

- Garber, D. 1983. Old evidence and logical omniscience in Bayesian confirmation theory. In Earman 1983, 99–132.
- Glymour, C. 1980. *Theory and evidence*. Princeton: Princeton University Press.
- Heck, R. 2003. Frege on Identity and identity statements. A reply to Thau and Caplan. *Canadian Journal of Philosophy* 33:83–102.
- Jeffrey, R. 1983. Bayesianism with a human face. In Earman 1983, 133–56.
- Larson, R., and P. Ludlow. 1993. Interpreted logical forms. *Synthese* 95:305–56.
- Leeds, S. 1979. Church's translation argument. *Canadian Journal of Philosophy* 9 (1): 43–51.
- Lewis, D. 1979. Attitudes *De Dicto* and *De Se*. *Philosophical Review* 88:513–43.
- Lewis, D. 1986. *On the plurality of worlds*. Oxford: Blackwell.
- Lewis, D. 1988. What experience teaches. *Proceedings of the Russellian Society* (University of Sidney) 13:29–57. Reprinted in Lewis 1990, 262–90.
- Lewis, D. 1990. *Papers in metaphysics and epistemology*. Cambridge: Cambridge University Press.
- Perry, J. 2001a. *Knowledge, possibility, and consciousness*. Cambridge, MA: MIT Press.
- Perry, J. 2001b. *Reference and reflexivity*. Stanford: CSLI Publications.
- Perry, J. 2003. Frege on identity, cognitive value, and subject matter. In *Building on Frege: New essays about sense, content, and concept*, ed. Newen, Nortmann, and Stuhlmann-Laeisz. Stanford: CSLI, 2003.
- Ramachandran, M. 1994. Frege's objection to the metalinguistic view. *European Review of Philosophy* 1:133–41.
- Skyrms, B. 1983. Three ways to give a probability assignment a memory. In Earman 1983, 157–61.
- Soames, S. 2005. *Reference and description. The case against two-dimensionalism*. Princeton: Princeton University Press.
- Stalnaker, R. 1978. Assertion. *Syntax and Semantics* 9. New York: Academic Press. Reprinted in Stalnaker 1999.
- Stalnaker, R. 1984. *Inquiry*. Cambridge, MA: MIT Press.
- Stalnaker, R. 1990. Mental content and linguistic form. *Philosophical Studies* 58. Reprinted in Stalnaker 1999, 225–40.
- Stalnaker, R. 1999. *Context and Content*. Oxford: Oxford University Press.
- Thau, M., and B. Caplan. 2001. What's puzzling Gottlob Frege. *Canadian Journal of Philosophy* 31:159–200.

The Southern Journal of Philosophy (2009) Vol. XLVII

A Euthyphronic Problem for Kitcher's Epistemology of Science

Jeffrey W. Roland
Louisiana State University

Abstract

Philip Kitcher has advanced an epistemology of science that purports to be naturalistic. For Kitcher, this entails that his epistemology of science must explain the correctness of belief-regulating norms while endorsing a realist notion of truth. This paper concerns whether or not Kitcher's epistemology of science is naturalistic on these terms. I find that it is not but that by supplementing the account we can secure its naturalistic standing.

Introduction

In a series of influential publications (1981, 1983, 1985, 1988, 1989, 1993), Philip Kitcher has advanced an epistemology of science that is remarkable for its putative applicability to mathematics as well as to the empirical sciences. That Kitcher's position purports to be naturalistic makes his combining the epistemology of mathematics with the epistemology of science generally all the more interesting. I have addressed Kitcher's epistemology of mathematics elsewhere (see Roland 2008b). Here I focus on Kitcher's epistemology of (nonmathematical) science—in particular, on whether or not Kitcher's epistemology of science is actually naturalistic. "Naturalism" means different things to different people.¹ I won't attempt to sort out which version of naturalism (if any) is the "correct" version. Instead I'll assess Kitcher's view in light of his own conception of naturalism (which he discusses generally in Kitcher 1983, 1988, and 1993, and most thoroughly in Kitcher 1992).

Jeffrey W. Roland is currently an assistant professor of philosophy at Louisiana State University. His research focuses primarily on issues at the intersection of epistemology, philosophy and foundations of mathematics, and philosophy of science.

Most important for present purposes is the commitment of Kitcher's naturalism to (KC₁) explaining the correctness of belief-regulating norms (i.e., norms governing belief formation and retention and (KC₂) a realist notion of truth, in the sense that what makes a statement (belief, etc.) true or false is independent of us and our cognitive activities (*modulo* statements about us and our activities. Call these *Kitcher's commitments*. I'll argue that Kitcher's epistemology of science fails to simultaneously keep both of these commitments but that this situation can be remedied by modifying his account.

I begin with a brief defense of my claim that Kitcher's naturalism is committed to KC₁ and KC₂ (§1). I then set out the central parts of Kitcher's epistemology of science (§2) and follow this by considering whether or not Kitcher's epistemology of science satisfies KC₁ and KC₂ (§§3 and 4). Given the particulars of Kitcher's position, the key question here concerns the (approximate) truth-conduciveness of unification. I find no reason to think that Kitcher's account can sustain the desired connection between unification and truth, and Kitcher's problems in this regard are brought into relief by considering their Euthyphronic dimension (§5) (see Wright 1992). I next offer a way of supplementing Kitcher's account so as to yield a unification-based, naturalistic epistemology of science (§6), but in a way Kitcher is not likely to endorse. I end (§7) by noting a possible confusion between epistemic and nonepistemic concepts of theoretical unification in an effort to pre-empt a potential objection.

1. Kitcher's Commitments

I argue that Kitcher is committed to KC₁ and KC₂ at some length elsewhere (Roland 2008b). Here I briefly revisit those arguments, rather than offer them in full.

On KC₁: Kitcher holds that "[t]he central problem of epistemology" according to naturalism "is to understand the epistemic quality of human cognitive performance, and to specify strategies through whose use human beings can improve their cognitive states" (Kitcher 1992, 74–75). We cognitively improve by deploying reliable cognitive processes, belief-forming and sustaining processes that tend toward true beliefs. Addressing the central problem of epistemology is a matter of "describing [cognitive] processes that are reliable" (75–76). So Kitcher's naturalism is concerned with the reliability (truth-conduciveness) of belief-forming and sustaining processes. Witness Kitcher's professed affinity for the epistemological views of Alvin Goldman (1976, 1979, 1986; see e.g., Kitcher 1992, 65) and Hilary Kornblith (1980; see e.g., Kitcher 1983, 18).

As understanding is the essence of explanation for Kitcher (1976, 1981, 1985, 1989), and satisfactorily addressing the

"central problem" of naturalistic epistemology requires understanding the reliability of belief-forming and sustaining processes, Kitcher's naturalism is committed to explaining the reliability of these processes. But such a commitment is also a commitment to explain the correctness of our belief-regulating norms. The norms governing belief-forming and sustaining processes are implicit in the processes themselves. We (typically tacitly) take conditions of application into consideration when forming beliefs and "deciding" whether or not to maintain them, and in so doing we (typically tacitly) deploy belief-regulating norms. The correctness of these norms depends on the reliability of the processes they ratify: a belief-regulating norm is correct to the extent that the belief-forming and sustaining processes it ratifies are reliable.²

On KC₂: Kitcher's affinity for reliabilism, especially as espoused by Goldman, provides a way of seeing Kitcher's commitment to KC₂. Goldman regards a realist conception of truth as essential to his epistemology (see 1986, ch. 7). He argues that it is crucial to his view that "truth ... not be an epistemic matter" in the sense that "what makes a [true (or false) statement] true (or false) is independent of our knowledge or verification" (143). Goldman endorses *verification-transcendent truth* according to which the truth value of a statement is independent of "our knowledge, or verification, of it (or even our *ability* to verify it)" (143; original emphasis).³ The realist conception of truth in play in KC₂ is precisely this verification-transcendent conception of truth.

Notice that this gives us a tight connection between KC₁ and KC₂. By Kitcher's own lights, the correctness of belief-regulating norms is understood in terms of the reliability of the belief-forming and sustaining processes ratified by those norms, and reliability requires a realist (verification-transcendent) conception of truth. So only a realist conception of truth is adequate for KC₁. In other words, given the way KC₁ gets cashed out for Kitcher, commitment to it results in commitment to KC₂.

2. Kitcher's Epistemology of Science

Kitcher frames the epistemology of science as an accounting of how we attain our epistemic goals, namely, "the goals of achieving a perspicuous description of nature and of attaining significant [i.e., nontrivial] truth" (Kitcher 1995b, 614). He "conceive[s] of science as a sequence of practices" (Kitcher 1989, 497), where a scientific practice⁴ *P* is represented by an ordered *n*-tuple, (L^P, K^P, Π^P, \dots) — L^P is the language employed by the scientific community in carrying on research, K^P is the set of statements about nature accepted by the scientific community, and Π^P is the set of explanatory schemata accepted

by the scientific community for the purpose of explaining members of K^P .⁶ If science is (represented by) the sequence of practices P_0, P_1, P_2, \dots , then the limit of scientific inquiry is the practice P_∞ toward which this sequence would tend were it to grow indefinitely into the future regulated by the norm(s) of acceptable inter-practice transition. Call P_∞ the *limit practice*.⁶

The limit practice figures prominently in Kitcher's epistemology of science:

Consider science as a sequence of practices that attempt to incorporate true statements (insofar as possible) and to articulate the best unification of statements (insofar as possible). As this sequence proceeds, certain features of the organization of beliefs may stabilize: predicates of particular types may be used in explanatory schemata and employed in inductive generalization; particular schemata may endure (possibly embedded in more powerful schemata). The "joints of nature" and the "objective dependencies" are the reflections of these stable elements. The natural kinds would be the extensions of the predicates that figured in our explanatory schemata and were counted as projectible in the limit as our practices developed to embrace more and more phenomena. Objective dependencies would be those recorded in the schemata that emerged in the limit of our practices. (Kitcher 1993, 172)

In recording objective dependencies in nature, correct explanatory schemata record the causal structure of the world: "The growth of science is driven in part by the desire for explanation, and to explain is to fit the phenomena into a unified picture insofar as we can. What emerges in the limit of this process is nothing less than the causal structure of the world" (Kitcher 1989, 500). Thus, according to Kitcher, the limit practice identifies natural kinds, describes the causal structure of the world, and yields significant truth.

Kitcher endorses a Kantian view of scientific progress according to which "the growth of scientific knowledge is governed by a principle of unification. Modifications of [practice] are correctly supported by pointing out that they would lead to a system of belief that is more unified" (Kitcher 1993, 172). Thus, for instance, if $P = \langle L^P, K^P, \Pi^P, \dots \rangle$ and $P' = \langle L^{P'}, K^{P'}, \Pi^{P'}, \dots \rangle$ are practices, with $L^P = L^{P'}$ and $K^P = K^{P'}$, and the explanations provided by P are more unifying than those provided by P' , then the transition from P to P' is acceptable or well supported.⁷ So the central norm regulating inter-practice transitions is the directive to unify.

Whether or not Kitcher's epistemology of science satisfies KC_1 and KC_2 depends on whether or not unification is (at least approximately) truth-conducive, where the relevant notion of truth is realist. If unification is truth-conducive, then

Kitcher's epistemology of science provides an explanation of the correctness of the prime belief-regulating norm of science (on Kitcher's view), the directive to believe the deliverances of your best (i.e., most unifying) theory.

3. Unification, Truth, and Explanatory Schemata

In considering the truth-conduciveness of unification we restrict our attention to the question of the (approximate) truth of causal pronouncements of P_∞ , that is, of causal statements in K^{P_∞} .⁸ After all, if the causal pronouncements of P_∞ fail to be (at least approximately) true, that's a sufficiently damning indictment of the truth-conduciveness of unification in the current circumstances. Given this, in order to vindicate Kitcher's epistemology of science we need an argument for

Vindicating Claim (VC). A causal statement σ is true just in case $\sigma \in K^{P_\infty}$.⁹

that respects KC_2 .¹⁰ The question is whether or not Kitcher can give such an argument. We can sharpen this question.

An argument for VC will be based on one of the following:

Base Claim 1. A causal statement σ is true just in case $\sigma \in K^{P_\infty}$ because a causal statement σ 's being in K^{P_∞} , that is, having a settled belief in σ at the end of scientific inquiry governed by the directive to unify, is constitutive of truth for statements like σ (viz., causal statements).

Base Claim 2. A causal statement σ is true just in case $\sigma \in K^{P_\infty}$ because the directive to unify is such that, for any true causal statement σ , following this directive will result in our having a settled belief in σ at the end of scientific inquiry (i.e., will result in $\sigma \in K^{P_\infty}$).

Now, since for Kitcher true causal statements are descriptions of objective dependencies in nature and such dependencies are captured by correct explanatory schemata (i.e., those schemata in Π^{P_∞}), we can reframe Base Claim 1 and Base Claim 2 in terms of correct explanatory schemata.

Correct Explanatory Schemata 1. A causal statement σ is true just in case $\sigma \in K^{P_\infty}$ because being a correct explanatory schemata—that is, being in Π^{P_∞} , the set of explanatory schemata at the end of scientific inquiry governed by the directive to unify—is constitutive of being an objective dependency in nature.

Correct Explanatory Schemata 2. A causal statement σ is true just in case $\sigma \in K^{P\infty}$ because there are objective dependencies in nature and, as a matter of fact, the directive to unify is such that those dependencies are captured by correct explanatory schemata at the end of scientific inquiry governed by that directive.

So arguing for VC requires Kitcher to accept either Correct Explanatory Schemata 1 or Correct Explanatory Schemata 2, and the question is whether or not he can accept either one while honoring KC_1 and KC_2 .

4. Kitcher on VC

Once we have Correct Explanatory Schemata 1 and Correct Explanatory Schemata 2, we can see that Kitcher more or less directly addresses VC. He intends his account of scientific progress to be neutral between Correct Explanatory Schemata 1 and Correct Explanatory Schemata 2. Citing his 1993 (169–73), he says that he remains “officially agnostic between the invocation of mind-independent non-Humean causation in nature, and a Kantian alternative, in which the causal structure of the world is that projected by our explanatory schemata in the limit of our attempts to unify the phenomena” and that it’s an option to “retain the non-Humean realist view about causation and contend that employing the principle of unification [is] the right way to identify the non-Humean causes” (1995a, 661).

This raises the issue of the (realist) truth-conduciveness of unification, that is, of (a realist version of) VC: “Why should non-Humean causes be disclosed by trying to unify the phenomena?” (Kitcher 1995a, 661). Kitcher entertains three answers to this question (see 1995a, 661).

- (I) Causal relations just are what result from our attempt to unify the phenomena, in the limit of inquiry.
- (II) A search for unification will disclose non-Humean causes.
- (III) We can’t rely on the principle of unification to disclose the genuine [i.e., non-Humean] causes.

Answers (I) and (II) correspond to Correct Explanatory Schemata 1 and Correct Explanatory Schemata 2, respectively, while (III) is essentially the denial of (II).

4.1 Accepting Correct Explanatory Schemata 1

Kitcher occasionally endorses (I) (i.e., Correct Explanatory Schemata 1). This is hinted at in the last sentence of the last

block quotation above and is more explicitly on display in this passage:

As the result of the demands of inquiry [namely, to unify the phenomena], demands that are not separable from the naive suggestion that science aims at the truth, we are prepared to project an order of nature. Laws are statements that play a particular role in the system that would emerge from an ideally extended inquiry. *Parallel characterizations apply for kinds and causal dependencies.* (Kitcher 1986, 214–15; emphases added)

According to the view expressed in this quotation, scientific inquiry proceeds under a directive to unify and truth is determined by what happens in the limit of inquiry; the attainment of truth via scientific theorizing cannot be separated from the directive to unify. So instead of laws, kinds, and causes that are independent of and revealed by our theorizing, we have laws, kinds, and causes that are the result of (and so dependent on) our theorizing, and by imparting these to nature we “project an order of nature” that is otherwise lacking. Hence, on this view, Correct Explanatory Schemata 1 holds and causal statements in $K^{P\infty}$ are true because being in $K^{P\infty}$ is what *makes* a statement true.

This conception of truth enables Kitcher’s account to at least nominally satisfy KC_1 . However, it plainly doesn’t satisfy KC_2 , as it obviously fails to be verification transcendent. On it, being true depends on belonging to a set the membership of which is determined by our epistemic practice, in particular by our best attempt at unification in the long run. As observed by Richard Miller (1995, 642), this is essentially a Peircean, pragmatist conception of truth for causal statements; simply being among the pronouncements of the limit practice makes a causal statement true.¹¹ But this pragmatic conception of truth is clearly inadequate for KC_2 . Indeed, any conception of truth for causal statements grounded in Correct Explanatory Schemata 1 is inadequate for KC_2 , since any such conception of truth will fail to be verification transcendent. So Kitcher can’t both accept Correct Explanatory Schemata 1 and honor KC_2 . It follows that Kitcher must reject (I) to maintain KC_2 .¹²

4.2 Arguing for Correct Explanatory Schemata 2

Let’s take stock. We had the disjunction of (I)–(III). Kitcher is most concerned to argue for the legitimacy of agnosticism between (I) and (II), but we’ve just seen that Kitcher must reject (I) in order to maintain KC_2 . This leaves us with the disjunction of (II) and (III). If (III) were correct Kitcher’s account would violate KC_1 , since then unification would not track causal truth and so would fail to be reliable. So Kitcher needs an argument

for (II), which he attempts by arguing against (III), thus undercutting what he takes to be the main motivation for rejecting (II). If this succeeds, then Kitcher has gone some way toward establishing (II), hence, also Correct Explanatory Schemata 2 and with it the naturalistic status of his epistemology of science. I contend that this attempt fails.

Kitcher argues against (III) indirectly, by trying to undermine a chief motivation for rejecting (II), namely, the so-called obsessive unifier. The worry is that if unification is our primary aim, then in our zeal to achieve this aim we might impose unity where there in fact is none. The result: At the end of the day our causal "truths" won't be true at all. In other words, unification needn't track non-Humean causes; (II) is false. This worry is explicitly addressed by Kitcher in two places.

In the first, he traces the intuition underlying the worry to David Lewis's challenge that the world might not be unified and articulates that challenge using the following three claims.¹³

- (Unif₁) On the unification approach to explanation, a factor F is explanatorily relevant to a phenomenon Φ if and only if there is a derivation δ , which instantiates a schema in $\Pi_{E(K^{\infty})}$, such that the conclusion of δ is (a description of) Φ and some premise of δ refers to F .
- (Unif₂) F is causally relevant to Φ only if F is explanatorily relevant to Φ .
- (PNU) It is possible (both metaphysically and epistemically) that there are F and Φ such that F is causally relevant to Φ and no derivation of (a description of) Φ instantiating a schema in $\Pi_{E(K^{\infty})}$ contains a premise referring to F .

Since Unif₁ and Unif₂ are taken to be necessary on the unification approach developed by Kitcher, PNU (which expresses the possibility that the world is not unified) must be false on that approach. After all, PNU claims that explanatory relevance and there being certain derivations come apart. But Unif₁ and Unif₂, if necessary, rule out that separability. Prima facie, the falsity of PNU doesn't seem very plausible, so Kitcher needs to say something to counter this implausibility.

He defends the claim that PNU is false by "rejecting the idea that there are causal truths that are independent of our search for order in the phenomena" and "[adopting] a different view of truth and correctness" (Kitcher 1989, 497). However, the view of truth he adopts is precisely the Peircean one that we've already rejected as inadequate for KC₂. But then it's clear that this position endorses (I) and rejects (II), and so both is not neutral between (I) and (II) and embraces a position that has already been set aside as inadequate.¹⁴

Kitcher (1995a) again addresses the obsessive unifier worry. There he indicates that he has "tried to scotch [this misunderstanding] elsewhere" (661). The "elsewhere" he cites concludes a discussion of the obsessive unifier with the revelation that "there is a sense in which the unification approach, in its new guise, can tolerate the possibility that the world might prove to be a messy place," because "while my version of the unification approach makes it constitutive of explanatory relevance that there be no basic explanatory (or causal) mechanisms that are not captured in the limit of attempts to systematize beliefs, it does *not* make it a necessary truth that this limit achieves any particular degree of unification" (Kitcher 1989, 499; emphasis added).

Let's grant for the sake of argument that the Kantian conception of scientific progress, according to which we in some sense project the order of nature, is not essential to the unification approach. In effect, this concedes to Kitcher the points made against his first response to the obsessive unifier worry and allows us to consider his (1995a) response free from that baggage. According to Kitcher (1995a), we needn't worry about the obsessive unifier because

The practitioners recognize that there are no true instantiations of a single schema that will cover all the cases, so they are forced to invoke many schemata if the entire range of phenomena is to be captured. Unification enjoins them to find the minimal set of schemata for generating all the phenomena; it neither insists that a single schema should be applied ... nor does it set any a priori limit to the size of the set of schemata that should prove necessary. The unificationist concedes that we might live in a very disorderly world, and asks only that we discover the minimal number of schemata for coping with it. (Kitcher 1995a, 661–62)

Given this response, the following two questions need to be answered. First, what grounds the ability of practitioners to tell that no single schema will do? Second, why should a conception of unification that is neutral with respect to how much unification we require stand a chance at being truth-conducive? I've argued elsewhere (2008a) that Kitcher's account does not provide adequate grounds on which to rule out that a single schema is sufficient to unify the full range of phenomena. I here restrict my attention to the second question.

If we concede that a practice needn't be unified to any particular degree in order to count as unifying the phenomena, then it seems that in our drive to unify we might well avoid the zealotry that gives rise to the vision of the obsessive unifier. This much seems right. But recall that we are currently engaged in trying to answer the worry that unification might not be a very good guide to (non-Humean) causal truth, in trying to undermine

(III) in favor of (II) in order to establish that a causal statement is true just in case it's in $K^{P\infty}$ (VC), where the operative conception of truth is verification transcendent (KC_2). If we allow a Kantian conception of scientific progress, with its Peircean conception of truth, then this worry can be addressed, but only at the price of violating KC_2 . If the Kantian move is rejected, then we're still at square one: we're looking for some reason to prefer (II) to (III), that is, we're looking for some reason to think that the causal pronouncements of the limit practice are likely to be (approximately) true in a realist sense. The problem is that if a practice needn't do much in the way of unifying phenomena to count as a unifying practice, it's hard to see what reason there could be for thinking that the limit practice is likely to yield (approximate) verification-transcendent truths. Kitcher offers no such reason. Indeed, when pressed on the connection between unification and truth he tends to revert to the Kantian approach.

Kitcher considers, without advocating, that one might hold that unification is (realistically) truth-conducive on the basis of the claim that nature is *in fact* systematically unified (as opposed to *necessarily* systematically unified) (Kitcher 1986, 211–12). The idea is that if it's truth about the world we're after in our theorizing, then the best (perhaps only) reason to think that theories constructed under the directive to unify are likely to be true is if the world is, in fact, unified. Put another way, since we read our theory off the world, pride of place for unified theories only makes sense if the world is unified. But even if one admits that the world is unified, why think that a theory that makes only slight inroads into that unified structure is likely to be (approximately) true? If nature is systematically unified and a practice P is unifying, but only slightly, it seems that P would likely get more—many more—things wrong than it would get right. Given this, do we really want to say that unification is truth-conducive? I think not. If unification is going to be truth-conducive, we need a much higher standard for being a unifying practice or theory than the one Kitcher sets (or, more to the point, doesn't set) in his attempt to ward off the bogey of the obsessive unifier.

4.3 A Coherentist Rejoinder

It might be tempting to invoke some sort of coherentism on Kitcher's behalf, the idea being that a unifying practice P will be one in which the set of statements K^P coheres and so, since coherence is truth-conducive, is likely to yield an approximately true account of the world. There are at least two problems with this strategy. First, it relies on theoretical unification yielding coherence of the right sort. Unification certainly yields some sort of coherence, but whether or not it's the sort that leads to

truth is an open question. This is so largely because it's unclear what kinds of coherence properties a system of beliefs has to have in order to be on balance sufficiently likely to be true. This leads to the second problem: The strategy under consideration relies on the claim that coherence is truth-conducive. This is a substantive claim, and one that has proved quite difficult to maintain (see, e.g., BonJour 1999). One could, of course, opt to be a thoroughgoing coherentist, adopting a coherence theory of truth in addition to the coherence theory of justification under consideration. But a coherence conception of truth violates KC_2 in the same way a Peircean, pragmatist conception of truth does: on both views what is true depends on our theorizing and so is not verification transcendent. For this and other well-known reasons,¹⁵ a coherence theory of truth is unacceptable.¹⁶ But, absent a coherence theory of truth, one who would defend Kitcher by an appeal to coherentism owes us a defense of the claim that coherence is truth-conducive as well as a defense of the claim that unification yields the right sort of coherence, both of which are highly nontrivial.

5. Kitcher and the Euthyphro Contrast

Kitcher's predicament with respect to VC can be illuminated by considering it in terms of what Wright (1992) calls the *Euthyphro contrast*. Equivalences such as VC can be given either a Socratic or a Euthyphronic construal. Recall that in Plato's *Euthyphro*, Euthyphro characterizes holiness as follows: An act is holy just in case it's pleasing to the gods. Call this *Euthyphro's Characterization*. Socrates responds to Euthyphro's Characterization by asking: Are holy acts holy because they please the gods, or do such acts please the gods because they are holy? (Hamilton and Cairns 1963, 174–79). Call this *Socrates' Question*. One who answers Socrates' Question by affirming that holy acts please the gods because they are holy, that is, by attributing an ability to detect holiness and a disposition to respond with pleasure to the gods, construes Euthyphro's Characterization Socratically. On the other hand, one who answers Socrates' Question by affirming that holy acts are holy because they please the gods, that is, by attributing the power to make an act holy to the gods (more specifically to their pleasure), construes Euthyphro's Characterization Euthyphronically.

As with Euthyphro's Characterization, VC admits both Socratic and Euthyphronic construals. Indeed, Base Claim 1 gives the Euthyphronic construal of VC and Base Claim 2 gives the Socratic construal of VC. In these terms, then, the crucial question for Kitcher is whether or not he can argue for either construal of VC without violating KC_2 .¹⁷

It's not hard to see that the Euthyphronic construal of VC (i.e., Base Claim 1) must be rejected by Kitcher. In the first

place, a Euthyphronic construal of VC is an antirealist construal and, hence, gives us a conception of truth that violates KC_2 . To paraphrase Wright, constitutive independence is what is wanted for realist truth (Wright 1992, 81), and the Euthyphronic construal of VC posits constitutive *dependence* of truth. In the second place, a Euthyphronic construal of VC would make VC a priori. On a Euthyphronic construal, VC is an analytic definition of the concept of truth for causal statements; it tells us in what being true consists for a causal statement. Analytic definitions are paradigmatically a priori. But Kitcher explicitly repudiates a priori epistemological principles: "Virtually nothing is knowable a priori, and, in particular, no epistemological principle is knowable a priori" (Kitcher 1992, 76). VC is clearly an epistemological principle. So Kitcher can't accept that it's a priori. Since it's a priori on the Euthyphronic construal, Kitcher can't avail himself of that construal of VC.

This leaves the Socratic construal of VC (i.e., Base Claim 2). Kitcher can't just help himself to the Socratic construal. The chief question is whether or not Kitcher can argue for VC while respecting KC_2 . With Base Claim 2 this question is clearly answered affirmatively, but then simply assuming that Kitcher is entitled to a Socratic construal amounts to begging the question in Kitcher's favor. So the operative issue is whether or not Kitcher can argue for Base Claim 2 without violating KC_2 . Let's grant for the sake of argument that the conception of truth deployed in Base Claim 2 suffices for KC_2 . (This is already implicit in the Socratic construal, but it doesn't hurt to be explicit.) According to Base Claim 2, the principal regulative norm of scientific inquiry (by Kitcher's lights) is such that by following it scientists will, in the long run, come to have true causal beliefs. What we need is an argument for this long-run success at tracking causal truth by unification. But this is exactly what we needed when we formulated VC, namely, an argument for the correctness of the belief-regulating norms of science. This raises a slightly different issue regarding question begging. Respecting KC_1 requires explaining the correctness of the belief-regulating norms of science, which in turn requires explaining the reliability of cognitive processes licensed by those norms. Kitcher's explanation of this reliability depends on an argument for VC, and in light of Kitcher's commitment to KC_2 this requires an argument for Base Claim 2 (i.e., an argument that Kitcher is entitled to the Socratic construal of VC). But the explanatory challenge raised by KC_1 arises anew in Base Claim 2: Base Claim 2 simply asserts the correctness of the belief-regulating norms of science, rather than explaining it. Kitcher's epistemology of science appears able to respect KC_2 only by giving up KC_1 (equivalently, respecting KC_1 appears to require sacrificing KC_2).

6. Naturalizing Kitcher

If we confine ourselves to the resources at Kitcher's disposal, it appears we can't simultaneously secure KC_1 and KC_2 . However, by supplementing Kitcher's epistemology of science we can remedy this situation. The idea is that the pronouncements of unifying theories are likely to be (approximately) true provided that the world is systematically unified and we require a unifying theory or practice to meet a sufficiently high standard of unification. So we adopt the view that nature is unified and give a direct argument in favor of (II).

As noted in §3, Kitcher (1986) considers something similar. Though he takes his epistemology of science to be consistent with the view that nature is systematically unified, Kitcher (1993, 169–73) finds this view, which he terms *strong realism*, epistemically problematic.

Ever since Hume, philosophers have faced the challenge of explaining how we are in a position to gain evidence for statements involving a family of notions—statements that identify causal relationships, statements that talk of objective explanatory dependence, statements that assert that a particular set of objects is a natural kind, statements that talk of natural necessities. The root problem seems to be that we have no semantical account of such statements that will fit into an epistemological account. (Kitcher 1993, 170)

So, according to Kitcher, strong realism gives us an access problem of the sort familiar from the Benacerraf problem for mathematical knowledge.¹⁸ Strong realism enables us to tell a nice story about the semantics of scientific discourse, a story that facilitates an account of the reliability of the modes of justification operative in science. But this story comes at the expense of a satisfying epistemology, since this requires some account of how we access the strongly realist world. This is what motivates Kitcher's preference for (I). But how apt is this comparison between the Benacerraf problem and our present situation?

The Benacerraf problem relies on (the apparent mismatch between) a causal theory of knowledge and the acausal nature of mathematical entities. If causal interaction of some sort is required for knowledge and we never causally interact with mathematical entities, then it looks as if mathematical knowledge is impossible. Humean worries about causal and other sorts of scientific knowledge are perhaps not so different in kind, but they seem very different in degree. Few will look askance at the claim that we are never in causal contact with mathematical entities; it's not so clear that the same can be said for the claim that we never causally interact with, say, natural kinds. Indeed,

naturalists are able to account for the reliability of scientific methods and, hence, the correctness of belief-regulating norms, *precisely because* they accept a primitive, realist notion of causation that underwrites reference to natural kinds (see Boyd 1980, 1989, 1991). This allows them access to the causal order of the world. The causal order of the world forms the *ontological ground* of inferential (inductive) practice in science. Consequently, kinds align with the causal order so that outcomes of induction on natural kinds tend to be (approximately) true.

Moreover, a causally ordered world is effectively a unified world. We can understand the claim that nature is systematically unified (roughly) as the claim that many events and states of affairs in nature that appear to be of different types are actually of the same type, so that the world is much less diverse than it initially appears to be. Kitcher's approach to explanation in science, and to the epistemology of science generally, aims at unifying the world from the top down by establishing relations of explanatory dependence from which we then read off the causal order of the world. But there is also a "bottom up" approach to the epistemology of science, and scientific explanation in particular, according to which causal order precedes and underwrites explanation and induces relations of explanatory dependence between events and states of affairs in the world. The unification of the world can then be read off of these explanatory dependencies.¹⁹ Of course, things aren't quite this simple, but the general idea should be clear: causal order and unification are two routes to the same place. So causal realism adds nothing substantive to the view that nature is unified.

Kitcher's epistemology of science proves unsatisfactory from the point of view of his own naturalism because its only way of linking unification and truth, so as to satisfy KC_1 , seems to be by interpreting truth in such a way that it fails to be verification transcendent and so violates KC_2 . Kitcher recognizes that he needs a realist conception of truth (hence his commitment to KC_2), but he can't see his way clear to adopt a sufficiently realist metaphysics to ground the requisite realist conception of truth. We can satisfy both KC_1 and KC_2 for Kitcher by adding to his epistemology of science a premise stating that the world is unified (or, alternatively, by adopting causal realism), but it's not clear that Kitcher would endorse this strong realist move. I haven't attempted to give substantial arguments for this additional premise. But the resources available for use in arguing for it—indeed arguments for it available in the literature, especially in its causal realism form—far outstrip those available for use in arguing for Base Claim 2 (the Socratic construal of VC).²⁰ This recommends the approach outlined in this section over Kitcher's as a way of vindicating a unification-based naturalism in the epistemology of science.

7. Concluding Remarks

In closing, I would like to address a potential point of confusion. Much of the foregoing discussion centers on the idea that unification is a regulative norm with epistemic import, in the sense of unifying theories tending to yield (approximate) truths. In short, if unification is an *epistemic* regulative norm, then a unifying theory T should be preferred over nonunifying or less-unifying theories. We've seen that unification is an epistemic norm in the presence of substantive metaphysical theses; unifying theories are more likely to be (approximately) true only when appropriate ontological grounds are available. Absent appropriate ontological grounds, an epistemology governed by a norm to unify will be in a poor position to explain the correctness of belief-forming and sustaining norms. There is, however, a sense in which unification figures in a regulative norm governing theory choice independent of metaphysics.

All else being equal, we should (and do) prefer theories that are simpler in the sense of being more cognitively tractable; the more difficult a theory is to understand, the less desirable it is. Very often the cognitive tractability of a theory is a matter of how well that theory integrates with extant well-established theories. One way a theory integrates well with established extant theories is by being presented in terms that are similar to those used in presenting the established extant theories. Similarity breeds familiarity, and theories that are cast in familiar terms are easier to grasp than those cast in unfamiliar terms. This sort of integration is a type of unification. If T is an established theory and T' integrates well with T , then $T + T'$ is unified in a way that $T + T''$, where T'' does not integrate well with T , is not. Thus, if T_1 and T_2 are theories concerning some subject S that are well confirmed and differ only in that T_1 is presented in familiar terms while T_2 isn't, then, by the foregoing remarks, we should prefer T_1 to T_2 . And this preference is the outcome of a regulative norm according to which we should accept the theory that exhibits greater unification with established extant theories than its competitors.

What's important to notice here is that this regulative norm is not an *epistemic* norm. Rather, it's a conventional or pragmatic norm that only enters into the process of theory choice after evaluations of rival theories on epistemic grounds have already concluded (in our example, T_1 and T_2 were both well confirmed *ex hypothesi*). Hence, the fact that unification plays a role in this (nonepistemic) regulative norm of theory choice in no way undermines my arguments against the purportedly epistemic regulative norm that has been the chief concern of much of this paper. That norm has the potential to be an epistemic norm, but, as I've argued, realizing this potential depends on the presence of an appropriately robust ontology. Kitcher

seems to want to avoid commitment to such an ontology, but without it his claim to an epistemology of science that satisfies KC_1 and KC_2 appears to fail. Consequently, so too does his claim to an epistemology of science that is naturalistic by his own lights.²¹

Notes

¹ For a sample of comparative views, see King 1994, Kornblith 1994, Maddy 2005, Rosenberg 1996, and Stroud 1996.

² Given that we aim to have true beliefs.

³ For more on the shortcomings of this and similar (verificationist or epistemic) conceptions of truth, see Goldman 1986, ch. 7. For an extended discussion of this issue, see Wright 1992.

⁴ In this paper, practices are what Kitcher calls *consensus practices* (see Kitcher 1993, 87–89).

⁵ Kitcher's practices sometimes consist of only these three components (see, e.g., Kitcher 1989). Other times (e.g., Kitcher 1983, 1993) Kitcher's practices incorporate additional components. For present purposes, these three will be sufficient.

⁶ For simplicity, I assume that there is a unique limit practice of a given sequence of practices.

⁷ For details on Kitcher's account of comparative unifying power, see Kitcher 1981, 1989.

⁸ See Boyd 1990 for a realist defense of approximate truth.

⁹ I ignore possible worries about the completeness of the limit practice, i.e., about whether or not all causal truths make it into $K^{P\infty}$.

¹⁰ VC itself gives Kitcher KC_1 .

¹¹ See also Kitcher 1986, 213–14 and 1989, 497–98. Kitcher (1986) appears to endorse this conception of truth for all empirical statements, not just causal statements.

¹² I ignore a problem often raised for views that claim something we want happens in the long run or at the end of inquiry, viz., the problem of connecting what happens in the long run with where we stand now. In Kitcher's case, the challenge would take the following form: Tell us why getting things (approximately) right in the long run should give us confidence that we have things (approximately) right along the way, in particular right now.

¹³ See Kitcher 1989, 498–99. I paraphrase and update the claims appearing there.

¹⁴ Even if we ignore that (I) has already been rejected, this argument begs the question against the proponent of (III). Kitcher argues against (III) with the intent of establishing the disjunction of (I) and (II), but in so doing he deploys one of the disjuncts in play, viz., (I).

¹⁵ See, e.g., those canvassed in chapter 3 of Kirkham 1992.

¹⁶ Note that even leading coherentists about justification reject coherence truth; indeed, they endorse a realist conception of truth. (I have in mind Keith Lehrer and Laurence Bonjour. See, e.g., Lehrer 1974, 1986 for the former's views. Bonjour now rejects coherentism about justification, but for his coherentist views see, e.g., Bonjour 1985.

¹⁷ In what follows, I rely on Wright's discussion of Euthyphronic contrasts in Wright 1992, especially ch. 3 and its appendix, though I

don't follow that discussion in every detail.

¹⁸ Kitcher comments on this similarity to the Benacerraf problem. See Kitcher 1993, 170, n. 57.

¹⁹ For more on this, see Kitcher 1985, 1989. The "top down" and "bottom up" characterization of the situation comes from the former.

²⁰ See, in particular, Boyd 1973, 1979, 1980, 1985a, 1985b, 1990, 1991 and Kornblith 1993.

²¹ Thanks to Richard Boyd, Jon Cogburn, Kevin Elliot, Harold Hodes, Richard A. Shore, Zoltán Szabó, and an anonymous referee for helpful discussions and comments.

References

- BonJour, Laurence. 1985. *The structure of empirical knowledge*. Cambridge, MA: Harvard University Press.
- . 1999. The dialectic of foundationalism and coherentism. In *The Blackwell guide to epistemology*, ed. John Greco and Ernest Sosa, 117–42. Malden, MA: Blackwell Publishing.
- Boyd, Richard. 1973. Realism, underdetermination, and a causal theory of evidence. *Nous* 7:1–12.
- . 1979. Metaphor and theory change. In *Metaphor and thought*, ed. Andrew Ortony, 481–532. Cambridge: Cambridge University Press.
- . 1980. Scientific realism and naturalistic epistemology. *Proceedings of the Philosophy of Science Association*, Vol. 2. Philosophy of Science Association, 613–62.
- . 1985a. Lex Orandi est Lex Credendi. In *Images of science: Scientific realism versus constructive empiricism*, ed. Paul M. Churchland and Clifford A. Hooker, 3–34. Chicago, IL: University of Chicago Press.
- . 1985b. Observations, explanatory power, and simplicity: Toward a non-Humean account. In *Observation, experiment, and hypothesis in modern physical science*, ed. Peter Achinstein and Owen, 47–94. Cambridge, MA: MIT Press.
- . 1989. How to be a moral realist. In *Contemporary materialism: A reader*, ed. Paul K. Moser and J. D. Trout, 297–343. London: Routledge.
- . 1990. Realism, approximate truth, and philosophical method. In *Scientific theories. Minnesota studies in the philosophy of science*, Vol. 14, ed. C. Wade Savage, 355–91. Minneapolis: University of Minnesota Press.
- . 1991. Realism, anti-foundationalism, and the enthusiasm for natural kinds. *Philosophical Studies* 61:127–48.
- Goldman, Alvin. 1976. Discrimination and perceptual knowledge. *Journal of Philosophy* 73:771–91.
- . 1979. What is justified belief? In *Empirical knowledge: Readings in contemporary epistemology*, ed. Paul K. Moser, 171–92. Rowman and Littlefield Publishers.
- . 1986. *Epistemology and cognition*. Cambridge, MA: Harvard University Press.
- . 1992. Epistemic folkways and scientific epistemology. *Liaisons: Philosophy meets the cognitive sciences*, 155–75. Cambridge, MA: MIT Press.

- Hamilton, Edith, and Huntington Cairns, eds. 1963. *Plato: The collected dialogues*. Number 71 in Bollingen. Princeton, NJ: Princeton University Press.
- King, Jeffrey C. 1994. Can propositions be naturalistically acceptable? In *Philosophical naturalism. Midwest studies in philosophy*, Vol. 19, ed. Peter A. French, Theodore E. Uehling, Jr., and Howard K. Wettstein, 53–75. Notre Dame, IN: University of Notre Dame Press.
- Kirkham, Richard L. 1992. *Theories of truth*. Cambridge, MA: MIT Press.
- Kitcher, Philip. 1976. Explanation, conjunction, and unification. *Journal of Philosophy* 73:207–12.
- . 1981. Explanatory unification. *Philosophy of Science* 48:507–31.
- . 1983. *The nature of mathematical knowledge*. New York: Oxford University Press.
- . 1985. Two approaches to explanation. *Journal of Philosophy* 82:632–39.
- . 1986. Projecting the order of nature. In *Kant's philosophy of physical science*, ed. Robert E. Butts, 201–35. Dordrecht: D. Reidel Publishing Company.
- . 1988. Mathematical naturalism. In *Essays on the history and philosophy of modern mathematics. Minnesota studies in the philosophy of science*, Vol. 11, ed. William Aspray and Philip Kitcher, 293–328. Minneapolis: University of Minnesota Press.
- . 1989. Explanatory unification and the causal structure of the world. In *Studies in explanation. Minnesota studies in the philosophy of science*, Vol. 13, ed. Philip Kitcher and Wesley Salmon, 410–505. Minneapolis: University of Minnesota Press.
- . 1992. The naturalist's return. *Philosophical Review* 101:53–114.
- . 1993. *The advancement of science*. New York: Oxford University Press.
- . 1995a. Author's response. *Philosophy and Phenomenological Research* 55:653–73.
- . 1995b. Précis of *The Advancement of Science*. *Philosophy and Phenomenological Research* 55:611–17.
- Kornblith, Hilary. 1980. Beyond foundationalism and the coherence theory. *Journal of Philosophy* 72:597–612.
- . 1993. *Inductive inference and its natural ground*. Cambridge, MA: MIT Press.
- . 1994. Naturalism: Both metaphysical and epistemological. In *Philosophical naturalism. Midwest studies in philosophy*, Vol. 19, ed. Peter A. French, Theodore E. Uehling, Jr., and Howard K. Wettstein, 39–52. Notre Dame, IN: University of Notre Dame Press.
- Lehrer, Keith. 1974. *Knowledge*. Oxford: Clarendon Press.
- . 1986. The coherence theory of knowledge. *Philosophical Topics* 14:5–25.
- Maddy, Penelope J. 2005. Three forms of naturalism. In *The Oxford handbook of philosophy of mathematics and logic*, ed. Stewart Shapiro, 437–59. Oxford: Oxford University Press.
- Miller, Richard W. 1995. The advancement of realism. *Philosophy and Phenomenological Research* 55:637–45.
- Roland, Jeffrey W. 2008a. Kitcher and the obsessive unifier. *Philosophy and Phenomenological Research* 77:493–506.

- Roland, Jeffrey W. 2008b. Kitcher, mathematics, and naturalism. *Australasian Journal of Philosophy* 86:481–97.
- Rosenberg, Alexander. 1996. A field guide to recent species of naturalism. *British Journal for the Philosophy of Science* 47:1–29.
- Stroud, Barry. 1996. The charm of naturalism. *Proceedings and addresses of the American Philosophical Association*, Vol. 2 of 70. American Philosophical Association, 43–55.
- Wright, Crispin. 1992. *Truth and objectivity*. Cambridge, MA: Harvard University Press.