes from now. However, it does not seem to be feasible to precisely determine objective probabilities in most of the everyday cases where we are trying to increase our degree of knowledge. For example, it does not seem feasible in one of Goldman’s favorite applications of the Goldman/Shaked theorem: viz., the evaluation of evidence presented in a court of law. Murderers just do not exhibit the high degree of uniformity that fair coins do.³²

Fallis (2002) presents this as a criticism of Goldman’s (1999) hope for accurate initial likelihoods to result in an increase in net V-value via something like an objective Bayesian analysis, but the problem with “chanciness” and objective probabilities presents a problem for Goldman’s (1999) use of objective likelihoods and for Lewis’ (1980) Principal Principle. We may be able to constrain our subjective priors via the chance of coin tosses, but most areas of interest concern for us require vastly more complicated probabilistic determinations. That said, I still claim that Lewis’ (1980) Principal Principle, or something like it, will be the best option for saving Goldman’s (1999) view. This may be a case where something is just better than nothing. There has to be something to constrain the initial priors other than the hope that individuals begin the probabilistic calculations from extremely rationally advantageous positions. For a

³¹ Meacham, Christopher J. G. (2005). “Three Proposals Regarding a Theory of Chance,” in *Philosophical Perspectives*, 19 (2005): 281-307.

³² Fallis (2002), p. 11-12.

robust epistemological view to start with this assumption is rather unusual. Unless Goldman (1999) claims that only a small percentage of individuals' are really in a position to benefit from (i) veritistic social practices in general, and/or (ii) increases in their V-value with respect to reporting events.

Kitcher (2002) considers another side of Goldman's (1999) argument when he questions how his view goes about producing V-values for collective doxastic agents. The problem isn't one of baseline judgment aggregation, as Goldman (1999) suggests that we could take a mean of collective V-values, such that if four individuals were trying to find the best route to San Jose, and we were considering their initial credence in p, that driving is the "best route," and their credence change over time, we might get a distribution such as:

t_1	t_2
S ₁ DB(p) = .75	S ₁ DB(P) = .90
S ₂ DB(p) = .90	S ₂ DB(P) = .50
S ₃ DB(p) = .20	S ₃ DB(P) = .40
S ₄ DB(p) = .50	S ₄ DB(P) = .75

To get an aggregate account of their initial credences, we might, as Goldman (1999) suggests, take the mean. So that at t_1 the group mean in p is 0.5875, and the group mean at t_2 after applying a social practice, π , is 0.6375. In my example, the collective V-value is increased by application of social practice, π ; however, in this case, the increase via the application of π may only be incidental. S₁, S₃, and S₄ had V-value increases, but there was nothing inherent about the social practice that I applied – nor will there be anything inherent in any social practice that Goldman (1999) endorses – that clearly prevents a net V-value loss after the application of π .³³ Goldman (1999) attempts to use the Bayesian model as a general model for how individual and aggregated V-value can also increase on application of specific practices, but, as I hope to have showed, his use of the Bayesian model fails, thus the problem with aggregation remains.

³³ Goldman (1999) also suggests taking the root mean square, or a weighted average, instead of just the mean, but this will not prevent individual decreases in V-value over time, via application of a practice, π , from decreasing the net V-value.

Kitcher (2002) points out that this model of aggregation will also suffer from what he calls the “tyranny of ignorance.” If we cannot say anything definitive about the truth-producing capabilities of any particular social practice, then the employment of that practice will be subject to both human error and human ignorance – not to mention the less desirable traits of epistemic arrogance, malevolence, and/or hubris. We will not have anything to protect collective V-values from any subject’s ignorance about (i) evidence or (ii) proper implementation of social practice, π . More importantly, Goldman’s (1999) view comes alarmingly close to having the same problems with epistemic taint as consensus-based models of social epistemology. This is beyond undesirable for a veritistic view, and Goldman (1999) offers no alternative here. Kitcher writes (2002):

Were we to offer a natural generalization of his measure of V-value, taking the collective V-value of a question to be measured by the number of people for whom that question is interesting, we'd have an obvious threat of the tyranny of ignorance. Goldman's distaste for that outcome is clearly expressed in the distancing of his own views from pragmatism (to which I've already referred). If mere aggregation of raw preferences for various types of research projects won't do, what will? Goldman doesn't tell us. In my judgment, the question is too important to ignore and it can be addressed without leading to the ills of vulgar democracy. Note first that, as soon as one gives up on the idea that the agenda for inquiry is set for us by nature, then it's a quite serious question whether the current social processes for framing our research accord with the promotion of collective V-value.³⁴

The “ills of vulgar democracy” to which Kitcher (2002) alludes is here aimed at Goldman’s (1999) view that groups and individuals are not only warranted in choosing their own “questions of interest” – whatever they be – but that neither groups nor individuals have a method for weeding out the baseline problem of ignorance. Thus we can pick which questions that we are interested in, in terms of the truth, but our prior ignorance about such questions may never be flushed out of the calculus by applying a selected social practice, π , to our prior V-value or $DB(p)$ because Goldman (1999) offers us no way of preventing the corruption of the social practice that we are implementing. If this is the case, we can have

³⁴ Kitcher (2002), p. 194.

a situation where our own ignorance is reflected back to us via our truth-producing methods as the truth. And this is problematic indeed.

4.3.3 Concluding Remarks

Kitcher's (2002) comments on Goldman's (1999) veritistic social epistemology really highlight how difficult it is to have our social epistemology both ways: (i) rooted in the actual behaviors and practices of doxastic agents, (ii) rooted in a methodology that can prevent those practices from being subject to epistemic taint. It is beyond the scope of this chapter to offer a wholesale ameliorative approach for Goldman's (1999) project. This is not because I agree with Kitcher's (2002) claim that Goldman (1999) would be better off arguing for an "ideal conditions" view of truth-production than one that so haphazardly situates the activity of truth-production with individuals and/or groups. (And, really, Kitcher (2002) would not be a fan of an ideal view either.) It is because Goldman (1999) does not seem to have come to any consensus on how to deal with the issue of radical subjectivity, ignorance, and general epistemic taint. His method ends up prioritizing doxastic subjects who are already in possession of fairly accurate information about the world (i.e., accurate objective likelihoods), and it is not clear that such knowing subjects even exist. Or at least it is not clear enough such that he should rest his epistemological view on their existence. Additionally, and paradoxically enough, Goldman (1999) allows radical subjectivity in terms of questions of interest, selected social practices, and methods for evaluating collective V-value. It is almost as if he ignores the issue of subjectivity at the most significant level of analysis and adds subjectivity at the least significant level. In any case, we can realize the importance of Goldman's (1999) approach to social epistemology while also realizing that it's time to move beyond it. As any good naturalist should know, we learn best from our mistaken hypotheses. So in the next chapter, I will co-opt Goldman's (1999) mistaken hypotheses and try to shed a more agreeable light on veritistic social epistemology – but perhaps not in a manner that Goldman (1999) himself would support. As Kitcher (2002) notes, Goldman's (1999) "attempt to resist pragmatism...seems misplaced."³⁵ I agree – and I hope to show that a pragmatic analysis is actually the best way to save the veritistic social epistemological project.

³⁵ Kitcher (2002), p. 192.

Veritistic Social Epistemology in a Pragmatist Framework

5.1 Pragmatism, Realism, and Convergence

It may be tempting to claim that the classical pragmatists would have been sympathetic to deflationism both metaphysically and epistemologically (See: Price (2011), Brandom (2011)), but this is a mistake. The urge to deny both meaningful ontological objects and correspondence has motivated “pragmatists” of a “radical” stripe since the initial publication of *Pragmatism*. I say “radical” in the only sense relevant to this project – and in the same sense expressed by James (1909, 1978) – this refers exclusively to those who aim to do epistemology (or “epistemology-as-hermeneutics”) without appeal to ontology and truth.¹ James (1909, 1978) went to great lengths in *The Meaning of Truth* to correct the misinterpretations about truth that arose from *Pragmatism*, with the “radical” view being his primary grievance. For the classical pragmatists, realism is still “real” realism, it just not meant in the sense often argued for by correspondence theorists who put their metaphysics first. James (1909, 1978) and Dewey (1917) particularly objected to the idea of “truth” as a *property of reality* instead of a property of a belief or an idea. If truth is a property of reality, then the knower would simply need to get into the correct “correspondence relationship” with the external world to be said to be “in possession” of truth. But James (1909, 1978) argues that this makes “truth” a constituent property of the world itself and there’s not much to motivate this view: “Realities are not *true*, they *are*; and beliefs are true *of* them.”² To assume otherwise is to assume that we can say something definitive about our ontological commitments without doing the appropriate “leg work,” or without engaging in the process of justifying our beliefs in a scientific, or broadly naturalistic, manner. This is why epistemology must come first in our discovering the constituent features of the external world.

¹ James (1909,1978) writes: “A distinction is sometimes made between Dewey, Schiller and myself, as if I, in supposing the object’s existence, made a concession to popular prejudice which they, as more radical pragmatists, refuse to make. As I myself understand these authors, we all three absolutely agree in admitting the transcendency of the object (provided it be an experienceable object) to the subject, in the truth-relation. Dewey in particular has insisted almost ad nauseam that the whole meaning of our cognitive states and processes lie in the way they intervene in the control and revaluation of independent existences or facts. His account of knowledge is not only absurd, but meaningless, unless independent existences be there of which our ideas take account...” ([9] 175).

² James (1909, 1978), p. [106] 272.

Similar to the semantic deflationist project, classical pragmatists also want to argue that “truth” is primarily a property of an activity. In this case, the activity is the process of verifying ideas or beliefs (because of this, “truth” is also a linguistic property as well, although it is not a *solely* linguistic property). But this isn’t the entire story. The activity still has to track on to the world somehow – and particularly in a way where we can say something definitive about the ontological objects of belief – that is, we can say we have achieved something like “satisfaction” with our current beliefs and/or ideas. We experience what can only be called emotional or psychological “satisfaction” when our ideas are true, or when they agree with what is “really real.” Thus when pragmatists claim that the “truth” is what relieves us from a state of anxiety, it is because we have confirmed something universally (or absolutely) true about our belief, i.e., we confirm that it “fits” or “agrees” with the world in some non-arbitrary manner. James (1909, 1978) writes:

The concrete pointing and leading are conceived by the pragmatist to be the work of other portions of the same universe to which the reality and the mind belong, intermediary verifying bits of experience with which the mind at one end, and the reality at the other, are joined. The ‘satisfaction,’ in turn, is no abstract satisfaction *überhaupt*....It is the *inherent relation to reality* of a belief that gives us that specific *truth*-satisfaction, compared with which all other satisfactions are the hollownest humbug. The satisfaction of *knowing truly* is thus the only one which the pragmatist ought to have considered.³

The question is now: how do we know when we have achieved “truth-satisfaction”? James (1909, 1978) takes his cues from Peirce (1877) on truth and the continuity of inquiry (as well as from a contentious debate with Bertrand Russell) and argues: “Reality means experienceable reality, both it and the truths men gain about it are everlastingly in process of mutation—mutation towards a definite goal.”⁴ Both James (1909, 1978) and Dewey (1938) took Peirce’s (1877) convergence view and method of inquiry seriously. However, we also know that Dewey’s (1938) attempt to formalize Peirce’s (1877) notion of convergence by wedding it to a strictly frequentist theory of probability only further grounded claims of

³ James, William (1909, 1978). *Pragmatism and The Meaning of Truth*. Cambridge, Massachusetts: Harvard University Press, pp. 270-272.

⁴ James (1909, 1978), p. 107. Also see: Russell, Bertrand (1908) “Transatlantic ‘Truth’” in *Albany Review*, or Russell (1909), “Transatlantic ‘Truth’” in *Edinburgh Review*. For a discussion of both, see: Berkeley, Hastings, “The Kernel of Pragmatism” in *Mind*, New Series, Vol. 21, No. 81 (Jan., 1912), pp. 84-88. It is important to note that this (1912) account also suffers from many of the same interpretative issues as Russell’s (1908) own account.

theory idealization (See: Levi 2006).⁵ The activity of verification relies on the end-state of convergence and allows the pragmatist to defer concrete metaphysics to the end of inquiry. Without a robust conception of convergence the pragmatist theory of truth would truly be reduced to instrumentalism. I have argued for a version of pragmatist-inspired naturalized epistemology that requires all traditionally “naturalized” projects to be *at minimum* methodologically “social.” This argument, as I hope to have shown, is far from new. However, the issue with most views of this stripe is that truth gets downgraded in favor of something like broad consensus or consensus-based “objectivity.” While the pursuit of capital-T “truth” may seem a burdensome leftover of traditional epistemology, I argue that without truth in *this sense* we cannot avoid the crippling epistemic taint of our social epistemological practices. Thus I claim that if naturalized social epistemology looks to pragmatism for its initial inspiration, it ought to look again for a resolution to the problem of truth. In the last chapter, I outlined one way we might consider a truth-oriented, or veritistic, social epistemology offered by Goldman (1999). I also hope to have shown how Goldman’s (1999) approach, while well motivated, was ultimately unsuccessful. So in this chapter I will offer my own account – one that I argue *respects* consensus-based approaches to social epistemology while also aiming for truth in the capital-T sense.

5.2 Pragmatism and the Convergence Thesis: A Bayesian Approach

A better framework for veritistic social epistemology comes first from taking the pragmatist convergence thesis seriously. While some philosophers (e.g., Hookway (2004)) object to convergence on the grounds of idealization, this objection most likely arises from problems inherent to “frequentism” as the formal apparatus of belief-change. The model proposed by the pragmatists, particularly in Dewey’s (1938) *Logic*, does not offer realistic *epistemic* constraints on belief-change or “belief-updating.” To reiterate an important point made by Fuhmann (2006), the potential goal is for all doxastic agents to be able to exercise the same selection function:

⁵ It is clear that Peirce had an evolving view of probability that started with probabilities as “propensities” and moved toward something akin to subjective Bayesianism without the Bayesian apparatus, i.e., priors could be subjectively assigned but should mirror something like objective chances, frequencies, etc. It is clear that Peirce did settle on a view of probability as frequencies, or “frequentism,” for the pragmatic account of truth. Levi (2004) has a very good discussion of this in Misak, Cheryl (ed.), “Beware of Syllogism: Statistical Reasoning and Conjecturing According to Peirce,” in *The Cambridge Companion to Peirce*, Cambridge, United Kingdom: Cambridge University Press, p. 257-286.

If *X* and *Y* could be brought to exercise the same choices—if, in other words, they would eventually adopt the same selection function...then given sufficient exposition of evidence, their states of belief will indeed converge in the limit.⁶

But how should we choose a selection function? And even if we can choose an appropriate selection function, how can we reasonably apply said function to groups and/or collective doxastic agents? The latter is a theoretical *and* a practical question. Even if individuals can be compelled to see the selection function as a normative requirement for belief-updating, how can we reasonably track the activities of individuals as “collectives” or “voluntary groups”?⁷ In this section, 5.2, I will argue for employing a Bayesian selection function over the frequentist one offered by the classical pragmatists. In the following section, 5.3, inspired by a claim made by List and Dietrich (2014), I argue that any Bayesian aggregation method, once epistemically justified, is suitable for use by a community of inquirers. I argue that voluntarily formed collective doxastic agents only ought to employ a principle of *epistemic disagreement*, or peer disagreement, to select their Bayesian aggregation method. Here I appeal to recent work done by Christensen (2013) and Elga (2007, 2010) on Bayesian epistemology and epistemic disagreement. In section 5.4, I will briefly suggest one potential alternative, also offered by List and Dietrich (2014), for aggregating Bayesian judgments – *supra-Bayesian pooling*. In the last section, 5.5, I will argue that our new selection function will allow for a richer notion of the pragmatist theory-practice divide as it pertains to social epistemological practices and offer some preliminary thoughts on how we might “truth-track” those practices on a large scale.⁸ I argue that a better selection function will give a new understanding to the pragmatic “method of science” such that we can allow our social epistemic practices to “run wild” in terms of fecundity, diversity, epistemic goals, societal implications, etc., without having to also constrain *what counts* as a knowledge-producing activity.

5.2.1 Scoping Methods: A Subjectivist-Objectivist Approach

⁶ Fuhmann (2006).

⁷ Tracking the activities of involuntary groups does present an additional obstacle to this view, e.g., prison populations, conscripted military service personnel, etc. I would argue that capturing data on involuntary collectives or groups would constitute a “tainted epistemic case” due to potential conflict between the role of epistemic normativity and authority in belief-updating.

⁸ I do not want to backtrack too much by reiterating my claim that all naturalized epistemological practices are social epistemological practices – I only note this to emphasize that the richer “practices” in question in terms of the theory/practice divide will also be distinctly *social* practices.

Bayesians typically come in two stripes: objective Bayesians and subjective Bayesians. While subjective Bayesianism is better represented in the philosophical literature, objective models are far more favored by Bayesian statisticians, i.e., those actively engaging in Bayesian or Bayesian-style computations. The philosophical problem for objective Bayesians, the *epistemic* problem, is that, unlike subjective Bayesians, objectivists must provide and explain non-arbitrary, non-informative (or objective) prior probabilities – and this is no small feat. (I will say that weighing current models, e.g., EP-calibrated priors, reference priors, Jeffreys priors, etc., make very clear the difficulty of finding consensus on that particular front). The upshot is that objective priors help prevent irrationality and/or “epistemic taint” from being compounded upon belief-updating. The approach that I prefer will bring us back to the discussion of Lewis’ (1980) Principal Principle from Chapter 4. As Weisberg (2011) and other objectivists have noted, there is actually a continuum between subjective and objective Bayesian approaches such that any particular view can be further along (in one particular direction) on that continuum.⁹ As a subjectivist at heart, I argue for a subjectivist-objectivist Bayesian view that uses objective probabilities understood as *objective chance* as an initial normative constraint on prior probability assignment. That is, for an epistemic agent to be considered a “rational updater” they have to also consider objective facts about the world as constraints on their beliefs. You can see immediately that this is the same direction that Goldman (1999) sought to move in. Goldman’s (1999) view, however, tried to meld subjective credence assignment with the assumption that an epistemic agent would begin with fairly rational priors. There is just no good reason to believe that this would be the case. We can begin similarly by arguing that agents should be able to assign credences according to their current epistemic state, or by “their own lights,” and then require the satisfaction of certain normative requirements for those belief assignments to be considered rational. These normative requirements would look something like Lewis’ (1980) Principal Principle *in addition to* the principle of conditionalization and what I call, “General Non-Partitioned Reflection.” The former items will not be terribly radical for a subjectivist with objectivist leanings, as it aims to achieve an initial objective, synchronic constraint on belief assignment. The latter requirement, General Non-Partitioned Reflection, is my interpretation of what van Fraassen (1995) was trying to

⁹ Weisberg, Jonathan (2011). “Varieties of Bayesianism” in *Handbook of the History of Logic, vol. 10*, eds. Dov Gabbay, Stephan Hartmann. Amsterdam: North-Holland Publishing Company/Elsevier, pp. 477-552.

achieve with General and Special Reflection. This is also a normative *synchronic* constraint on belief assignment that aims to scope an agent's initial credences. A subjectivist approach in conjunction with these three principles is what will constitute the *selection function* for our pragmatist-inspired theory of truth:

- (i) The Principle of Conditionalization;
- (ii) The Principal Principle;
- (iii) General Non-Partitioned Reflection + R-Reflection

I will start by assuming that Lewis' (1980) Principal Principle will do a lot of the heavy lifting in terms of dealing with the objectivity of our initial credences. In the last chapter I argued that Goldman (1999) would do well to adopt this principle to deal with the problems inherent to his own view. I argued that he could (1) adopt the Principal Principle, or (2) advocate for the convergence thesis. What I did not suggest for Goldman (1999) – but do suggest for my own view – is that we ought to use (1) *in conjunction with other normative constraints* to achieve (2). One problem with relying on the Principal Principle exclusively is that we just don't know that many fine-grained "chancy" outcomes. But what if we used it in conjunction with something like General Non-Partitioned Reflection? What if we used it in conjunction with Dutch book arguments? Similar to what Wheeler and Williamson (2011) attempt to do with "Evidential Probability-Calibrated Objective Bayesianism," we can front-load the normative requirements for belief with what I call "objective scoping methods," such that convergence of belief is actually attainable. The Principal Principle is one scoping method. Can we say the same thing about Dutch books or Dutch strategies? What about Reflection?

5.2.2 Dutch Books, Diachronic Coherence, and Revising the Reflection Principle

First, I would argue, as van Fraassen (1984) does, that an epistemic agent who is vulnerable to Dutch *strategies* as opposed to Dutch books, is still having issues with the coherence of their initial credence assignment. Consider the following series of bets: Let H be the hypothesis, and E , the agent's "future attitude to the hypothesis"; (Bet 1) pays 1 if $(\sim H \& E)$ and costs $P(\sim H \& E)$, (Bet 2) pays x if $\sim E$ and costs $xP(\sim E)$, and (Bet 3) pays y if E and costs $yP(E)$.¹⁰ Now it may be useful to note the setup. In the

¹⁰ Van Fraassen (1984) notes: "Here the probability of $\sim E$ equals 1 minus the probability of E . The number x is the usual conditional probability of $\sim H$ given E ...finally y is x minus the subjective probability the customer will have for the hypothesis, when and if E becomes true" (240).

synchronic case, the Dutch *book*, the bets are offered simultaneously and the epistemic agent presumably sees the bets as fair. In the diachronic case, the Dutch *strategy*, the agent is again offered a set of bets as well as a *third* bet, offered later, which taken all together are regarded by the agent as fair. Van Fraassen (1984) takes it that we can assume that an agent who is vulnerable to a diachronic Dutch book “has an initially state of opinion or practice of changing his opinion, which together constitute a demonstrably bad guide to life”, and shows above all else that the agent is, at the outset, entertaining *incoherent* beliefs.¹¹ But why? What is the relationship between my initial credence and my future credence?

So let’s consider a bicycle-racing event that will take place tomorrow at 2 p.m. Let H = the proposition that Cyclist 2 will win the race. Let’s assume that I will consider taking a 2:1 bet on Cyclist 2 tomorrow at 2 p.m. as a fair bet, where $P(E) = 0.5$. Now the bookie asks me what my subjective probability will be for my future attitude that Cyclist 2 will lose. Let’s say that my subjective probability that $P(\sim H \& E) = 0.4$. Given that $x = 0.4/0.5 = 0.8$ and $y = x - 3/5 = 1/5$, the cost of the series of bets will be as follows: (Bet 1) will cost 0.4, (Bet 2) will cost $0.4 [x(1-P(E)) = (0.8)(0.5) = 0.4]$, and (Bet 3) will cost $1/10 [yP(E) = (1/5)(0.5) = 1/10]$. Note that the total cost of all of the bets $[x + yP(E)] = 0.8 + (1/10)$. Now, if I am *incorrect* about my future attitude, I will receive x , or 0.8, for a total win of 0.8. If I am *correct* about my future attitude, I will receive $1/5$, or 0.2 on my third bet plus the fair price for my bet against Cyclist 2 winning the race, $3/5$, or 0.6. Again, the total win equals 0.8. Now we can see my present predicament. Either way my \$1 bet will result in a \$0.02 net loss. Or as van Fraassen (1984) puts it: “It turns out that I shall have been the loser come what may.”¹²

What we are supposed to take from our vulnerability to a Dutch strategy is that my initial credence on the day before the race does not cohere in a rationally acceptable way to my credence on the day of the race. And this motivates van Fraassen (1984) to posit an additional synchronic constraint on my present beliefs – the *Reflection* principle. He claims: “To satisfy the principle, the agent’s present subjective probability for proposition A , on the supposition that his subjective probability for this proposition will equal r at some time later, must equal this same number r .”¹³ The broader version of this

¹¹ *Ibid.*

¹² Van Fraassen (1984), p. 241

¹³ Van Fraassen (1984), p. 244

is called the “General Reflection Principle,” and argues: “My current opinion about event E must lie in the range spanned by the possible opinions I may come to have about E at later time t, as far as my present opinion is concerned.”¹⁴ Or, rather, the agent’s present beliefs ought to correspond to a range of future beliefs that he *thinks* he might have.

First, van Fraassen’s (1984, 1995) Reflection principle is too strong a constraint on initial credence assignment.¹⁵ This is because we can show cases, i.e., Arntzenius’ (2003) Shangri-la case, where the epistemic agent violates Reflection and we do not find the agent’s beliefs to be irrational.¹⁶ However, we do still have reason to consider something *like* Reflection – a principle that preserves our intuitions about initial credence assignment and diachronic coherence – as an epistemic desideratum. That is, diachronic coherence may not be required for rationality but may still be considered epistemically desirable. This will work best under something like *ceteris paribus* conditions. The motivations behind van Fraassen’s Reflection principle seem to be rooted in having a *synchronic* constraint that is concerned with diachronic coherence. Rather, it is concerned with how we see our present selves as connected to our future selves. Recall that General Reflection says that we ought to have the beliefs now that we *think* that we will have in the future. Aside from exotic cases, this simply seems to suggest that *I* as the *subject* of belief ought to allow the (presumably) reasonable “me” of the future to constrain the “me” now. This seems rather benign. We can reasonably imagine ourselves as having a variety of very boring beliefs that we also think we will have in the future. If we did not, in some sense, think our two “selves” related, we could not explain why we think of ourselves as dynamic epistemic *subjects*. That is, we think ourselves capable of having beliefs both now and in the future; and we think that these beliefs (both “forward-looking” and “backward-looking”) belong somehow to a continuous, identifiable “*I*”.¹⁷ Need these beliefs cohere? Perhaps not necessarily, but we assume in most cases that they do. This

¹⁴ Van Fraassen (1995), p. 12

¹⁵ This is an argument also made by Weisberg (2007) in “Conditionalization, Reflection, and Self-Knowledge,” *Philosophical Studies*, 135: 179-197.

¹⁶ See: Arntzenius (2003)

¹⁷ I am borrowing the terms “forward-looking” and “backward-looking” coherence from Christensen (2000). Roughly, this means coherence with respect to the future and with respect to the past. Christensen denies that we can have a coherence principle that links the present and future as an epistemic requirement. I think that we can for the above-mentioned reason—we think of ourselves as the same epistemic subject over time. However, I will leave questions concerning why we can assume we are the same subject over time to those concerned with issues of personal identity. This may beg some fairly substantive questions against my claim here, but I’m not sure that this cannot be somehow circumvented.

assumption will be enough to suggest that a diachronic coherence principle is desirable. This is because, to emphasize the above point, we only need to say that the agent *thinks* this to be the case. In any case, assumptions are, of course, not arguments. So let me explain why I think that this assumption is well founded.

Most of our beliefs do not look to be of the sort considered by special and/or exotic cases, e.g., the Sleeping Beauty Problem, and the focus on exotic cases has made a hard wall to climb for synchronic constraints like Reflection. The best way to see why this is the case is to consider the example offered by Christensen (1991):

There is an unusual psychedelic drug, call it LSQ, with the property that those under its influence, while fairly normal in most respects, believe very strongly that they can fly. Suppose that our agent is quite sure that she has just swallowed a hefty dose of LSQ, and someone asks her, “What do you think the probability is that you’ll be able to fly in one hour, given that you’ll then take the probability that you can fly to be .99?” I take it as obvious that the answer mandated by Reflection (“.99, of course!”) is ridiculous.¹⁸

So what is the worry with this case? Christensen’s (1991) response misrepresents what exactly the agent will believe about LSQ (although he does represent the belief as understood by van Fraassen’s (1984) Reflection perfectly). This is because the question posed to the drug-taker is (not to be too legalistic) a *leading* question. This is because most likely the agent will not be able to reasonably foresee his or her future belief about the probability of flying, even given that he or she knows at *present* certain facts about the drug. Rather, the most plausible response is that in the present the agent’s future beliefs about LSQ and flying will be that he or she is unsure about what she will believe in the future. In the future I may believe that I can fly due to my taking LSQ. However, this does not mean that in the present, even if I was explicitly told that I would believe this, that I will believe *in the present* that this is a foreseeable option for me. At best, I will probably *believe that* in the future I will be under the influence of a drug that makes me believe I can fly. I take it that the moral here is that we can easily show violations of Reflection if not enough weight is being put on what is meant by the “foreseeable” future. Consider an abridgment of the previous example, where instead of believing that I can fly, I will believe that I am a staunch Republican

¹⁸ See: Christensen (1991) for a more explicit discussion of this example.

(instead of, let's say, a bleeding-heart Democrat). You may tell me that my taking LSQ will result in this belief, but this does not mean that I will take you seriously. Most likely, I will believe that *in the future* I will be under the influence of a drug that makes me have the corresponding belief, i.e., that I am a staunch Republican. But Reflection dictates that what is foreseeable for me should be determined by the beliefs that I hold at present. And an emphasis on this will yield a different result than in the original LSQ example.

The problem with exotic cases is that it assumes that I can foresee myself as capable of believing just about anything in the future. However, if that were the case, then *any* normative requirement for diachronic coherence would seem quite useless. A more sympathetic reading of Reflection would insist that it is my *present* beliefs about my future foreseeable beliefs that can account more reasonably for what I think myself capable of believing. This does not beg the question against diachronic coherence so much as restate one of the more intuitive elements of van Fraassen's (1995) argument for Reflection: We think of ourselves as rational epistemic agents, and we can, to some extent, foresee ourselves, according to our own subjective position, as being reasonable in the future. It seems to me that that the moral of cases like the LSQ case merely reinforce the idea that in the realm of the foreseeable, not everything will be up for grabs. However, even if we throw out cases that seem to exploit the notion of the foreseeable future, can we still make a case for Reflection? We can make a case if we take foreseeable future belief as a set of beliefs, or a bundle. I will, for the moment, call this a "*Reflection-like*" principle. Consider the following example from Maher (1992): Currently you believe that after 10 drinks you will not be able to drive safely home. However, you also believe that after 10 drinks you *will* believe that you can drive safely home. Reflection clearly tells us that I ought to believe now what I think I will come to believe in the future. Thus I am *wrong* to believe that after 10 drinks I cannot drive safely home.¹⁹ Of course, this would be extremely unwise. The problem is that this reading of Reflection is extremely misleading. It implies that my beliefs work on the world in a 1-to-1 relation as opposed to in "belief-bundles." Even if I consider that in the future I will come to have probability 1 in the belief, "I am okay to drive after 10 drinks," it seems irrational to believe it now (as *now* I believe the opposite and would have to say that I am

¹⁹ See Maher (1992), a similar example also appears in Christensen (1991)

currently mistaken about my present belief). This irrationality stems from the idea that my future possible beliefs must match up directly to some exclusive evidential scenario in the world.

But what if my beliefs were a package deal? Although it may appear that I will have a high probability in the future belief, “I am okay to drive after 10 drinks”, this belief will have a non-zero *conditional* probability and will be subsumed under a more general belief about my future state. Say, the belief in the proposition: “My beliefs after 10 drinks are entirely untrustworthy.” Given that I assign higher probability to the *general* belief, I will probably accept a particular belief, *r*, even though at present it seems contradictory for me to do so. That is because my particular belief, *r*, is *inert* (it has a very low, non-zero probability conditional on my acceptance of the general belief in the future). Thus my general belief that in the future I will believe myself untrustworthy does constrain my present beliefs – and rightfully so. My future possible beliefs work in conjunction with what I think, in general, the world will be like. I assume that the world will be *like* a place where I cannot represent my beliefs accurately, i.e., presumably because I am intoxicated. Thus the various *de dicto* temporal descriptions (which will include any relevant subdivisions of time, but not exhaustive or exclusive at any *t*) can be *bundled* into *general* beliefs about what the world at a particular *t* will be like. For example, consider any foreseeable future beliefs that I may have about myself after 10 drinks (e.g., “I cannot drive home safely,” “My motor skills are less exact”). These particular beliefs may serve as evidence in many possible worlds, even those worlds where I am not at a dinner party and/or not presently worried about my alcohol consumption. These particular beliefs will have a very low, non-zero probability at *any* possible world where I think that I am, i.e., any possible world of which I think I will be a member.²⁰ In most cases, they will be irrelevant (or *inert*) beliefs that I have *about* the possible worlds where I am drinking and/or I am worried about my driving or my motor skills. However, if at some particular *t*, I have a relevant general belief – say, that “My beliefs after 10 drinks are untrustworthy” – then the probability of the particular belief will increase *conditional* on my acceptance of the more general belief. This is what ought to constrain what I believe at present. The mistake is to assume that we ought to pick out one particular belief, *r*, at a time. Here is a very natural point to jump ship from van Fraassen’s (1995) Reflection principle. This is because it will

²⁰ There may be a worry here of having to evaluate general beliefs in irreducibly *de se* terms. I’m not quite sure as of yet how to handle this, so I will put it off and assume that there is some way to resolve this issue.

now be entirely untenable to say that our evidential scenarios are mutually exclusive with regards to future possible worlds. (It might be important to note that Weisberg offers a very good, intuitive argument against partitioning which I will simply footnote here).²¹

So I have made the following assumption and following arguments about rational agents: (1) we think of ourselves as a continuous, or *dynamic*, epistemic subjects. This means that we already do make some assumptions with regards to forward-looking and backward-looking coherence. Thus, we should have a principle that at least captures these assumptions as epistemically desirable; (2) we should not discount the possibility of “*Reflection-like*” principles by exploiting the idea of what will count as a foreseeable future belief via exotic cases. That is, any foreseeable future belief ought to, from the agent’s own position, actually *be* foreseeable; and, (3) we should consider the individual beliefs we are “reflective of” as grouped together and subsumed under a **general belief** that is representative of what we think the world will be like at some given *t*. A **particular belief**, *r*, isolated from a general belief will merely give us a *de dicto* temporal description that can potentially serve as evidence in many possible worlds. Thus it is not informative (we accept it as having a very low probability) at any particular world *unless* a relevant general belief is also accepted. If we have no reason to assign a higher probability to the general belief, i.e., to accept it, then we can say of any particular, *r*, that it is a belief of which we ought not be “reflective.”

5.2.3 General Non-Partitioned Reflection

For the moment, I will call this “Reflection-like” principle **General Non-Partitioned Reflection (GNR)**. Where our future foreseeable degree of belief in *r* should match our current degree of belief in *r*, where *r* constitutes non-partitioned *de dicto* temporal descriptions of the world. These temporal descriptions of the world should be considered as broad pieces in an agent’s holistic belief state. They are, more or less, the beliefs that an agent takes to be true of the world regardless of the epistemic state

²¹ Weisberg (2007) makes the following argument against partitioning: “Here’s my guess: the partitioning assumption results from a confusion between the epistemic paths an agent may take and the information she learns on those paths. While the possible histories I may encounter between now and *t* do form a partition of the space of possibilities, the information that I may glean along those histories needn’t form a partition. To make this point vivid, we can visualize an agent’s epistemic history as ticker-tape, where each cell of the tape corresponds to a time and contains information the agent conditionalizes on at that time...Now, the set of possible tapes certainly forms a partition, since an agent must undergo exactly one tape. But the contents of the tapes—the conjunction of cell-contents—needn’t obviously form a partition, for the reasons already given” (184).

of any general belief, p , or any particular “defeater” or “inert” belief, r . This is where a “Reflection-like” principle will be more useful than van Fraassen’s (1995) Reflection. For example, if we can imagine ourselves believing, “Any U.S. state capital will be a city within the federal republic of the United States,” Reflection requires that we believe this at present as well. However, what happens in the case where we can also, and simultaneously, imagine ourselves believing: “I’m unaware of the territorial status and boundaries of the component parts of the United States”? How do we follow Reflection in cases where a more broadly defined belief may or may not invalidate a more specific belief? Furthermore, although we are luminous (See: Williamson (2000)) to both beliefs, we may be unaware of *if and how* that broad belief we are holding invalidates a more specific belief that we also may hold. In which case, Reflection implies coherence in terms of an agent’s holistic belief states where coherence may not hold. This would only be diachronic coherence *by decree*. Problems with Reflection arise from how we frame the beliefs about which we ought to be “reflective.” Although I do not want to suggest that we can simply convert every case where we violate Reflection into a different case, i.e., just create a general belief under which we subsume any particular “violator” belief – this would be too *ad hoc*. What we need is a principle that does not suffer from the problems that van Fraassen’s (1995) Reflection principle suffers from, and specifically, does not make unnecessarily narrow assumptions about how epistemic agents might *order beliefs*.

For example, in the case described above, the general belief, “My beliefs after 10 drinks are untrustworthy,” will hold at any **centered world** (i.e., *centered* world, designated by a world, agent, and time) even if I also currently hold the potential defeater belief(s), “I do not drink,” or “I do not like drinking.” However, I argue this interpretation of “belief-bundling” would require an additional account of **belief-ordering** – taking particular beliefs, or defeater and inert beliefs, as classes of “higher-order beliefs,” where higher-order beliefs describe beliefs that either have (a) *inert* beliefs, or beliefs which have a very low, non-zero *conditional* probability at non-centered worlds, and (b) *defeater* beliefs, where an agent may hold both p and *not-p* as *de dicto* temporal descriptions at any given *centered* world. (E.g., “Obamacare is a federal scam on American taxpayers,” and “The Affordable Care Act is federal relief for American taxpayers”). This interpretation would also require a non-arbitrary way to operationalize the relationship between general and particular beliefs such that a particular belief can provide probabilistically significant evidence at *centered worlds* and remain virtually inert at non-centered worlds.

Here I think it's useful to think of this "operation" in terms of an evaluation of an undirected graph, where the nodes are both general and particular (defeater, inert) beliefs and the edges are the probabilistic relations between beliefs. The edges denote belief strength (i.e., degree of belief) by their respective length and/or proximity to other nodes, and the relationship *between* unconnected nodes denotes beliefs at non-centered worlds. Although we can quickly see that neither representation will help us with the difference between belief *types*. Two general beliefs that are only minimally related in terms of an agent's holistic belief state may look similar to general beliefs that are only very minimally related to particular beliefs. For example, consider the potential nodes representing general beliefs:

- "The city of Denver sits at a rather high-elevation."
- "Orange juice is beneficial when you have a mild cold."

Graphically, this may look very similar to the relationship between a general belief and a particular (inert, non-zero probability) belief that serves as evidence at a centered world – assuming the graph represents an agent's holistic belief state at both centered and non-centered worlds. For example, consider the following potential nodes representing a general belief and a particular belief:

- "My beliefs after 10 drinks are untrustworthy."
- "I cannot drive home safely."

But how would we represent the increase in probability of the *second* belief, "I cannot drive home safely," conditional on an agent having the first belief? In this case, the graph helps us to visualize the problem described above: We need a subjective principle *in addition to (GNR)* that gives an agent reason to increase the probability of a given particular belief, r , conditional on having the appropriate general belief, p . This will protect our new Reflection-like principle, (GNR), against exotic cases like Christensen's (1991) LSQ case as well as against the problem that our beliefs about the world seem not to be as fine-grained as van Fraassen's (1995) Reflection principle suggests. So what would this additional principle look like? And how would we apply it in cases where we are also employing (GNR)?

5.2.4 Scoping (GNR): Defeater Beliefs, Inert Beliefs, and "R-Reflection"

Like Reflection, (GNR) is a normative requirement for belief updating and one of the three objective scoping methods for our subjectively determined beliefs. Unlike van Fraassen's (1995) case for Reflection, I will not attempt to argue here that (GNR) is in any way entailed by the Principle of

Conditionalization. Reflection's most meaningful role is its normative appeal for agent's interested in diachronic coherence, and that the same can be said for (GNR). Unfortunately, (GNR) alone isn't enough to address the issue of higher-order beliefs (or "belief bundles") because it only requires the use of the Reflection principle from the perspective of non-partitioned r beliefs, or "defeater" or "inert" beliefs. For this reason, I propose the adoption of (GNR) be coupled with the adoption of a further normative principle, "R-Reflection," that asks an agent to be "reflective" about *employing* Reflection. Ultimately, I presented an explanation of R-Reflection in the section above (Section 5.2.2), but not the *justification* for it. **R-Reflection** can be formulated as follows: Any foreseeable future belief ought to, from the agent's own position, actually *be* foreseeable. We should consider the beliefs we are "reflective of" as grouped together and subsumed under a **general belief** that is representative of what we think the world would be like at some given t . A **particular belief**, r , isolated from a general belief will merely give us a *de dicto* temporal description that can potentially serve as evidence in many possible worlds. This belief is not informative (we accept it as having a very low probability) at any particular world *unless* a relevant general belief is also accepted. That is, if I believe that having taken LSQ will make me believe that I can fly in the near future (approx. 1 hour), then my inert belief r = "I will believe that I can fly," while my general belief p = "I will be under the influence of a drug that makes me believe I can fly," or even more broadly, p = "I will not be able to represent my beliefs about the world accurately in one hour." The justification for R-Reflection follows from accepting that it is possible to have general beliefs that are, so to speak, "unhinged" from particular beliefs, and can be ordered or arranged into different belief types. In some cases, like the LSQ case, the particular belief will serve almost no evidential purpose at t if the general belief, p , is accepted. This only works because the beliefs are not necessarily linked in an agent's holistic belief state. Due to this, (GNR) coupled with R-Reflection should be sufficient to avoid the problems sketched by Christensen (1991) and Maher (1992) against Reflection in addition to avoiding the unnecessary complication of trying to partition foreseeable future beliefs. Along with the Principle of Conditionalization and Lewis' (1980) Principal Principle, these two additional principles will serve as the objective scoping methods for our Bayesian *selection function*.

5.3 A Note on Bayesian Aggregation

If the Principle of Conditionalization, the Principal Principle, (GNR), and R-Reflection are our objective scoping methods and the combination of those methods constitutes our selection function, we are still left with the issue of how to handle these methods as applied to groups and/or collective doxastic agents. Bayesian methods in epistemology are generally applied to individuals, not groups. So how can our selection function now work for *social* epistemology? In the last chapter, I considered the aggregation methods used by Goldman (1999) in his account of veritistic social epistemology. His account prioritized weighted voting methods as the veritistically best social practice in the case concerning the weather bureau, but he did not suggest that there were any universally applicable social practices (or Bayesian aggregation methods) to increase collective credence for all cases. Goldman (1999) is right to move away from suggesting such a social practice or method is uniquely possible. As List and Dietrich (2014) argue in their work on probabilistic judgment aggregation, even in the case of weighted voting (or weighted linear voting), we can assign many different probabilistic “pooling functions” dependent on “the context and the intended status of the collective opinions.”²² Different intentions – epistemic and otherwise – dictate how we go about selecting a pooling function (or probabilistic aggregation function). They write:

At least three questions are relevant:

- Should the collective opinions represent a *compromise* or a *consensus*? In the first case, each individual may keep his or her own personal opinions and adopt the collective opinions only hypothetically when representing the group or acting on behalf of it. In the second case, all individuals are supposed to take on the collective opinions as their own, so that the aggregation process can be viewed as a consensus formation process.
- Should the collective opinions be justified on *epistemic* or *procedural* grounds? In the first case, the pooling function should generate collective opinions that are epistemically well-justified: they should ‘reflect the relevant evidence’ or ‘track the truth,’ for example. In the second case, the collective opinions should be a fair representation of the individual opinions. The contrast between the two approaches becomes apparent when different individuals have different levels of competence, so that some individuals’ opinions are more reliable than others.’ The epistemic

²² List, Christian and Franz Dietrich (2014). “Probabilistic Opinion Pooling” in *Oxford Handbook of Probability and Philosophy* (forthcoming).

approach then suggests that the collective opinions should depend primarily on the opinions of the more competent individuals, while the procedural approach might require that all individuals be given equal weight.

- Are the individuals' opinions based only on *shared information* or also on *private information*?

This, in turn, may depend on whether the group has deliberated about the subject matter before opinions are aggregated. Group deliberation may influence individual opinions as the individuals learn new information and become aware of new aspects of the issue. It may help to remove interpersonal asymmetries in information and awareness.²³

In the first case, our use of “collective opinion” or aggregated judgment will depend on whether the “group decision” will override individual belief and/or opinion. We can surely imagine a scenario in which the aggregated decision *qua decision* will only come into play if the individual is representing the group and not him or herself. Alternatively, we can also imagine a scenario in which the individual subsumes his or her belief beneath the collective opinion and adopts that opinion as his or her own. As List and Dietrich (2014) point out, there is some ambiguity in how we employ collective opinion making, and I would further suggest that this ambiguity can only be resolved on a case-by-case, contextual basis. What the collective opinion is *for* is actually a meta-epistemological issue to be determined prior to invoking that decision. I would also argue that how individuals form their beliefs in relation to the activities of the group, as in the third case listed above, is also a meta-epistemological issue, but not for the same reasons. Exactly how an individual can or cannot be influenced by the activity of collective opinion making relies in large part on the numerous ways in which our cognitive processes can be manipulated – including manipulated by reducing our prior biases and prejudices. This is an issue for epistemology broadly speaking, but only insofar as epistemology (social or traditional) can pull conclusively from successful research in psychology, neuroscience, and cognitive science. But this is not within the scope of this project. Instead I am primarily concerned with the second case, or how we ought to justify our collective opinions. I would argue however that our procedural concerns are not distinct from our epistemic concerns, unless we have some prior notion of what constitutes a “competent” individual. Aside from being able to analytically negotiate the relevant information, and engage in what I would *very tentatively* call “normal cognitive

²³ List (2014), p. 2.

processing,” all subjects should be given equal weight in our selected aggregation method. Thus I take it that what we need is an aggregation method that is justified on broadly epistemic grounds. As List and Dietrich (2014) note, the hope is that our collective opinions will be well-justified or will “track the truth,” and my hope is that they track the truth in the same way that our individual beliefs do (via the Bayesian methods outlined above). But this is not to say that we should hope to arrive at a *unique* method that will allow us to collectively track the truth. And I don’t want to suggest a veritistic social practice that allegedly produces an increase in collective credence without also suggesting that there is a systematic way of subjecting the method itself to epistemic evaluation. That is, if the method isn’t unique, then we need a principled way to decide between methods that we think can collectively track truth. Additionally – as good epistemic agents should be good Bayesians – I also think that whatever methods we do decide on for collective truth-tracking should be broadly Bayesian methods. In the former case, a principled theory of “epistemic disagreement” will handle the issues of context and intent laid out by List and Dietrich (2014) above. I argue that this is one moderately successful, non-dogmatic way of mediating the meta-epistemological issue of deciding what the method is for. I will come back to this in further detail in the next section, 5.4. In the latter case, I think it’s OK to use whatever methods of aggregating Bayesian judgments are most useful for our *epistemic* purposes, e.g., Goldman’s (1999) use of weighted voting schemes. As List and Dietrich (2014) argue, and I agree, it’s not the function itself that matters most, *it’s the purpose of the function* that matters most. And it is in this sense that we can all be “good Bayesians.”

5.4 Epistemic Disagreement and Epistemic Peers

David Christensen (2007) claims that epistemic disagreement or peer disagreement fills a critical gap in how we arrive at and justify our own beliefs. He writes: “We live in states of epistemic imperfection because we do not always respond to the evidence we have in the best way. Given that our epistemic condition consists in imperfect responses to incomplete evidence, part of being rational involves taking account of these sources of imperfection.”²⁴ The good news, he claims, is that epistemic disagreement can help us improve upon our own epistemic imperfections. It’s fairly apparent that engaging in epistemic disagreement, or *doxastic* disagreement, can beneficially challenge our current credences, and, in some cases, lead us to revise upon them as well. For example, if I am having a disagreement with a friend over

²⁴ Christensen, David. “Epistemology of Disagreement: The Good News,” in *Philosophical Review*, 116 (2007), 187-217, p.187.

whether or not dogs make good house pets, I might defer to her belief that they do not make good house pets because I know that she is afraid of dogs and also because I am interested in being (or appearing to be) a sympathetic friend. Assuming we both provided two distinct evidential scenarios (e.g., cases where dogs bite, and cases where they do not), my belief revision, or credence change, was not made by our coming to some rational consensus about canine behavior, or by our evaluating the same evidence (or evidential scenarios), it was made out of deference to a friend that I thought beneficial to the relationship at the time. That is, my reasons for credence change seem decidedly non-epistemic. Now we might say this example only distinguishes epistemic disagreement from other types of disagreement, or even other types of socio-behavioral practices, but that would be incorrect. Epistemic disagreement constitutes any type of disagreement between epistemic agents where what is *at stake* in the disagreement is something that “*S believes*,” whether or not the content of that belief is veritistic.

Improving the *veritistic* imperfections of our beliefs via epistemic disagreement can only happen if our method of epistemic disagreement is principled in some meaningful way as to avoid non-veritistic credence changes as a result of epistemic disagreement. This also might aid us in circumventing those cases where our change in credence is based solely on what might be beneficial to us for non-veritistic reasons. But how are we to go about this “epistemic self-improvement” via a principled epistemic disagreement?²⁵ First, it’s important to note that although epistemic peers may engage in rational disagreement (say, involving the same evidential scenario) there is no reason to think that they *will* come to some conclusion about credence revision for either party. There can – *and often will be* – something like epistemic stalemate where neither party is swayed by the other’s interpretation of the same evidence or where one party will not budge from employing non-veritistic reasons. But it’s fair to remain unconcerned about those cases and to focus on cases where credence change is more likely, or are, as Christensen (2007) notes, cases where credence changes is actually “called for.”²⁶ Let’s consider a case where two friends are on a road trip and have accidentally become lost. Both individuals consult the same road map, and have the same background knowledge of the geographical area, but they do not reach the same conclusion about which local roads to take to find the main highway in the fastest amount

²⁵ I borrow the term “epistemic self-improvement” from Christensen (2007). He uses it to cover both our credence improvements via additional (good) evidence as well as our improved credal attitudes, or how we can improve our *responses* to the evidence available to us, p.187.

²⁶ Christensen (2007), p. 189.

of time. Let's further complicate this scenario by disallowing what Christensen (2007) calls a "live-and-let-live" attitude and assume that there is something at stake for these two individuals other than their selecting the fastest route to the main highway.²⁷ For example, perhaps the time of arrival is of some serious importance. (This will also help to prevent the scenario I sketched above where a credence change is made for non-*veritistic* reasons.) Now what should these two epistemic peers do? Let's assume that one friend has a .45 degree of belief in the proposition, "Taking Evergreen Road to the highway will be faster," while the other friend has a .65 degree of belief in the proposition, "Taking Quarry Road to the highway will be faster." Notice that I have taken for granted that the two friends are in fact *epistemic peers* – that is, they regard each other as equally capable of adequately evaluating and analyzing the evidence before them. If that is the case, it would be reasonable for them to both employ something like the Principle of Symmetry and raise and/or lower their current credences accordingly, and, for this particular example, to conclude that either route would be equally successful (via a .50 credence in both alternatives).

If the two friends are not epistemic peers – and they are not epistemic peers on *veritistic* grounds – then a different principle of resolution will have to be employed. In this case, for example, the case where one friend has (a) evidence that is unavailable to the other friend, or (b) analytical or evaluative skills or capabilities that the other friend does not possess, then we have a case of epistemic *superiors* and epistemic *inferiors* with regard to the case at hand. The importance of epistemic superiority and epistemic inferiority lies only in the fact that an epistemic agent is in possession of more (or less) accurate evidence in relation to a given *P*. We do not want to assume that a case of epistemic superiority or inferiority would constitute a blanket case of "deferral to epistemic authority," or even a normatively mandated change in one's current credences. This is because, as mentioned above, we cannot discount that there are often blurry lines between *veritistic* and non-*veritistic* grounds for believing in a given *P* and that there are no hard-and-fast rules for determining which beliefs are operating in any given case. It would be easy to imagine in a case where an epistemic agent is in possession of more accurate evidence related to a given *P*, but opts to base a belief in that *P* on non-*veritistic* grounds. In this case, it will do no good for the other epistemic agent to change their credence according to that peer, or to "defer to

²⁷ Christensen (2007) takes the idea of a "live-and-let-live" attitude in epistemic disagreement from the work of Adam Elga, p. 190.

epistemic authority,” when we cannot rule out such cases. We can also see how easily this line of thought leads us back into making traditional epistemological assumptions about objectivity and belief-formation! Thus let’s assume that there will be a degree of subjectivity at work here at all times – which is par for the course when negotiating individual beliefs. Let’s also assume that epistemic agents engaging in disagreement are also interested in generating what Goldman (1999) would call the “best veritistic outcome” and have some minimal commitment to truth-seeking.

So what do we do in the case of epistemic superiors and inferiors? In much of the literature on epistemic disagreement (e.g., Elga (2005, 2007), Christensen (2007), Frances (2012)), it is generally agreed upon that in the case of epistemic peers, and in the case of epistemic superiors/inferiors, disagreement should lead us to reevaluate and in many cases readjust our current credences. Namely, there is something about disagreement with reasonable peers that should make us *less* confident in our own beliefs regardless of whether or not we ultimately readjust our current credences. But we cannot deal with the case of epistemic superiors/inferiors in the same way that we dealt with epistemic peers. The Principle of Symmetry would be markedly less useful in the case of epistemic superiors/inferiors because we do not have *evidential symmetry*. A more nuanced approach will be necessary. The approach described above as “deferral to epistemic authority” is also too blunt of a tool to handle cases where (a) an epistemic superior employs non-veritistic reasons for believing *P*, or in cases where (b) an epistemic superior’s additional evidence toward *P* is not accepted by the other party. In these cases, it is not even clear that we can successfully give more weight to the epistemic superior’s credences without saying something substantive about the quality of the evidence itself –and given the social epistemic framework that I’ve argued for here – this would be an impossible thing to do. There is no “bird’s-eye-view” with regard to evidence. There is no unchallengeable point of complete objectivity from which we can say anything conclusive about the *epistemic content* of an agent’s evidence for a given *P*. Our principle of epistemic disagreement for epistemic peers, e.g., the Principle of Symmetry, is clear but not a guiding principle, and I think any principle for non-epistemic peers (i.e., epistemic superiors, etc.) would beg the question against the objectivity of a presumed “epistemic superior’s” belief content. To move forward with our initial concern of having a principled way of deciding between broadly Bayesian methods, we need to be a bit more comfortable with the circularity in our own evaluative processes.

Epistemic disagreement isn't a "once-around" method, it will need to look closer to how reflective equilibrium is used in moral philosophy and applied ethics, i.e., we will have to make some intuition-based assumptions to start. There ought to be no "epistemic deference" because disagreement between epistemic agents should be a constant negotiation between our intuitions about particular cases, some agreed upon epistemic guidelines (or principles) and our approach to evidence and evidence quality. How we weigh the evidence of an epistemic superior or how we proceed from the evidential symmetry of epistemic peers should be the result of fruitful disagreement, implementation of principles and/or guidelines, and evaluation of evidence, not of a "principled" renouncing of our own epistemic autonomy.²⁸ And while this is beyond the scope of this project, constructing a method of reflective equilibrium for epistemic disagreement is the best way I can see to avoid these problems with epistemic superiority and inferiority. In the meantime, the Principle of Symmetry coupled with a *critical stance* toward views of non-epistemic peers seems a good, moderate way to proceed in terms of methods of epistemic disagreement.

5.5 What's Next?: Mapping The Epistemic Terrain

In this chapter, I argued for a form of veritistic social epistemology that fits broadly within a pragmatist framework. I also argued for revised belief updating principles to be employed within that framework coupled with what I consider to be a reasonable stance toward epistemic disagreement – one that allows for reflective equilibrium amongst disagreeing epistemic peers and/or epistemic superiors and inferiors – as the basic foundation for a pragmatist-inspired veritistic social epistemology. There are many things left to say and many details for which I must give an I.O.U. in order to proceed with the conclusion of this project. Additionally, the variety of veritistic social epistemology for which I've argued depends heavily upon the justification of epistemic communities and the ability of those communities to successfully network together, and this feature demands both the theoretical justification that I have provided along with an explanation of its practical application. This social epistemological project cannot be a solely theoretical agenda. Thus, for this project to have real-world weight, there needs to be some account of how epistemic communities can track the work of other epistemic communities along with the broader networks of which they are members.

²⁸ I borrow the term "epistemic autonomy" from Coady (2012) and with many shared concerns about relinquishing our own epistemic authority to those of an expert or epistemic superior.

Goldman (2010) introduced a type of systems-level epistemology that he calls “systems-oriented social epistemology,” that attempts to bridge this gap by applying the “social” part of social epistemology to the activities of formal institutions. As with a good portion of this project, I use Goldman’s (2010) seminal work in social epistemology as a jumping off point for my own analysis and this will also be the case for my final chapter. In the final chapter, I briefly explain Goldman’s (2010) view and propose moving beyond it. That is, I argue that we need to stop thinking of social epistemology as solely the “sanctioned” knowledge-producing activities of formal institutions (courts of law, research universities, think-tanks, governmental entities, etc.). We need to also track and examine the “unsanctioned” knowledge-producing activities of marginalized epistemic communities and the *informal institutions* that may exist within those communities. We need to think less rigidly about the external construct, or social and cultural “sanctioning” of communities of knowledge-producers by existing institutions operating as voices of authority, e.g., research institutions, governmental institutions, etc., and think more broadly about how we can capture, or *truth-track*, the activities of both formal and informal institutions. Although varying communities may not share a single systems-level language or shared set of data points – or even *community-level* language or a shared set of data points – the goal will still be to track their activities via some recognized system of outputs, such as markers of produced written work (peer-reviewed or not), i.e., citations. The secondary goal will be to map those outputs and establish *relations* to other institutions doing similar work to create relational networks. Here I appeal to current work being done in data visualization, scientific citation studies, and information science. Here I also relinquish the broader terms of “community” and “system” to focus on the more fine-grained terms of *formal institution*, *informal institution*, and *network*. I do this in order to make the next chapter as clear and practically oriented as possible: “communities” and “systems,” broadly speaking, do not so obviously produce publications and/or citations. The institutions within communities and within systems produce such citations and thus are the preferred level of analysis for “knowledge outputs.” It is this work that the social epistemological view that I advance here aims to track, map, and network.

Chapter VI

A Network-Level Epistemology: Representing Large-Scale Data

6.1 Systems-Oriented Social Epistemology and the Problem of Scale

There is a lingering issue that this project may suffer from a problem of scale. That is, whatever scale of social analysis we're considering at present, be it testimonial relationships between two epistemic subjects, or the social workings of large-scale social institutions, we must inevitably continue to up the ante, so to speak, in terms of how we interpret the veritistic project outlined in Chapter 5. This criticism is unavoidable. It is the unavoidable result of attempting to overlay a truth-oriented epistemological process on the social landscape in which we live. Ultimately, the goal is to at least convincingly suggest that such overlaying is possible from our more fine-grained social interactions, e.g., relying on testimony, to our more large-scale social interactions, e.g., the activities of research institutions, courts of law, etc. However, it is yet unclear whether there is a natural end to how broadly we can slice up our social world or how minutely we can slice up our social interactions. At the upper limit, this problem is particularly acute due to the ubiquity of web-based software that allows our social interactions to circumvent the mundane issues of time, space, and sheer number of inquirers who can meaningfully participate in a shared (digital) space. For this reason, at the upper limit, we have ask if we are just to assume that our units of social of analyses would naturally fluctuate upwards for the unforeseeable future?¹ What if we further categorize our social units of analysis by whether they are formal or informal institutions? That is, what if we start to breakdown just what it is that we *mean* by these upper units of social analysis? This could mean the difference between the potentially manageable difficulties of evaluating large-scale inter-governmental political organizations, such as the G8, and the potentially unmanageable problem of evaluating more amorphous social units like non-governmental communities of hackers. This might also mean rethinking the critique made in the last chapter that argues that Goldman's (2010) almost exclusive focus on "formal institutions" is too rigid.

Goldman (2010) suggests a "systems-oriented" approach to social epistemology which evaluations formal institutions on an *ad hoc* basis. There may be good reason for this. If we start with

¹ At the lower limit, there is perhaps a case to be made about the role of mirror neurons and shared affective stances, but this is neuroscientific terrain is well beyond the scope of my project here. Suffice to say, that we can also probably assume that there will continue to be fluctuations at the lower limit as well.

what we know, and then move further along the chain of “social evidence,” won’t we eventually end up analyzing informal institutions because they are systems of greater complexity? That is, won’t we eventually move from analyses of fundamental social relationships (e.g., testimony) to analyses of formal institutions (e.g., courts of law) and eventually move toward analyses of informal institutions (e.g., communities of hackers) as well? To his credit, Goldman (2010) *does* discuss everything from testimonial relationship, to the activities of juries, to digital spaces for inquirers, like Wikipedia. He also evaluates these social spaces according to their ability to maximize a given utility for both individual doxastic agents and collective doxastic agents, and much like his general view of veritistic social epistemology, relies on community-derived selection methods to maximize given utility. There are, of course, better methods than others; and, although the determination of “best methods” may be relative to the epistemic norms of a given culture or locale, this looks like a sound start for what will ultimately be a very epistemological endeavor.² So does Goldman (2010) really warrant the critique I leveled against his “systems-oriented” view at the end of Chapter 5?

The short answer is yes. The longer answer is that Goldman (2010) does not appear have any interest in validating informal institutions at all. Particularly, informal institutions that do not, and perhaps even *cannot*, have the checks-and-balances that are seemingly inherent to the workings of more formal institutions (e.g., editorial staffs, boards of directors, peer-review committees). For example, in the above-mentioned Wikipedia case, he argues, following Sunstein (2006), that what initially appears to be a reasonable example of “public knowledge” as generated by a “community of inquirers,” is more likely the result of a few dictatorial Wikipedia entry editors than it is of multiple, independent inquirers trying to justify the same bit of knowledge. Furthermore, he argues, that because successful communities of inquirers, or the successful application of the “wisdom of crowds,” involves the use of aggregation by averaging or majority voting, Wikipedia fails to qualify on those grounds:

The examples of the wisdom of crowds involve aggregation, i.e., either averaging or taking a majority vote of the independent viewpoints. Wikipedia entries, by contrast, are rarely determined in this fashion. Entries are usually edited by single individuals and the form of an entry at each

² Goldman (2010), p. 18. Pagination from *pre-print*.

moment is a function of whoever was the last person to edit it before you looked at it. The last editor can therefore be a self-appointed dictator (Sunstein 2006: 158).³

There are two things here that I take immediate issue with: (1) it is not clear that the “wisdom of crowds,” in the sense described by Goldman (2010) above as “averaging” or “majority voting,” is anything more than brute consensus unless further conditions are detailed, and (2) it is not clear that the criticism that Goldman (2010) makes against Wikipedia paints an accurate picture of how “wisdom of the crowds” can be used to achieve something over and above brute consensus. That is, it might be that *iterative consensus* is the underlying epistemic feature of this web-based platform. For example, in the description of the above Wikipedia dictator, it may be the case that consensus can be reached by multiple, diverse, and independent inquirers by way of many revisions over a long period of time. The community works asynchronously so that any particular dictatorial revision is only “dictatorial” in the most meaningless sense. Moreover, the number of editors (or “dictators”) who do and/or possibly will work on any one entry cannot be determined - perhaps estimated by a separate probabilistic analysis – out the outset of inquiry. So that any particular Wikipedia entry has a few dictatorial entries does not mean that the entire enterprise of Wikipedia as an example of “public knowledge” has hopelessly failed.

Goldman (2010) extends this criticism also to the workings of the blogosphere, and particularly to those elements within the blogosphere that look to supplant well-recognized formal institutions like print journalism. To this, Goldman (2010) claims:

When there are no longer conventional journalistic enterprises, which hire reporters to investigate matters that require months of research, who will undertake this investigation?...The matter might be formulated in terms of the epistemological metaphor of “foundations” of knowledge...Unless we are content to let bloggers fabricate whatever comes into their heads, we need initial stories to be based on first-hand observation, or searching interviews with people who have observed the relevant incidents (or ongoing practices) first-hand...Traditionally, investigative reporters are the people paid and trained to unearth such facts...How would bloggers serve this function? So it is doubtful that the blogosphere, qua social system, can adequately replace the traditional media in terms of epistemic outcomes. The blogosphere “free-rides” on the conventional media by picking

³ Goldman (2010), p. 14.

up their reportage and commenting on it. But if all of the conventional media disappear, including news-gathering agencies of all sorts (newspapers, wire services, and so on), how will the blogosphere supplant them with unpaid amateurs (Goldman, 2008)?⁴

Where this view fails, and where Goldman (2010) also suffers from a failure of imagination, is that the epistemic norms of truth-seeking can be agreed upon at the outset of any specific inquiry, and particular agreed upon epistemic constraints can go a long way in preventing some of the issues that he outlines here. Although this may or may not be case with Wikipedia, it does not indict all informal projects of this type and it certainly does not indict the entire blogosphere.

There is something going on with Goldman's (2010) analysis of systems-oriented social epistemology that appears only to validate the formal systems and formal institutions already in place as bearers of public knowledge. His critique of Wikipedia and his disdain for the blogosphere seem to only highlight an underlying distrust of the "social" part of social epistemology. My initial concern here was one of *size and scale of our social epistemological analysandum* and a brief look at Goldman's (2010) approach raises a few concerns about starting with the formal institutions we do know well and trying to move into the greater orders of complexity that make up informal institutions over time. The first problem, well illustrated by Goldman's (2010) response to the "problem" of the blogosphere, is that informal institutions may be too easily dismissed as useless rabble-rousing instead of considered on the grounds of their own distinct realms of inquiry. The second problem, a meta-theoretical one, is that even if informal institutions are not dismissed, their truth-seeking activities may fail to be captured as such before they are either dismissed or co-opted into more established formal institutions. No matter what the scale, we must have some way of capturing the activities of both formal and informal institutions without arbitrarily restricting the content of their respective realms of inquiry at the *outset* of inquiry.

6.2 A Network-Level Epistemology

It is here that our social epistemology must meet the real world – in bridging our theoretical construction of truth-making with ways of indexing, mapping, and evaluating the activities of various communities of inquirers, whether they be within formal institutions or informal ones. A systems-oriented epistemology misses the truth-gathering potential of informal institutions, or informal communities of

⁴ Goldman (2010), p. 15-16.

inquirers, in favor of chronicling the activities of institutions that have already been sanctioned, so to speak, by our current social, cultural, political apparatus. Research institutions, courts of law, and governmental organizations are good places to start, but they are a long way from capturing the sheer breadth and scope of contemporary inquiry. For this reason, I propose thinking of our social epistemology not as systems but as *networks*. Systems are delineated entities; they have predetermined boundaries that exist at the start of any inquiry or evaluation. Networks, on the other hand, may consist of interconnected systems, but it also may consist of independent, shared, and/or relational data. Our social epistemology should be networked and it should also find a way to capture from all communities of inquirers, if so desired. But how would this work?

To answer this, I look finally to work done in scientific citation studies, cluster studies, and cluster mapping. It is important to note that information science and computer science are at the forefront of mapping the dynamics of scientific activities. It is a discussion that is broad and complex enough to require its own evaluation and its own philosophical discussion such that a thorough analysis is beyond the scope of what I can do here. That said, I conclude my project by giving a very brief introduction to the basic elements of mapping dynamics with an overview of scientific citation work and some preliminary thoughts as to how it can be used for the future of veritistic social epistemology.

6.2.1 Scientific Citation Studies and Scientific Citation Mapping

In Scharnhorst, Börner, and van den Besselaar's (2012) *Models of Science Dynamics: Encounters Between Complexity Theory and Information Science*, they claim that the project of mapping scientific activities is well underway and now moving in the direction of mapping analysis or explanation.⁵ Seminal work done by Garfield (1955) in scientific citation studies created a path for "networks of publications and their citation patterns, word use, collaborating researchers, [and] topics in e-mail threads [to be] measured, analysed and visualized over time," such that now we have "maps of science" as created by ISI Thompson Reuters Web of Knowledge, Elsevier's Scopus, and the NSF funded project, "Mapping Science."⁶ Garfield (1955) initially argued that the way scientists were expected to engage in

⁵ See: Scharnhorst, Börner, and van den Besselaar (2012) *Models of Science Dynamics: Encounters Between Complexity Theory and Information Science*, New York, New York: Springer, p. xii.

⁶ *Ibid.*

progressive science post-World War II required spending too much time looking for the citation information for past work. He (1955) writes:

It is too much to expect a research worker to spend an inordinate amount of time searching for the bibliographic descendants of antecedent papers. It would not be excessive to demand that the thorough scholar check all papers that have cited or criticized such papers, if they could be located quickly. The citation index makes this check practicable. Even if there were no other use for a citation index than that of minimizing the citation of poor data, the index would be well worth the effort required to compile it.⁷

Garfield (1955) proposed creating a citation index according to categories of ideas, or what he called “association-of-ideas.”⁸ The citation index serves to index the scientists’ “thoughts” instead of indexing only discreet “concepts” within broader texts. In this context, “thoughts” are the primary arguments or intellectual ideas within the publication and “concepts” are the broader topics. The index aims to fill a gap in need between the authors of works and the scientific researchers looking for the history of any particular line of inquiry. This gap is filled by connecting significant publications in a given field, or for a given publication, to the publications that they subsequently influenced. As Garfield (1955) describes in the following example, the citation index anticipates and identifies what many scientific researchers may never find through many hours of tedious study on the scholarly influence of even *a single significant publication*, let alone multiple significant publications within a given field or specialty:

For example, [an] article on information theory, if thoroughly indexed, might have included an entry under reading devices for the blind.

Yet if this were done, our periodical indexing services would clearly become hopelessly overloaded with material that is not necessary to lead us to the micro unit – the entire article or one of its major sections... It would require an army of indexers to read the articles and identify the exact subject matter of every paragraph or sentence... Were an army of indexers available, it is still doubtful that the proper subject indexing could be made.”⁹

⁷ See: Garfield, Eugene. (1955). “Citation Indexes for Science: A New Dimension in Documentation through Association of Ideas” in *Science* 122 (3159), pp. 108-111.

⁸ Garfield (1955), p. 108.

⁹ Garfield (1955), p. 110.

The citation index attempts to weed out the noise in standard periodical indexes. To do this, the index creator must establish an initial reference list for a publication, or set of publications, by identifying the categories of ideas, or “association-of-ideas” relevant to that publication and creating a bibliography of all subsequent publications that contain the initial citation. After a bibliographic list is created, the initial reference list is cross-referenced again with those citations for how well their information relates to the initial list. That is why scientific citation indexes are more useful than standard indexes and that is also why they have been useful in subsequent attempts to “map” contemporary science.

To this, Scharnhorst and Garfield (2010) suggest in “Tracing Scientific Influence,” that large-scale scientific citation indexes – now including communication between authors, email exchanges, and citation linkage, as well as other non-bibliographical information – can be generated by computer programs that create the initial scientific citation index and connect them in networks, mapping both bibliographic and non-bibliographic information.¹⁰ The result being that information scientists can map science in greater and more inclusive detail, linking active areas of inquiry with institutional relationships and insights about the socio-scientific motivations of inquiry. This has led Scharnhorst, Börner, and van den Besselaar (2012) to ask the secondary question of who’s engaging in scientific inquiry (“who are the actors driving the development of science”), where, and why:¹¹

Earlier large-scale maps concentrated on scientific communications as manifested in papers and their citation interlinkage (Scharnhorst and Garfield 2010). Partly, this was due to the fact that unique author names are hard to determine because of same names, name variants and misspellings. So, a large part of bibliometrics and scientometrics analyses texts (titles, keywords, words, references). Some automated techniques have partly solved this problem, at least on a higher level of aggregation. In maps of scientific communication, authors as well as institutions can now be made visible with a higher reliability. To explain the networks in which researchers are linked (by publishing or communicating), current research in social-psychology and sociology of science becomes relevant. Resume analysis, ethnographic observations, and interviews were presented as ways to gain access to local motivations and behavior the collective effect of which

¹⁰ Scharnhorst and Garfield (2010), p. 1-2.

¹¹ Scharnhorst, Börner, and van den Besselaar (2012), p. xii.

is reflected in the large-scale global maps of science.¹²

6.2.2 “Open” and “Closed” Network-Mapping for Social Epistemology

This question about the current actors driving scientific activity is not a benign one and something needs to be said about how certain citation indexes and maps can be influenced by the formal institutions that produce the citations in the first place. In this section, I will consider a few of the data visualization projects created by *Places & Spaces: Mapping Science* from the Cyberinfrastructure for Network Science (CNS) Center at the Department of Information and Library Science at the School of Informatics and Computing at Indiana University. The *Places & Spaces* project aims to “inspire cross-disciplinary discussion on how to best track and communicate scholarly activity and scientific progress on a global scale.”¹³ The breadth and depth of these maps serve as constructive examples of how the “actors” of scientific inquiry (excluding the mapmakers as “actors” of meta-inquiry) are significant to the results of inquiry and/or the knowledge produced. In “closed” systems, as generally represented by formal institutions, the content created by the system, including the raw data that may be generated by it, is regulated by who’s allowed to participate in the data gathering (or “epistemic evidence” collecting) and what policies and procedures govern the data that is sanctioned for use by the system itself. In “open” systems, there is a lot more “epistemic noise,” in that the participants in any particular inquiry cannot be meaningfully restricted and the epistemic evidence they collect does not have to be sanctioned for use in the process of inquiry.¹⁴ What Scharnhorst, Börner, and van den Besselaar (2012) picked up on is the issue that most of the current scientific maps were created using formal institutional resources (including employing institutionally-supported methodologies and/or philosophies) as well as significant restrictions on which “actors” of scientific inquiry count in terms of mapping content. Consider the following *Places & Spaces* maps:

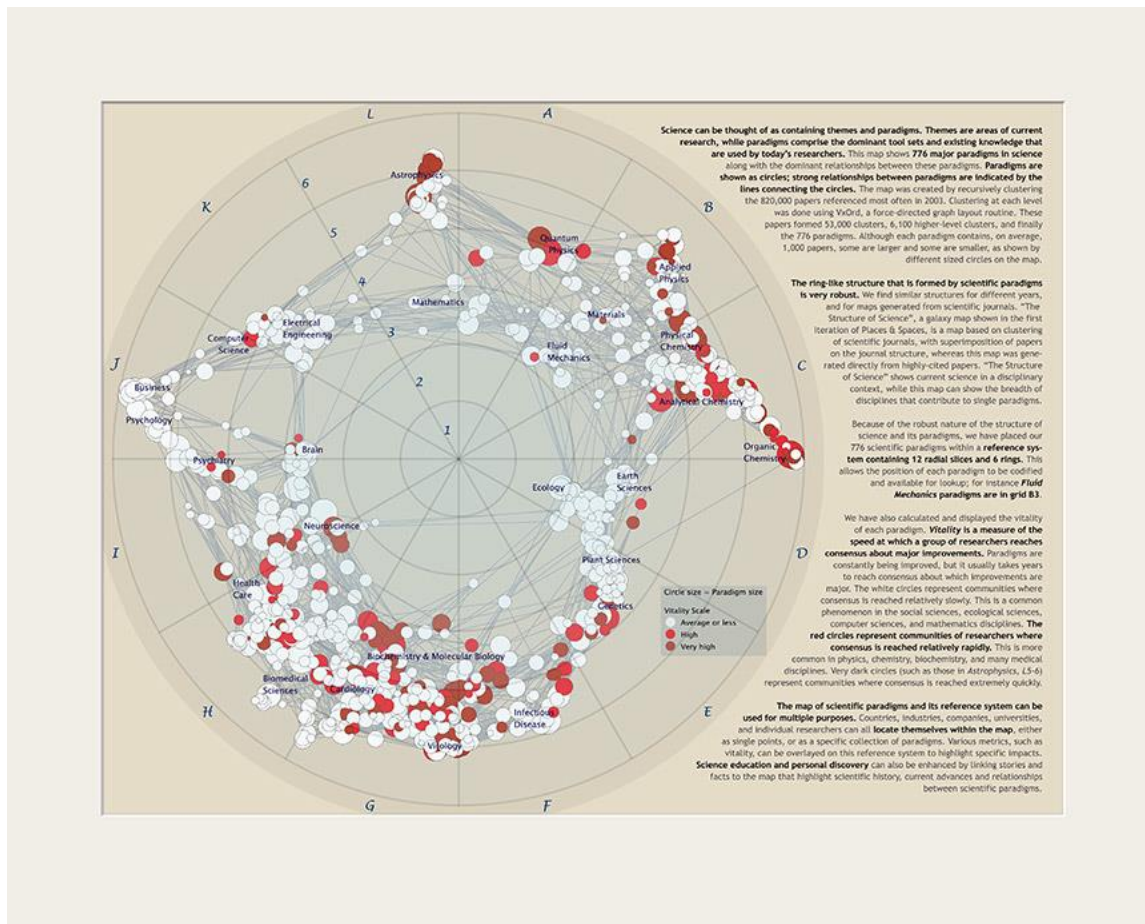
¹² Scharnhorst, Börner, and van den Besselaar (2012), p. xii-xiii.

¹³ “Places & Spaces: Mapping Science” at the Cyberinfrastructure for Network Science Center in the Department of Information and Library Science at the School of Informatics and Computing at Indiana University. http://scimaps.org/what_is_a_science_map.html. January 2015.

¹⁴ I do not attend to make a meaningful parallel between “epistemic noise” and “data noise.” I do want to insist that “epistemic noise” is (a) “information,” strictly-speaking, and (b) able to be parsed by both human and machine.

Figure 1.

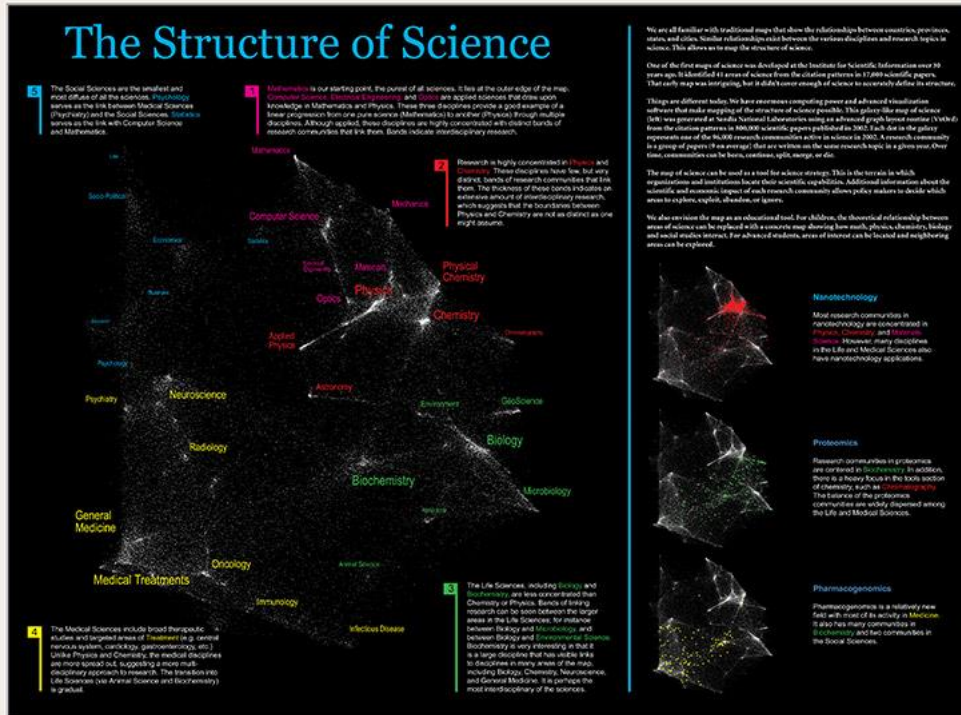
Map of Scientific Paradigms by Kevin W. Boyack and Richard Klavans



Kevin W. Boyack and Richard Klavans create maps of science that can be used for planning and evaluation on the national, corporate, and personal levels. Science itself can be thought of as containing themes and paradigms: themes are current areas of research, while paradigms comprise the dominant tool sets and existing knowledge that are used by current researchers. To visualize these scientific paradigms, Boyack and Klavans used the VxOrd graph layout tool to recursively cluster the 820,000 most important papers referenced in 2003, resulting in 776 paradigms. The most dominant relationships between paradigms were also calculated and are shown as lines between paradigms. The map of scientific paradigms constitutes a reference system that can be used for multiple purposes. Countries, industries, companies, and individual researchers can all locate themselves within the map, either as a single point or as a specific collection of paradigms. Science education and discovery can also be enhanced by linking to the map stories and facts that exemplify content and relationships between scientific paradigms.

Figure 2.


The Structure of Science by Kevin W. Boyack and Richard Klavans



Kevin W. Boyack and Richard Klavans share a deep interest in the mapping of science as a platform for planning and evaluation on the national, corporate, and personal levels. Through a multistep process, this galaxy-like map of science was created from citation patterns in 800,000 scientific papers published in 2002. Each dot represents one of 96,000 active research communities. Over time, communities can be born, grow, split, merge, or die. By coupling coefficients between papers, using the VxOrd layout algorithm and a modified single-link clustering routine, Boyack and Klavans were able to calculate the placement of research communities. For example, communities made up solely of papers in biochemistry journals show up in the biochemistry section, while communities that are evenly split between biochemistry and chemistry journals show up midway between biochemistry and chemistry. In 2005, this was the most comprehensive and most accurate literature map ever generated.

Figure 3.

Science-Related Wikipedia Activity by Bruce W. Herr II, Todd M. Holloway, Elisha F. Hardy, Kevin W. Boyack, and Katy Börner



Developed by research programmers Bruce W. Herr II and Todd M. Holloway, graphic designer Elisha F. Hardy, and information scientists Kevin W. Boyack and Katy Börner, this map shows the structure and dynamics of the English Wikipedia based on 659,388 articles and their editing activity. The similarity of each article-article pair was calculated as the number of shared links to other articles. The resulting similarity matrix was read into VxOrd to generate the base map layout. An invisible 37 x 37 half-inch grid was drawn underneath the network and filled with relevant images from key articles. Overlaid are 3,599 math, 6,474 science, and 3,164 technology articles. They are color-coded in blue, green, and yellow, respectively, with all other articles appearing in grey. Exactly 8,181 articles are in one category, 2,348 in two, and 73 in all three categories. The four corners show smaller versions of the map with articles size-coded according to article edit activity (top left), number of major edits from January 1st, 2007, to April 6th, 2007 (top right), number of bursts in edit activity (bottom right), and the number of times other articles link to an article (bottom left). These visualizations serve to highlight current trends and predict future editing activity and growth in Wikipedia articles related to science, technology, and mathematics. Interactive Wikipedia maps are available at <http://scimaps.org/web/maps/wikipedia/>.

Boyack and Klavans' (2006) project (Fig. 1) maps scientific "themes" and "paradigms" where "themes" are described as and mapped according to areas of current scientific research, and "paradigms" are described as and mapped according to "dominant tool sets and existing knowledge that are used by today's researchers."¹⁵ The map represents 776 significant scientific paradigms as topic clusters and the most significant relationships between those paradigms as the lines between the topic clusters. The map used a recursive algorithm for large datasets, or "recursive clustering," for the approximately one million papers most referenced in 2003 in order to create both higher and lower-order level clustering.¹⁶ The researchers used VxOrd, a forced-directed graph layout routine to generate 53,000 clusters, 6,100 higher-order clusters, and ultimately 776 paradigms.¹⁷ The topic cluster colors on Boyack and Klavans' (2006) map represents paradigm "vitality," or "the speed at which a group of researchers reaches consensus about major improvements."¹⁸ From the map, the white circles represent scientific communities where consensus is reached slowly; these communities include the social sciences, ecological sciences, computer sciences, and mathematical sciences, the red circles represent communities where consensus is achieved relatively quickly, these communities include physics, chemistry, biochemistry, and the medical sciences.¹⁹ The very dark circles represent communities where consensus is achieved very quickly, such as astrophysics, physical chemistry, virology, and neuroscience.²⁰

Prior to Boyack and Klavans (2006), Boyack and Klavans (2002), in Fig. 2., created a similar map by mapping 800,000 scientific publications from 2002. In this case, they mapped the activity of scientific communities with "communities" defined as groups of papers (approx. 9) writing about the same research

¹⁵ Boyack, Kevin W. & Klavans, Richard. 2006. *Map of Scientific Paradigms*. Courtesy of Kevin W. Boyack and Richard Klavans, SciTech Strategies, Inc. In "2nd Iteration (2006): The Power of Reference Systems," *Places & Spaces: Mapping Science*, edited by Katy Börner and Deborah MacPherson. <http://scimaps.org>. They describe their mapping project as follows: "Boyack and Klavans used the VxOrd graph layout tool to recursively cluster the 820,000 most important papers referenced in 2003, resulting in 776 paradigms. The most dominant relationships between paradigms were also calculated and are shown as lines between paradigms. The map of scientific paradigms constitutes a reference system that can be used for multiple purposes. Countries, industries, companies, and individual researchers can all locate themselves within the map, either as a single point or as a specific collection of paradigms. Science education and discovery can also be enhanced by linking to the map stories and facts that exemplify content and relationships between scientific paradigms" (<http://scimaps.org>).

¹⁶ *Ibid.*

¹⁷ *Ibid.*

¹⁸ *Ibid.*

¹⁹ *Ibid.*

²⁰ *Ibid.*

topic in a given year; merging, splitting, growing, and dying scientific communities are also represented on the map.²¹ Boyack and Klavans (2006) used overlapping information between the (2002) map and the (2006) map, however, the raw data generated different visualizations about the activity of science: “‘The Structure of Science’ shows current science in a disciplinary context, while [‘Map of Scientific Paradigms’] can show the breadth of disciplines that contribute to single paradigms.”²² Notice, however, how both of Boyack and Klavans’ (2002, 2006) scientific maps are drawing upon overlapping, if not identical, data sets. Compare this to Fig. 3, *Science-Related Wikipedian Activity*, where the researchers show the dynamics almost 700,000 English language Wikipedia articles and editing activity in math, science, and technology to highlight current scientific activity and predict future scientific activity as related to Wikipedia entries, edits, and shared links. The content of these data sets are “open,” meaning that the underlying data may or may not be sanctioned by formal research institutions or formal research publications *and* that are no policies or prohibition on who can contribute to the data other than what is currently practiced and accepted by the larger user-community, i.e., a community-derived standard. This is in contrast to the activity we see represented in Fig. 1-2, which both utilize what I call “closed” data sets, or data that is sanctioned by formal research institutions and/or formal research publications. As Scharnhorst, Börner, and van den Besselaar (2012) mention, if these current scientific maps supposed to allow us to ask the who, where, and why part of contemporary scientific activity, then it is in the relation between shared “closed” data sets and the representations of the activities of scientific communities that we should start. How is it that “closed” systems contribute to knowledge? And where does that leave “open” systems like Wikipedia?

6.2.3 Can “Closed” Systems Contribute to Knowledge?

It seems clear from an evaluation of Figures 1-2 that “closed” data sets and the “closed” systems from which they arise surely do contribute to knowledge in a meaningful sense. The issue is that closed systems may not accurately gauge the margin of error for any particular line of inquiry (or hypothesis or set of hypotheses) because their data sets are fundamentally exclusionary. Maps of current scientific

²¹ Boyack, Kevin W., and Richard Klavans. 2005. *The Structure of Science*. Courtesy of Kevin W. Boyack, Sandia National Laboratories and Richard Klavans, SciTech Strategies, Inc. In “1st Iteration (2005): The Power of Maps,” *Places & Spaces: Mapping Science*, edited by Katy Börner and Deborah MacPherson. <http://scimaps.org>.

²² *Ibid.*

activity do not include activity that is unsanctioned by formal research institutions and formal research publications. And although current scientific maps do include non-bibliographic information, they do not include, for example, work done by scientists who are not currently associated with formal research institutions, or who have publications that are not peer-reviewed by formal research publications. There is, of course, good reason for this. Current scientific research is regulated by industry-specific policies and procedures that aim to protect scientists (e.g., institutional backing for research ideas), research subjects (e.g., ethics committees, mandatory protocols training), and the public alike (e.g., governmental oversight) from potential abuses of poor, questionable, or dangerous scientific methodology. We only need to take a cursory look through the history of modern science to see why this is necessary, e.g., the Tuskegee experiments (1932), the “Monster Study” (1939), the Nazi “medical” experiments, the Milgram experiments (1963), the Willowbrook Hepatitis B experiments (1955-1970), etc. That said, for the larger social epistemological view that I argued for in Chapter 5 to work, our notion of scientific communities will also have to shift from solely closed systems to both open *and* closed systems. This will inevitably introduce quite a bit of epistemic noise into analyses of scientific activity, and, in particular, it will necessarily introduce epistemic noise into our attempts at scientific mapping – but we *need* the data produced by informal institutions, or open systems, in order to engage in successful truth-tracking. We can hope to eliminate this epistemic noise both theoretically and technologically by (a) adopting a belief-stance, as communities of inquirers, toward belief-updating which utilizes our revised Bayesian selection function and principle of epistemic disagreement for first-order epistemic scrutiny, and (b) employing scientific mapping tools and data visualizations, like the examples shown above, for second-order epistemic scrutiny.

6.2.4 Network-Mapping and Emerging and Waning Clusters

The future of mapping science will continue to illustrate and anticipate what we already see happening: running cluster studies over several years shows us emerging and waning topic clusters within scientific activity, very broadly construed. What we want is for current emerging and waning topic clusters to be predictive for *future* emerging and/or waning topic clusters, thus providing insight into which scientific activities, which fields of inquiry, and which field-specific hypotheses are growing or dying according to the cluster-map. The problem, for now, is that citations are the backbone of scientific activity

mapping and citations also have a “half-life,” or a measure of how long the citation continues to be cited. What determines the half-life of citation may have nothing to do with how significant the publication is or was, it may be the result of a defunct publisher or the result of lost or destroyed archives. But we might ask, much like the social epistemological view that I argued for, isn't possible that topic clusters are diachronic and that they can be shown to be true or false over long periods of time much like any proposition P ? I.e., the topic is a function of the frequency with which it is cited and its relevance to the truth of certain scientific hypotheses? Unfortunately, even if we can resolve that issue (which is rather tricky, historically speaking), there is the additional problem that topic clusters themselves, unlike citations, are the results of subjective cognitive labor. That is, if we take citation studies to be of the first-order, then second-order citation studies, in terms of assigning topics to clusters of citations (or other bibliographic and non-bibliographic information) represents second-order cognitive labor. Thus we would need to parse clusters for relevance qua cluster and we now we have a regress problem! As I have argued throughout most of this project, our attempt to weed out subjectivity at the start of inquiry – even at the level of grouping similar citations together – is hopelessly flawed. Because the bottom line for all inquiry is that it must start somewhere. Topic cluster mapping is no different in that respect. Our best hopes are to take what we can know and move forward with our analysis of scientific activity in good faith.

6.3 Representing Large-Scale Data: Concluding Remarks

The social epistemological project that I have argued for has taken new shape in this chapter – and the practical application of this project is riddled with as many issues as the theoretical one of the last chapter. That is no surprise. The data visualization tools that we use to represent scientific activity and scientific communities only reflect the same prejudices, social relations, markers of prestige, etc., that are part of inquiry and part of the communities themselves. The hope here is twofold. First, that truth-making and truth-tracking are continually in process and that as communities of inquirers we can improve upon our current states of belief via the methodological processes of scientific inquiry coupled with some normative requirements on belief-making and belief-updating. The other hope is that technology, and specifically the ability to use technology to map and track our scientific communities in increasingly public ways, allows the Peirce's theoretical “community of inquirers” to actually see itself and engage with itself as such – a community. Our knowledge-producing activities do not happen in a theoretical vacuum.

They happen in very real, very human spaces. We are now in a position to merge these two spaces and make our epistemology truly social.

6.4 A Final Note on the “Social” Part of Social Epistemology

In the first chapter, I suggested taking “social epistemology” to be a cluster of views about how social practices, behaviors, and institutions, etc. inform both knowledge-production and the “truth-making” process. I am not committed to any particular understanding of the “social” part of social epistemology full stop. This is because *all* epistemology is essentially social epistemology and it matters little whether or not we come to any consensus about how to best understand what we mean by the “social” part of social epistemology – this is not the same as saying that we can operate with absolutely *no working conception* of the “social.” If we can show that social context is a necessary part of answering epistemological questions of any variety, then all such discussions merely become matters of emphasis or degree. In this sense, what constitutes the “social,” i.e., different social communities, varying social behaviors, etc. will have different meanings in different social contexts. Thus what we broadly mean by “social” or “social context” should remain flexible and fluid enough to accommodate a multitude of such communities and practices. This is particularly important because the socio-behavioral makeup and practices of digital spaces and digital communities requires that we cast the net of the “social” so much wider than we ever have previously. That is why I have insisted in this project on focusing on epistemic constraints for socially-constituted *methods* not on epistemic constraints for, or epistemic justification of, the belief-states of community members. Now there is a concern that this version of social epistemology is too narrow, focusing on aggregating doxastic attitudes by method, and not on challenging the socially-constructed practices of science that too often serve to reify the social, cultural, and political biases that makes scientific practice so epistemologically fraught. This critique is further complicated by the fact that I have also insisted on the value of “epistemic communities” but have left those communities mostly prescription-less. It may even be fair to ask whether I have simply taken the steam out of “social epistemology” in general by divorcing it from the social practices that make it up. How can my view here be considered *meaningfully social*?

Making normative prescriptions about social practices in science, particularly as they pertain to social interactions and/or relations between individuals within a given epistemic community is not my aim

here – although it is a worthy one. In my view, to engage in social epistemology one need only see the world through a naturalist lens and be committed to what I broadly conceive of as the “scientific method” – and, I hope, to make use of good Bayesian methods, like the “Bayesian selection function” outlined in Chapter 5. That being said, any questions under investigation within a given community will themselves be community-derived and thus may vary widely in their content and applicability to Bayesian analyses. Some questions arriving from the community may deal entirely with non-veritistic epistemic issues that I don’t consider here, such as issues of membership, authority, exclusion/inclusion, marginalization, etc. Those epistemic issues are non-trivially community-dependent. And although this might read as unwillingness on my part to get my hands dirty by fleshing out what constitutes social and behavioral interactions that are *conducive* to “good science,” I do take my version of social epistemology as merely starting closer to the ground floor. Making normative prescriptions and/or recommendations about social interactions and behaviors require us to assume that we can augment the role of bias and prejudice in (broadly) scientific inquiry by collectively affirming our own fallibility and agreeing upon fruitful ways of changing our interactions such that they don’t stand in the way of successful inquiry. But, following Solomon (2001), the idea of meaningfully challenging and changing individual biases seems idealized and overly optimistic at best. We can easily imagine a community that is too hostile to accusations of individual bias and/or prejudice to even consider a change of culture or practice. A more modest approach starts with the adoption of a coarser-grained apparatus, i.e., a way of weeding out of biases and/or prejudices in prior belief-states at the level of *method*. This approach combined with increasing ways to visualize epistemic communities utilizing the Bayesian selection function (or the “scientific method” more generally) may eventually allow us to further analyze epistemic communities with high probabilities, or seeming consensus, around a given hypothesis, p , (or, correspondingly, *low probabilities* around a given p) and evaluate the socio-cultural conditions *that may have allowed* for such consensus. This would be the start of a formidable research project within social epistemology – one that makes use of emerging tools in digital humanities to more meaningfully identify epistemic communities and the socio-cultural elements of community practice may be actively helping or harming the pursuit of successful inquiry. It would hopefully pick up where this project must conclude for now.

Bibliography

- Alcoff, Linda and Elizabeth Potter, eds. *Feminist Epistemologies*. New York: Routledge (1993).
- Alston, William P. *Beyond "Justification": Dimensions of Epistemic Evaluation* Ithaca: Cornell University Press (2005).
- Almeder, Robert. *Harmless Naturalism: The Limits of Science and the Nature of Philosophy*. Chicago: Open Court Publishing (1998).
- Anderson, Elizabeth. "Uses of Value Judgments in Science: A General Argument, with Lessons from a Case Study of Feminist Research on Divorce" in *Hypatia*, Vol. 19, No. 1, Feminist Science Studies (Winter, 2004) (2004), pp. 1-24.
- Anthony, Louise. "Symposium: Feminist Epistemology: Comment on Naomi Scheman" in *Metaphilosophy* 26 (3) (1995), pp. 191-198.
- Anthony, Louise. *A Mind of One's Own*. Boulder: Westview Press (1993).
- Arntzenius, Frank. "Reflections on Sleeping Beauty" in *Analysis*, 62, (2002), pp. 53-61.
- Arntzenius, Frank. "Some Problems for Conditionalization and Reflection" in *Journal of Philosophy*, 100, (2003), pp. 356-370.
- Berkeley, Hastings. "The Kernel of Pragmatism" in *Mind*, New Series, Vol. 21, No. 81 (Jan., 1912) (1912), pp. 84-88.
- Brandom, Robert. *Making It Explicit: Reasoning, Representing, and Discursive Commitment*. Cambridge: Harvard University Press (1998).
- Brandom, Robert. "Pragmatics and Pragmatisms," in *Hilary Putnam: Pragmatism and Realism*, eds., James Conant and Urszula M. Zeglen. New York: Routledge (2002).
- Brandom, Robert. *Between Saying and Doing: Toward an Analytic Pragmatism*. New York: Oxford University Press (2008).
- Brandom, Robert. *Perspectives on Pragmatism: Classical, Recent, and Contemporary*. Cambridge: Harvard University Press (2011).
- Carnap, Rudolf. "Empiricism, Semantics, and Ontology" in *Meaning and Necessity: A Study in Semantics and Modal Logic*. Chicago: University of Chicago Press (1950).
- Coady, C. A. J. "Testimony and Observation" in *American Philosophical Quarterly* 108 (2) (1973), pp. 149-55.
- Coady, David. *What to Believe Now: Applying Epistemology to Contemporary Issues*. Oxford: Wiley-Blackwell (2012).
- Code, Lorraine. *What Can She Know? Feminist Theory and the Construction of Knowledge*. Ithaca: Cornell University Press (1991).

- Code, Lorraine. "Taking Subjectivity Into Account" in *Feminist Epistemologies*, eds. Linda Alcoff and Elizabeth Potter. New York: Routledge (1993).
- Chalmers, David J. "Ontological Anti-Realism" in *Metametaphysics: New Essays on the Foundations of Ontology*, eds. David J. Chalmers, David Manley, Ryan Wasserman. New York: Oxford University Press (2009).
- Christensen, David. "Clever Bookies and Coherent Beliefs" in *The Philosophical Review*, Vol. 100, No. 2, (Apr., 1991) (1991), pp. 229-247.
- Christensen, David. "Diachronic Coherence versus Epistemic Impartiality" in *The Philosophical Review*, Vol. 109, No. 3, (Jul., 2000) (2000), pp. 349- 371.
- Christensen, David. "Epistemology of Disagreement: The Good News" in *Philosophical Review*, 116, (2007), pp. 187-217.
- Davidson, Donald. "On the Very Idea of a Conceptual Scheme" in *Proceedings and Addresses of the American Philosophical Association* 47, (1973), pp. 5-20.
- Davidson, Donald. *Inquiries Into Truth and Interpretation*. New York: Oxford University Press (1984).
- Dennett, Daniel. *Darwin's Dangerous Idea: Evolution and the Meanings of Life*. New York: Simon & Schuster (1995).
- Descartes, René. *Meditations on First Philosophy: With Objections and Replies (Cambridge Texts in the History of Philosophy)*. Cambridge: Cambridge University Press (1996).
- Dewey, John. *The Quest for Certainty: A Study of the Relation of Knowledge and Action*. New York: Minton, Balch and Company (1929).
- Dewey, John. *Logic: The Theory of Inquiry*. New York: Holt, Rinehart & Winston (1938).
- Dewey, John. (1917). "The Need for a Recovery of Philosophy" in *The Essential Dewey, Vol. 1: Pragmatism, Education, Democracy*, eds. Larry A. Hickman and Thomas M. Alexander. Bloomington: Indiana University Press (1998).
- Dorr, Cian. "Sleeping Beauty: In Defense of Elga" in *Analysis*, 62 (2002), pp. 292-296.
- Elga, Adam. "Self-locating Belief and the Sleeping Beauty Problem" in *Analysis*, 60, (2000), pp. 143-147.
- Elga, Adam. "On Overrating Oneself... And Knowing It" in *Philosophical Studies* 123, (1-2), (2005), pp. 115-124.
- Elga, Adam. "Reflection and Disagreement" in *Noûs*, Vol. 41, No. 3, (2007), pp. 478-502.
- Elga, Adam. "How To Disagree About How to Disagree" in *Disagreement*, eds. Ted Warfield and Richard Feldman. New York: Oxford University Press (2010).

- Fallis, Don. "Goldman on Probabilistic Inference," in *Philosophical Studies*, 109: 223-240, (2002).
- Fallis, Don. "Toward an Epistemology of Wikipedia" in *Journal of the American Society for Information Science and Technology*, 59, 10, (2008), pp. 1662-74.
- Flanagan, Owen. "Varieties of Naturalism," in *Oxford Handbook of Religion and Science*, ed. Philip Clayton. New York: Oxford University Press (2006).
- Frances, Bryan. "Discovering Disagreeing Epistemic Peers and Superiors" in *International Journal of Philosophical Studies*, 20 (1), (2012), pp. 1-21.
- Fuhrmann, André. "Is Pragmatist Truth Irrelevant to Inquiry?" in *Knowledge and Inquiry: Essays on the Pragmatism of Isaac Levi*, ed. Erik J. Olsson. Cambridge: Cambridge University Press (2006).
- Fuller, Steve. "Provocation on Reproducing Perspectives: Part 3" in *Social Epistemology: A Journal of Knowledge, Culture and Policy* 2 (1) (1988), pp. 99-101.
- Fuller, Steve. *Social Epistemology*. Bloomington: Indiana University Press (2002).
- Garfield, Eugene. "Citation Indexes for Science: A New Dimension in Documentation Through Association of Ideas" in *Science* 122 (3159) (1955), pp. 108-111.
- Gillies, Donald. "The Duhem Thesis and the Quine Thesis," in *Philosophy of Science: The Central Issues*, eds. Curd and Cover. New York: W. W. Norton & Company (1998).
- Gillies, Donald. "Varieties of Propensity" in *British Journal for the Philosophy of Science* 51 (4) (2000), pp. 807-835.
- Godfrey-Smith, Peter. *Complexity and the Function of Mind in Nature*. Cambridge: Cambridge University Press (1998).
- Godfrey-Smith, Peter. "Dewey, Continuity, and McDowell," in *Naturalism and Normativity*, eds. David MacArthur and Mario de Caro. New York: Columbia University Press (2010).
- Goldman, Alvin I. *Liaisons: Philosophy Meets the Cognitive Sciences*. Cambridge: The MIT Press (1992).
- Goldman, Alvin I. *Knowledge in a Social World*. New York: Oxford University Press (1999).
- Goldman, Alvin I. "Why Social Epistemology is *Real* Epistemology" in A. Haddock, A. Millar, and D. Pritchard, eds. *Social Epistemology*. Oxford: Oxford University Press (2010).
- Goldman, Alvin I. "Systems-Oriented Social Epistemology" in *Oxford Studies in Epistemology, Vol. III*, eds. Tamar Gendler and John Hawthorne. New York: Oxford University Press (2010).
- Green, Thomas Hill. *Hume and Locke*. New York: Crowell (1968).
- Haack, Susan. *Evidence and Inquiry: A Pragmatist Reconstruction of Epistemology*. Amherst: Prometheus Books (1993).

- Haack, Susan and Konstantin Kolenda. "Two Fallibilists in Search of the Truth," in *Proceedings of the Aristotelian Society, Supplementary Volume*, Vol. 51 (1977), pp. 63-104.
- Hájek, Alan. "Interpretations of Probability" in *The Stanford Encyclopedia of Philosophy (Spring 2010 Edition)*, ed. Edward N. Zalta <http://plato.stanford.edu/archives/spr2010/entries/probability-interpret/> (2010).
- Hájek, Alan. "'Mises Redux'—Redux: Fifteen Arguments against Finite Frequentism" in *Erkenntnis*, Vol. 45, No. 2/3, *Probability, Dynamics, and Causality* (Nov., 1996), pp. 209-227.
- Harding, Sandra. *Whose Science? Whose Knowledge?: Thinking from Women's Lives*. Ithaca: Cornell University Press (1991).
- Hookway, Christopher. "Truth, Reality, and Convergence," in *The Cambridge Companion to Peirce*, ed. Cheryl Misak. Cambridge: Cambridge University Press (2004).
- Hume, David. (1748). *An Enquiry Concerning Human Understanding*, ed. Tom L. Beauchamp. New York: Oxford University Press (2000).
- James, William. "Pragmatism's Conception of Truth" in *Pragmatism: A New Name for Some Old Ways of Thinking*. New York: Longman, Green, and Company (1907).
- James, William. *The Meaning of Truth*. New York: Longman, Green, and Company (1909).
- James, William. *Pragmatism and the Meaning of Truth*. Cambridge: Harvard University Press (1978).
- James, William. *Essays in Radical Empiricism*. Charleston: BiblioLife Books (2009).
- Joyce, James M. "The Development of Subjective Bayesianism" in *Handbook of the History of Logic, Volume 10: Inductive Logic*, eds. Dov Gabbay, Stephan Hartmann and John Woods. New York: Elsevier (2010).
- Kant, Immanuel. *Critique of Pure Reason*, trans. Werner S. Pluhar, intro. Patricia W. Kitcher. Cambridge: Hackett Classics/Hackett Publishing Company (1996).
- Keller, Evelyn Fox and Helen E. Longino, eds., *Feminism and Science*. New York: Oxford University Press (1996).
- Kim, Jaegwon. "What Is 'Naturalized Epistemology'?" in *Philosophical Perspectives*, Vol. 2, (1988), pp. 381-405.
- Kitcher, Philip. "The Naturalists Return," in *The Philosophical Review*, Vol. 101, No. 1 (Jan., 1992), pp. 53-114.
- Kitcher, Philip. *The Advancement of Science: Science Without Legend – Objectivity Without Illusions*. New York: Oxford University Press (1993).
- Kitcher, Philip. "Veritistic Value and the Project of Social Epistemology," in *Philosophy and Phenomenological Research*, Vol. LXIV, No. 1, January (2002), pp. 191-198.
- Kitcher, Philip. *Preludes to Pragmatism: Toward a Reconstruction of Philosophy*. New York: Oxford University Press (2012).

- Kornblith, Hilary. "A Conservative Approach to Social Epistemology," in *Socializing Epistemology: The Social Dimensions of Knowledge*, ed. Frederick F. Schmitt. New York: Rowman & Littlefield (1994).
- Kuhn, Thomas. *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press (1962).
- Kuhn, Thomas. *The Essential Tension: Selected Studies in Scientific Tradition and Change*. Chicago: University of Chicago Press (1977).
- Kuhn, Thomas. *The Road Since Structure: Philosophical Essays, 1970-1993*. Chicago: University of Chicago Press (2000).
- Lackey, Jennifer and David Christensen, eds. *The Epistemology of Disagreement: New Essays*. New York: Oxford University Press (2013).
- Latour, Bruno and Steven Woolgar. *Laboratory Life: The Construction of Scientific Facts*. Princeton: Princeton University Press (1986).
- List, Christian and Franz Dietrich. "Probabilistic Opinion Pooling" in *Oxford Handbook of Probability and Philosophy* (forthcoming).
- Levi, Isaac. *Decisions and Revisions: Philosophical Essays on Knowledge and Value*. Cambridge: Cambridge University Press (1984).
- Levi, Isaac. "Seeking Truth," in *Belief and Meaning: Essays at the Interface*, eds. W. Hinzen and H. Rott. Munich: Hansel-Hohenhaus (2002).
- Levi, Isaac. "Beware of Syllogism: Statistical Reasoning and Conjecturing According to Peirce," in *The Cambridge Companion to Peirce*, ed. Cheryl Misak. Cambridge: Cambridge University Press (2004).
- Lewis, David. "A Subjectivist's Guide to Objective Chance" in *Studies in Inductive Logic and Probability*, ed. Richard C. Jeffrey. Oakland: University of California Press (1980).
- Longino, Helen E. *Science as Social Knowledge: Values and Objectivity in Scientific Inquiry*. Princeton: Princeton University Press (1990).
- Longino, Helen E. "Subjects, Power and Knowledge: Description and Prescription in Feminist Philosophies of Science" in *Feminist Epistemologies*, eds. Linda Alcoff and Elizabeth Potter. New York: Routledge (1993).
- Longino, Helen E. "The Essential Tensions: Phase Two—Feminist, Philosophical, and Social Studies of Science" in *A Mind of One's Own: Feminist Essays on Reason and Objectivity*, eds. Louise M. Antony and Charlotte Witt. San Francisco: Westview Press (1993).
- Longino, Helen E. "The Fate of Knowledge in Social Theories of Science" in *Socializing Epistemology: The Social Dimensions of Knowledge*, ed. Frederick F. Schmitt. New York: Rowman & Littlefield (1994).
- Longino, Helen E. *The Fate of Knowledge*. Princeton: Princeton University Press (2002).

- Loux, Michael J. and Wm. David Solomon. "Quine on the Inscrutability and Relativity of Reference" in *Notre Dame Journal of Formal Logic*, Vol. XV, No. 1, January 1974 (1974), pp. 16-24.
- Maher, Patrick. "Diachronic Rationality" in *Philosophy of Science*, Vol. 59, No. 1, (Mar., 1992) (1992), pp. 120-141.
- Margolis, Joseph. *Reinventing Pragmatism: American Philosophy at the End of the Twentieth Century*. Ithaca: Cornell University Press (2002).
- Meacham, Christopher J. G. "Three Proposals Regarding a Theory of Chance" in *Philosophical Perspectives*, 19 (2005), pp. 281-307.
- Meacham, Christopher J. G. "Sleeping Beauty and the Dynamics of *De Se* Beliefs" in *Philosophical Studies* 138 (2) (2008), pp. 245-269.
- McDermott, John J, ed. *The Writings of William James: A Comprehensive Edition*. Chicago: University of Chicago Press (1978).
- McDowell, John. *Mind and World*. Cambridge: Harvard University Press (1996).
- Millstein, Roberta L. "Interpretations of Probability in Evolutionary Theory" in *Philosophy of Science* 70 (5), (2003), pp. 1317-1328.
- Nelson, Lynn Hankinson. "Epistemological Communities," in *Feminist Epistemologies*, eds. Linda Alcoff and Elizabeth Potter. New York: Routledge (1993).
- Nelson, Lynn Hankinson and Jack Nelson, eds. *Feminist Interpretations of W. V. Quine*. University Park: Penn State University Press (2003)
- Olsson, Erik, ed. *Knowledge and Inquiry: Essays on the Pragmatism of Isaac Levi*. Cambridge: Cambridge University Press (2006).
- Pihlström, Sami. "Peirce's Place in the Pragmatist Tradition," in *The Cambridge Companion to Peirce*, ed. Cheryl Misak. Cambridge: Cambridge University Press (2004).
- Peirce, C. S. "The Fixation of Belief" in *Popular Science Monthly* 12, November 1877, (1877), pp. 1-15
- Peirce, C. S. "How To Make Our Ideas Clear" in "Illustrations of the Logic of Science II," in *Popular Science Monthly*, Vol. 12, Jan. (1878).
- Peirce, C. S. *Values in a Universe of Chance: Selected Writings of Charles S. Peirce (1839-1914)*, ed. Philip P. Wiener, New York: Doubleday Anchor Books (1958).
- Price, Huw. "Naturalism and the Fate of the M-Worlds" in *Aristotelian Society Supplementary Volume* 71 (1) (1997), pp. 247-268.
- Price, Huw. "Naturalism Without Representationalism" in *Naturalism in Question*, eds. Mario De Caro and David Macarthur. Cambridge: Harvard University Press (2004).

- Price, Huw. "Truth as Convenient Friction," in *Naturalism and Normativity*, eds. Mario De Caro and David Macarthur. New York: Columbia University Press (2010).
- Price, Huw. *Naturalism Without Mirrors*. New York: Oxford University Press (2011).
- Putnam, Hilary. *Reason, Truth, and History*. Cambridge: Cambridge University Press (1981).
- Putnam, Hilary. *Realism with a Human Face*. Cambridge: Harvard University Press (1990).
- Putnam, Hilary. *The Collapse of the Fact/Value Dichotomy and Other Essays*. Cambridge: Harvard University Press (2002).
- Putnam, Ruth Anna. "Dewey's Epistemology," in *The Cambridge Companion to Dewey*, ed. Molly Cochran. Cambridge: Cambridge University Press (2010).
- Quine, W.V.O. *Ontological Relativity and Other Essays*. New York: Columbia University Press (1968).
- Quine, W.V.O. "Norms and Aims" in *The Pursuit of Truth*. Cambridge: Harvard University Press (1960).
- Quine, W.V.O. (1951). "Two Dogmas of Empiricism" in *From a Logical Point of View: Nine Logico-Philosophical Essays*. Cambridge: Harvard University Press (1980).
- Rorty, Richard. *Philosophy and the Mirror of Nature*. Princeton: Princeton University Press (1979).
- Rorty, Richard. *Contingency, Irony, and Solidarity*. Cambridge: Cambridge University Press (1989).
- Rorty, Richard. "Pragmatism, Davidson and Truth" in *Objectivity, Relativism, and Truth, Philosophical Papers, Vol. 1*, Cambridge: Cambridge University Press (1991).
- Rosenthal, Sandra. "Pragmatic Natures" in *The Future of Naturalism*, eds. John R. Shook and Paul Kurtz. New York: Humanity Books (Prometheus Books) (2009).
- Russell, Bertrand. (1908). "Transatlantic 'Truth'" in *Albany Review*, 2, 393-410, reprinted as "William James's Conception of Truth in *Philosophical Essays*. New York: Longman, Green, and Company (1910).
- Scharnhorst, Andrea and Eugene Garfield. "Tracing Scientific Influence" in *Dynamics of Socio-Economic Systems*, Vol. 2, No. 1, (2010), pp. 1-33.
- Scharnhorst, Andrea, Katy Börner, and Peter van den Besselaar. *Models of Science Dynamics: Encounters Between Complexity Theory and Information Science*. New York: Springer (2012).
- Schmitt, Frederick, ed. *Socializing Epistemology: The Social Dimensions of Knowledge*. New York: Rowman & Littlefield (1994).

- Sellars, Wilfrid S. "Empiricism and the Philosophy of Mind" in *Minnesota Studies in the Philosophy of Science* 1, (1956), pp. 253-329.
- Solomon, Miriam. "A More Social Epistemology" in *Socializing Epistemology: The Social Dimensions of Knowledge*, ed. Frederick F. Schmitt. New York: Rowman & Littlefield (1994).
- Solomon, Miriam. *Social Empiricism*. Cambridge: A Bradford Book: The MIT Press (2001).
- Stich, Stephen and Richard Nisbett. "Justification and the Psychology of Human Reasoning" in *Philosophy of Science* (1980), 47, pp. 188-202.
- Stich, Stephen. "Naturalizing Epistemology: Quine, Simon, and the Prospects for Pragmatism," in *Philosophy and Cognitive Science*. Cambridge: Cambridge University Press (1993).
- Stich, Stephen. *The Fragmentation of Reason: Preface to a Pragmatic Theory of Cognitive Evaluation*. Cambridge: The MIT Press (1990).
- Talbott, William J. "The Case for a More Truly Social Epistemology," in *Philosophy and Phenomenological Research*, Vol. LXIV, No. 1, January (2002), pp. 199-206.
- Talbott, William J. "Bayesian Epistemology," *The Stanford Encyclopedia of Philosophy* (Fall 2013 Edition). Edward N. Zalta (ed.) (2001, 2008 revised).
<http://plato.stanford.edu/archives/fall2013/entries/epistemology-bayesian/>
- Turrisi, Patricia Ann, ed. *Pragmatism as a Principle and Method of Right Thinking: The 1903 Harvard Lectures on Pragmatism*. Albany: State University of New York Press (1997).
- Van Fraassen, Bas C. "Belief and the Will" in *Journal of Philosophy*, 81, (1984), pp. 235-256.
- Van Fraassen, Bas C. "Belief and the Problem of Ulysses and the Sirens" in *Philosophical Studies*, 77, (1995), pp. 7-37.
- Weisberg, Jonthan. "Conditionalization, Reflection, and Self-Knowledge" in *Philosophical Studies*, 135, (2007), pp. 179-197.
- Weisberg, Jonathan. "Varieties of Bayesianism" in *Handbook of the History of Logic*, Vol. 10, eds. Dov Gabbay and Stephan Hartmann. Amsterdam: North-Holland Publishing Company/Elsevier (2011).
- Wheeler, Gregory and Jon Williamson, "Evidential Probability and Objective Bayesian Epistemology" in *Handbook of the Philosophy of Science, Vol. 7: Philosophy of Statistics*, eds. Prasanta S. Bandyopadhyay and Malcolm Forster. New York: Elsevier (2011).
- Williamson, Jon. *In Defence of Objective Bayesianism*. New York: Oxford University Press (2010).
- Williamson, Timothy. *Knowledge and Its Limits*. New York: Oxford University Press (2000).
- Williamson, Timothy. "What Is Naturalism?" *The New York Times*. [New York] September 4, 2011, Online Edition, The Opinionator:
<http://opinionator.blogs.nytimes.com/2011/09/04/what-is-naturalism>.

Wohlsen, Marcus. *Biopunk: DIY Scientists Hack the Software of Life*. New York: Penguin Group (2011).

Network Analysis Maps – Bibliographic Information

Boyack, Kevin W. and Richard Klavans. *The Structure of Science* in “1st Iteration (2005): The Power of Maps,” *Places & Spaces: Mapping Science*, eds. Katy Börner and Deborah MacPherson. <http://scimaps.org>

Boyack, Kevin W. and Richard Klavans. *Map of Scientific Paradigms* in “2nd Iteration (2006): The Power of Reference Systems,” *Places & Spaces: Mapping Science*, eds. Katy Börner and Deborah MacPherson. <http://scimaps.org> (2006).