Philosophy and Phenomenological Research
Vol. LXXVII No. 2, September 2008

Kitcher and the Obsessive Unifier*

JEFFREY W. ROLAND
Louisiana State University

Introduction

Philip Kitcher has advanced an account of scientific progress[1] in which explanatory unification[2] plays a central role. According to that account, one of the two basic ways science progresses is explanatorily, where explanatory progress is a matter of increasing unifying power. A natural worry about this sort of view is that too enthusiastically seeking unification will lead us to impose artificial structure on the world, thus yielding an incorrect view of the world and its goings on. Kitcher has addressed this worry under the heading *obsessive unifier*.[3] In this paper, I argue that his response to the obsessive unifier worry is unsatisfactory. I further suggest a remedy to the obsessive unifier worry, but one that Kitcher is not likely to endorse.

1. The Worry

Kitcher views scientific practice at a time t as representable by an ordered n-tuple among whose components are the language L employed by the scientific community in carrying on research at t, the set of statements K about nature accepted by the scientific community at t, and the set of explanatory schemata Π accepted by the scientific

* Thanks to the following people for helpful comments: Richard Boyd, Jon Cogburn, Kevin Elliott, Harold Hodes, Richard A. Shore, and Zoltán Szabó. An earlier version of this paper was read at the 2006 Central Division Meeting of the American Philosophical Association. Thanks to the Programming Committee for that meeting as well as the session participants and my commentator, John T. Roberts.

1 See Kitcher (1986, 1993).

2 See Kitcher (1981, 1985, 1989).

3 See Kitcher (1989, pp. 494–499) and Kitcher (1995, pp. 661–662).

community at t.[4] An argument schema for Kitcher is an ordered triple consisting of a *schematic argument*, a set of *filling instructions* for that argument, and a *classification* for that argument. To help fix ideas, here is one of Kitcher's favored toy schemata—*COMMON DESCENT*:[5]

> (CD_1) G and G' are descended from a common ancestor G_0.
>
> (CD_2) G_0 members had F.
>
> (CD_3) F is heritable.
>
> (CD_4) No factors intervened to modify F along the G_0–G G_0–G' sequences.
>
> (CD_5) Therefore, members of G and G' have F.
>
> *Filling Instructions*: F is replaceable by the name of a trait, and G and G' are replaceable by the names of groups of organisms.
>
> *Classification*: (CD_1)–(CD_4) are premises; (CD_5) is deduced from (CD_1)–(CD_4).

An *explanation* (with respect to K) is an instantiation of an argument schema the conclusion of which is a member of K. A schema is explanatory (with respect to K) if some instantiation of it is an explanation (with respect to K).[6]

Kitcher understands unifying power as a matter of the greatest number of explanations, differentiated by explananda, being generated from the least number of schemata, modulo stringency of the generating schemata.[7] Intuitively, the stringency of an argument schema is a matter of how difficult that schema is to instantiate. A schema's classification, filling instructions, and the structure of the schematic sentences that make up the schema all constrain what counts as a legitimate instantiation of that schema. For example:

---

[4] At places (e.g., (1989), Kitcher's practices only involve these three components and at other times (e.g., Kitcher (1993)) there are additional components involved. For present purposes, these three will suffice. Note that 'practice' as I use it in this paper corresponds to Kitcher's 'consensus practice' (see Kitcher (1993, p. 87–89)).

[5] See, e.g., (Kitcher 1993, p. 83).

[6] Notice this notion of being an explanation *with respect to K* differs from Kitcher's notion of being an explanation *acceptable relative to K* (see, e.g., Kitcher (1989, p. 434)). The latter requires that the premises of an explanation, as well as its conclusion, be drawn from K. Kitcher recognizes that this is an idealization which can't be quite right, since K sometimes expands in the course of giving an explanation. My notion of being an explanation with respect to K is comparable to the notion of being acceptable relative to K in the extended sense Kitcher discusses in connection with why-questions. (See Kitcher (1989, p. 435–436).)

[7] For details, see Kitcher (1981, pp. 515–522) and Kitcher (1989, pp. 430–435, 477–480).

Only four-premised arguments are candidates for legitimate substitution instances of *COMMON DESCENT*; only arguments where the name of a trait appears where '*F*' appears in *COMMON DESCENT* and names of groups of organisms appear where '*G*' and '*G*′' appear in *COMMON DESCENT* are candidates for legitimate substitution instances of *COMMON DESCENT*; only arguments which have the same form as *COMMON DESCENT*, a form determined by the structure of the schematic sentences that make up *COMMON DESCENT*, are candidates for legitimate substitution instances of *COMMON DESCENT*. As far as unification is concerned, the more stringent the schemata in Π, the better.

The unifying power of a practice is identified with the unifying power of the explanatory schemata it incorporates, and science progresses as unifying power increases. The obsessive unifier worry is that in our zeal to increase unifying power we might impose unity where there is none, with the result that at the end of the day the deliverances of science (i.e., the members of *K*) will have little claim to even approximate correctness. According to Kitcher (1995), we needn't worry about the obsessive unifier because

> [t]he 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. (p. 661)

As we'll see, it's not clear that Kitcher's practitioners are in a position to recognize any such thing.

2. The Problem

Consider the following argument schema:

> (GCD_1) God willed that G, G' be descended from a common ancestor G_0.
>
> (GCD_2) God willed that G_0 members have F.
>
> (GCD_3) God willed that F be heritable.
>
> (GCD_4) God willed that no factors intervene to modify F along the G_0–G G_0–G' sequences.
>
> (GCD_5) Whatever God wills obtains.
>
> (GCD_6) Therefore, members of G and G' have F.
>
> *Filling Instructions*: F is replaceable by the name of a trait, and G and G' are replaceable by the names of groups of organisms.

Classification: (GCD_1)–(GCD_5) are premises; (GCD_6) is deduced from (GCD_1)–(GCD_5).

This schema is related to *COMMON DESCENT* in an obvious way, and was obtained by substituting innocently modified versions of (CD_1)–(CD_5) into the schema:

(G_1) God willed that ϕ_1.

(G_2) God willed that ϕ_2.

.....

(G_{n-1}) God willed that ϕ_{n-1}.

(G_n) Whatever God wills obtains.

(G_{n+1}) Therefore, ϕ_n.

Filling Instructions: ϕ_1–ϕ_n are replaceable by schematic sentences that occur as the $n-1$ premises and conclusion, respectively, of an argument schema appearing in Π^P for some P.

Classification: (G_1)–(G_n) are premises; (G_{n+1}) is deduced from (G_1)–(G_n).

Call this schema *GOD*.

GOD is a meta-schema in the sense that its allowable substitutions are constituents of other schemata, i.e., *GOD* in some sense takes another schema (such as *COMMON DESCENT*) and uses it to form a new schema having the same schematic conclusion as the initial schema. It will be useful to think of *GOD* as an operation on argument schemata so that $GOD(\pi)$ is the schema obtained by putting the i^{th} schematic premise of π for ϕ_i ($1 \leq i \leq n-1$) and the schematic conclusion of π for ϕ_n in *GOD*.[8] Schemata to which *GOD* has been applied will be said to have been *redeemed*. For any practice $P = \langle L^P, K^P, \Pi^P, \ldots\rangle$, let P_{GOD} be the result of replacing Π^P in P with $GOD(\Pi^P)$—the set of schemata obtained by redeeming each $\pi \in \Pi^P$—i.e., $P_{GOD} = \langle L^P, K^P, GOD(\Pi^P), \ldots\rangle$.[9] A practice P that has been transformed into P_{GOD} will also be said to have been redeemed. Practices that have been redeemed have some interesting properties. First, we need some terminology.

---

[8] Though it's useful to think of *GOD* as an operation, it's not necessary to think of it as such. Practitioners of P_{GOD} might arrive at the schemata in $GOD(\Pi^P)$ antecedent to independent formulation of the schemata in Π^P.

[9] Technically, the language of P will probably need to be expanded in forming P_{GOD} to allow expression of the concept of God's willing. Kitcher has to allow for this sort of modification of language unless all conceptual progress in science comes down to refinement in the reference of an existing term or explicit definition of a new term using existing terms, which is implausible.

Given practices $P = \langle L^P, K^P, \Pi^P, \ldots\rangle$ and $P' = \langle L^{P'}, K^{P'}, \Pi^{P'}, \ldots\rangle$ with $K^P = K^{P'}$, P and P' are *explanatorily equivalent* if and only if Π^P and $\Pi^{P'}$ have the same *explananda set*, where the explananda set for a class of argument schemata Π accepted for the purpose of providing explanations relative to K is the set of $\sigma \in K$ such that σ instantiates the conclusion of some schema in Π. If Π^P and $\Pi^{P'}$ both (a) have the same cardinality and (b) are equally stringent, then P and P' are *unificationally equivalent*. We'll say that P is *unificationally superior* to P' if either (a.i) Π^P and $\Pi^{P'}$ have the same cardinality and (b.i) Π^P is more stringent than $\Pi^{P'}$ or (a.ii) Π^P has cardinality less than $\Pi^{P'}$ and (b.ii) Π^P and $\Pi^{P'}$ are equally stringent. We don't have a precise measure of two sets of schemata being equally stringent, but taking a notion of relative stringency[10] for granted, the following will suffice for present purposes. Let the *letter set* of an argument schema π be the set of schematic letters appearing in π. Then two sets of schemata Π and Π' are *equally stringent* if and only if, among the $\pi \in \Pi$ and $\pi' \in \Pi'$ that have the same letter set, there are exactly as many schemata from Π which are less stringent than schemata from Π' as there are schemata from Π' which are less stringent than schemata from Π. One set of schemata is *less stringent than* the other, (say, Π is less stringent than Π') if and only if, among the schemata in Π and Π' that have common letter sets, there are more schemata from Π which are less stringent than schemata from Π' than there are schemata from Π' which are less stringent than schemata from Π.

Intuitively, this terminology allows us to compare rival practices in terms of explanatory and unifying power.[11] Practices P and P' as above are rivals because they compete for the confidence of those who hold the same beliefs; this is the force of the constraint that $K^P = K^{P'}$. If P is explanatorily equivalent to P', then P and P' explain the same things. If P is unificationally equivalent to P', then neither P nor P' unifies K^P $(=K^{P'})$ better than the other: neither does with fewer explanatory schemata than the other and neither has the upper hand with respect to stringency of schemata deployed. If P is unificationally superior to P', then P does better than P' along one of the dimensions used to measure unifying power (viz., number or stringency of schemata).

Let's return to *GOD*. For any practice P, P_{GOD} is explanatorily equivalent to P. Further, P_{GOD} is unificationally equivalent to P as well. To see this, fix a P. Obviously, there are just as many schemata in

[10] See Kitcher's discussion of such a notion in (1989, pp. 479–480).

[11] Here 'explanatory power' must be understood independently of unifying power, i.e., in a non-Kitcherian sense.

$GOD(\Pi^P)$ as there are in Π^P. After all, viewed as an operation *GOD* is a bijection between Π^P and $GOD(\Pi^P)$. Moreover, $GOD(\Pi^P)$ is neither more nor less stringent than Π^P, since for any $\pi \in \Pi^P$, instantiating π naturally induces an instantiation of $GOD(\pi)$, and vice versa. So every schema in Π^P is neither harder nor easier to instantiate than its counterpart in $GOD(\Pi^P)$.

However, it looks as if there is a sense in which $GOD(\Pi^P)$ does a better job of unifying than Π^P does: any two schemata in $GOD(\Pi^P)$ are similar in that each has been redeemed, i.e., instantiates *GOD*. In Kitcher (1981), something quite like this sort of similarity counts toward the unifying power of a set of schemata. As Kitcher puts it:

> [I]nstead of merely counting the number of different [schemata] in a [set of schemata], we [should] pay attention to similarities among them. All the [schemata] in a [set of schemata] may contain a common core schema, that is, each of them may contain some schema as a subschema. The unifying power of a [set of schemata] is obviously increased if some (or all) of the schemata it contains share a common core schema. (1981, p. 521)

To be sure, *GOD* is not a subschema of schemata in $GOD(\Pi^P)$—it's more like a superschema. But the point Kitcher is making in this passage applies equally well to *GOD*. All schemata in $GOD(\Pi^P)$ are robustly similar in a certain way, and for that reason $GOD(\Pi^P)$ arguably unifies better than Π^P.

So P_{GOD} is either a formidable rival to any practice P that might be favored by Kitcher's account or a rival that bests any practice P that might be favored by Kitcher's account. Either way, Kitcher needs a principled way to rule out P_{GOD} as a legitimate rival practice if his account is to be defensible.[12] Call this the *God problem*[13]. The God problem apparently undermines Kitcher's response to the obsessive unifier worry. Do Kitcher's practitioners have grounds on which to disqualify P_{GOD}?

3. Responses and replies

A defender of Kitcher might attempt to model a response to the God problem on Kitcher's reply to a challenge raised by Larry Laudan.[14] According to Laudan, we can infer nothing about the truth or referential achievements of a scientific theory from its apparent explanatory or

[12] I here assume what I take to be uncontroversial, viz., that P_{GOD} is not acceptable to Kitcher (or many other parties to the debate under consideration).

[13] This problem develops a suggestion made independently by Richard Boyd (in conversation) and Richard Miller (in Miller (1995)).

[14] See, e.g., Laudan (1981).

predictive success, since an induction on the history of science reveals that current theory is likely to be wrong in various ways despite that apparent success. Kitcher responds to Laudan by arguing that the things which turned out to be wrong about past theories were dispensable parts of those theories, in the sense that those parts played no role in the apparently successful explanations and predictions of the theories. We are committed to the approximate correctness of only those parts of our theories that actually do some work for us; the idle parts of a theory can fail without taking the theory down with them.[15] This undermines Laudan's so-called *pessimistic induction* on the history of science, but does it help with P_{GOD}?

Kitcher's defender argues as follows: By Kitcher's reply to Laudan, we can safely and justifiably ignore those parts of a practice that do no explanatory or predictive work. Indeed, we can safely excise those idle parts; this is just what Kitcher needs to deal with P_{GOD}. The very same explanations and predictions are obtainable from P as are obtainable from P_{GOD}; so God doesn't make any explanatory or predictive difference, i.e., God is explanatorily and predictively idle in P_{GOD}. Thus we can safely excise God from P_{GOD}, and the result of doing this is just P, as desired.

This argument depends heavily on judgments of explanatory and predictive idleness. If we can't reliably make these sorts of judgments, then we won't be in any position to say what part of a theory can or cannot safely be ignored or excised. In replying to Laudan, Kitcher appeals to explanatory schemata in determining which parts of a practice are idle. He draws a distinction between *working posits* and *presuppositional posits*. The former are those entities referred to (if anything is) by terms appearing in successful explanatory schemata. The latter are those entities which must exist in order for instantiations of successful explanatory schemata (taken at face value) to be true, even though the explanatory schemata themselves make no reference to them. Kitcher argues that Laudan's pessimistic induction works, if at all, to show that presuppositional posits are not secure. Hence, Laudan's induction cuts against only those posits that aren't referred to in explanatory schemata. But Kitcher has already argued that only idle parts of a practice are in danger from Laudan's induction. Thus, the idle parts of a practice are just those not referred to in its explanatory schemata.

Though at first blush this seems promising, it doesn't help with the God problem. Solving the God Problem requires a principled reason to

[15] For a more complete presentation of Kitcher's response to Laudan, see (Kitcher 1993, pp. 141–149).

prefer P over P_{GOD}. These practices are explanatorily equivalent, so we can't choose P over P_{GOD} on the basis of explanatory superiority. And since P is not unificationally superior to P_{GOD}, we can't choose P over P_{GOD} on the basis of unificational superiority. The present suggestion is that we should look to successful explanatory schemata and judge, on the basis of what appears in those schemata, that God is a non-player in successful explanations and predictions. But which schemata should we examine in order to reach this conclusion? Not the schemata of P_{GOD}; we would get the wrong result, since God is a player in those schemata. If we look to the schemata of P to determine which elements are contributing in the right ways, then we've already privileged P over P_{GOD}, and for no (obvious) good reason. But we can't legitimately privilege P over P_{GOD} in the process of choosing between P and P_{GOD}. This shows that the sort of defense under consideration only works after some practice has become the practice to beat, the "frontrunner" practice among some group of rival practices. But this is just the question we're addressing: On what grounds can P_{GOD} legitimately be rejected as the practice to beat?

One might think that appealing to a principle of parsimony, such as Ockham's razor, would help here. God is superfluous, and that makes it legitimate to excise God from P_{GOD}. A bit of probing shows that this isn't as helpful as it might initially appear. How are we to understand the claim that God is superfluous? God is not unificationally superfluous, for reasons given above. In addition, judgments concerning unification are made on the basis of (intuitively) the ratio of argument schemata to explananda (modulo stringency). Since GOD is a single schema and P and P_{GOD} have the same explananda set, this ratio in the case of P_{GOD} will never be larger than that in the case of P, for any practice P. Similarly, God is not explanatorily superfluous; playing an explanatory role just is a matter of figuring in the most unifying argument schemata. Kitcher isn't in a position to argue that God is ontologically superfluous, since on his view judgments about what mechanisms and kinds there are depend on the argument schemata accepted as part of the prevailing practice. This is characteristic of his "top down" approach to explanation: "Top down approaches [to explanation] will attempt to provide an account of what theoretical explanation is, use this as a basis for underwriting talk about 'fundamental mechanisms,' and so proceed toward the identification of causes in particular cases" (Kitcher 1989, p. 430).[16] The difficulty with responding to the God problem by an appeal to parsimony is that in spelling out in what way P_{GOD} is less parsimonious than P we must

[16] See also Kitcher (1985, 1986).

make a judgment concerning God that is analogous to the judgment that God is idle in P_{GOD}, and making such judgments, for Kitcher, ultimately depends on having already settled on a stock of argument schemata. Indeed, we can see Kitcher's unificationism as an attempt to articulate a principle of parsimony: accept only those mechanisms and kinds that figure in the most unifying argument schemata. No wonder, then, that appeal to a principle of parsimony to defend Kitcher's unificationism takes us in a circle.

One might also think that the God problem can be rejected on grounds that schemata from $GOD(\Pi^P)$ aren't genuinely unifying on Kitcher's account because they fail to share a common core schema. It seems to me that one can say at least the following three things in reply to this.

The first is that Kitcher's account is not obviously committed to unifying schemata sharing a common core. As we noted in §2, Kitcher (1981) is explicit about the sharing of a common core schema contributing to unification. However, Kitcher doesn't mention core schemata in discussing his account of scientific explanation in either (1989) or (1993). He does introduce the notion of an *explanation extension* in these later works—roughly, an explanatory schema π is extended by another schema π' if π appears as a subschema of π'—which might do the same work as shared core schemata. But explanation extensions are inter-practice phenomena: extending some or all of the schemata in Π^P yields a new (unificationally superior) set of schemata, and substituting this new set for Π^P in P yields a new practice P'. If P appears early enough in the development of science, then the schemata in Π^P may well share no subschema. If there is no such shared schema, then the God problem will arise here even if it doesn't arise later in the sequence of practices. This gives us a sort of start-up problem.

Second, if we key on syntactic features of argument schemata in making unification judgments, as Kitcher does, the preference for common *core* schemata over common *shell* schemata seems unmotivated, if not outright wrong. A shell schema is a superschema or metaschema like *GOD*. Imagine we're given nine schemata, three of which have some core schema π in common, three of which have another core schema π' in common, and three of which have yet another core schema π'' in common. Suppose we're asked to partition the collection of schemata according to shared syntactic features, with the aim of minimizing the number of pieces in our partition. We would naturally partition the collection into three pieces, grouping them by shared core schema. Now imagine playing the same game with the redeemed counterparts of the nine schemata. In this case, we would end up with the trivial partition, every schema in the same piece, since syntactically

every schema in the collection of redeemed schemata has something in common. So if we're interested in unification, and unification is determined by or read off syntactic similarity, then it really looks like we should prefer shared shell schemata over shared core schemata. At best, the choice of shared core schemata over shared shell schemata is unmotivated. We need some principled reason to think that schemata that share a core unify better than schemata that share a shell. Such a reason can't be grounded in syntactic features of schemata, but it's not clear what else Kitcher has to work with in light of his "top-down" approach to scientific explanation.

Third, even if we accept that schemata must have common core schemata and that shared core schemata are preferable to shared shell schemata, P_{GOD} is still unificationally *equivalent* to P. After all, every core schema common to the members of Π^P will have a counterpart core schema common to the members of $GOD(\Pi^P)$; whenever π is a subschema of π', $GOD(\pi)$ is a subschema of $GOD(\pi')$.[17] Given this, partitioning Π^P and $GOD(\Pi^P)$ according to shared core schemata will result in partitions with an equal number of pieces each with an equal number of member schemata. Thus, even conceding the points of the two immediately preceding paragraphs, Kitcher still needs to face the God problem.

The solutions to the God problem we've so far considered have been unsuccessful. But a defender of Kitcher might hold that the God problem is a non-starter and so is not in need of a solution. The idea is that scientific theorizing aims at order or structure so as to facilitate our understanding of and getting along in the world, and since God's intentions are arbitrary introducing God brings chaos rather than order. So introducing God is incompatible with the aims of science and is *a priori* ruled out. There is no God problem after all. Ignoring that this move is probably not available to Kitcher, since he rejects the *a priori* (and most emphatically so in epistemology),[18] it is still unsatisfactory.

In the first place, this argument involves the same sort of circularity we confronted in some of the solutions already considered. On Kitcher's view, judgments concerning order depend on argument

[17] There is a wrinkle here in that the conclusion of an instantiation of $GOD(\pi)$ will be of the form "God wills that ϕ" rather than simply ϕ so that technically $GOD(\pi)$ is not exactly a subschema of $GOD(\pi')$. This can be ironed out by noting that there is still a perfectly good sense in which $GOD(\pi)$ is a subschema of $GOD(\pi')$. Also, P_{GOD}-practitioners would likely accept that God wills that ϕ if and only if ϕ, making a substitution of the latter for the former in $GOD(\pi')$ legitimate from their point of view.

[18] See, e.g., Kitcher (1992, p. 76).

schemata. The unification brought about by lowering the ratio of argument schemata to explananda reflects or imposes order. So if God figures in the most unifying argument schemata, then by Kitcher's own lights God contributes to order, not chaos. Moreover, anyone who seriously advocated P_{GOD} could respond to the claim that the arbitrariness of God's intentions leads to chaos in at least the two following ways. First, since our experience to date has been the outcome of God's will, and that experience has not been particularly chaotic, by induction we have no reason to expect that our future experience, though an outcome of God's will, will be particularly chaotic. Just because God might make things to go differently in the future is not a reason to expect that It will make things go differently in the future. Second, the theist might "go Cartesian" and invoke God's perfection, specifically the perfection of God's will and Its benevolence toward us. Given the perfection of God's will It will not act in a way inconsistent with Its aims, and given Its benevolence toward us one of Its aims is that our experience be uniform and stable enough to allow us to get along in the world. From this conception of God, and it is far from ad hoc, the theist might reasonably argue that God is a point of stability, rather than instability, in our experience.[19]

4. A proposal

The solution to the God problem just considered fails because: (i) it distinguishes between active and idle factors of successful explanatory practice on the basis of the explanatory schemata deployed by that practice; (ii) explanatory schemata are chosen solely according to unifying power; (iii) unifying power alone is insufficient to choose between $GOD(\Pi^P)$ and Π^P (hence, also between P_{GOD} and P); therefore, (iv) judgments of activeness and idleness are tied to unification in such a way that they cannot facilitate choosing between P_{GOD} and P. Solving the God problem requires a way of identifying active factors of successful scientific practice which does not rely so heavily on unification considerations. In closing, I'd like to suggest that adopting a robust causal realism would give us what we need.

If the world is systematically unified and a unifying theory or practice is one that meets a certain sufficiently high standard of unification, then it's likely that the pronouncements of unifying theories would be (approximately) correct. So it looks like we can answer the obsessive unifier worry by adopting the view that nature is unified. This isn't a

[19] I'm not, of course, endorsing this view. I'm just pointing out that it is open to one who's interested in defending the adoption of P_{GOD} on Kitcherian grounds.

novel idea; Kitcher considers something like this in (1986), and he there recognizes that his account of scientific progress is consistent with just such a view, which in (1993) he calls *strong realism*.[20] Kitcher is uncomfortable with this view because he takes it to be 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, p. 170)

Analogous to the infamous Benacerraf problem for mathematical knowledge, a strong realist picture allows us to make sense of our scientific discourse, but it is cold epistemological comfort, according to Kitcher, as it leaves us with no convincing story about how we access the strongly realist world of our theorizing.[21] Call this the *access problem*.

The access problem is what motivates Kitcher to give explanatory schemata such a central role in his account of scientific progress. Explanatory schemata are presented linguistically. Thus, for Kitcher, accessing the relevant features of the world is little more problematic than accessing language: ''Our recognition of an explanatory ordering precedes, and makes possible, our identification of causal relationships'' (Kitcher 1985, p. 637), and the explanatory ordering in question is reflected in the explanatory schemata we accept. But as we have seen, Kitcher's counsel on how to determine which explanatory schemata to accept opens him to the God problem. In other words, Kitcher's solution to the access problem is what makes him vulnerable to the God problem. Causal realism provides a solution to the access problem which does not give rise to the same vulnerability.

Adopting a robust causal realism affords us referential access to natural kinds, which in turn allows us to access the causal order of the world. The causal order of the world forms the *ontological ground*[22] of inferential (inductive) practice in science; induction on natural kinds results in (approximate) truths because kinds align with the

---

[20] See, e.g., pp. 169–173.

[21] Kitcher notes this affinity with the Benacerraf problem. See Kitcher (1993, p. 170, fn. 57).

[22] Cf. Millikan (1999).

ontologically robust causal order.[23,24] Thus, adopting causal realism amounts to adopting a solution to the access problem. Moreover, since a causally-ordered world is effectively a unified world, causal realism supports the claim that the world is systematically unified.

The claim that the world is unified can be understood as roughly the claim that situations and goings on in the world are connected in ways such that many events and states of affairs that appear to be of distinct types are, in fact, of the same type, so that there are many fewer types of phenomena in the world than there *prima facie* seem to be. On the view that the 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 be "read off" of these explanatory dependencies, but, unlike on Kitcher's view, those dependencies do not rely on unification considerations. Of course, things are a good bit more complicated than this, but I trust the general idea is clear: an ordered world and a unified world are two sides of the same coin. This being the case, causal realists arguably have no problem with P_{GOD}. They can make sense of correct explanations—successful explanatory practice—independently of unification considerations, which allows them to identify active factors in successful explanations without appealing to unification. This in turn allows them coherently to classify God as explanatorily idle. So it seems that where Kitcher is caught between the God problem and the access problem, the causal realist is free of both.[25]

I have not here canvassed all the ways in which Kitcher's view might be defended against the challenge presented by the God problem. Neither have I offered a conclusive argument in favor of causal realism. My purpose has been the more restricted one of pointing out that Kitcher's position is subject to a particular challenge, that some obvious responses to that challenge don't work, and that the challenge can be met by supplementing Kitcher's position in a way that Kitcher himself

---

[23] For more on this, see Boyd (1980, 1989, 1991), Kornblith (1993), and Roland (2007).

[24] This should not be taken to suggest that merely adopting a robust causal realism yields a map of the causal order of the world, or that mapping that order thereby becomes an easy matter, or even that there is a unique map of the causal order. Adopting a robust causal realism simply provides necessary resources for undertaking the arduous task of mapping the causal order.

[25] Notice that the solution to the God problem I've suggested need not be an abandonment of unificationism. I have only advocated adding a premise to Kitcher's view endorsing causal (strong) realism, which as I mentioned above Kitcher holds is consistent with his view. This solution does lead to an abandonment of Kitcher's "top down" approach to explanation, but if this approach were necessary to unificationism then Kitcher's view would not be consistent with strong realism.

recognizes as a possibility but rejects. If I have been successful in this much, then I have shown that we should not be too quick to dismiss causal realism in favor of Kitcher's top down approach to scientific explanation.

References

Boyd, Richard (1980). "Scientific Realism and Naturalistic Epistemology," in *Proceedings of the Philosophy of Science Association*, volume 2, Philosophy of Science Association, 613–662.

—— (1989). "How to be a Moral Realist," in *Contemporary Materialism*, edited by Paul K. Moser and J. D. Trout, Routledge, 297–343.

—— (1991). "Realism, Anti-Foundationalism, and the Enthusiasm for Natural Kinds," *Philosophical Studies*, 61: 127–148.

Kitcher, Philip (1981). "Explanatory Unification," *Philosophy of Science*, 48: 507–531.

—— (1985). "Two Approaches to Explanation," *Journal of Philosophy*, 82: 632–639.

—— (1986). "Projecting the Order of Nature," in *Kant's Philosophy of Physical Science*, edited by R. E. Butts, D. Reidel Publishing Company, 201–235.

—— (1989). "Explanatory Unification and the Causal Structure of the World," in *Studies in Explanation*, edited by Philip Kitcher and Wesley Salmon, volume XIII of *Minnesota Studies in the Philosophy of Science*, University of Minnesota Press, 410–505.

—— (1992). "The Naturalist's Return," *The Philosophical Review*, 101 (1): 53–114.

—— (1993). *The Advancement of Science*, Oxford University Press.

—— (1995). "Author's Response," *Philosophy and Phenomenological Research*, 55: 653–673.

Kornblith, Hilary (1993). *Inductive Inference and Its Natural Ground*, MIT Press.

Laudan, Larry (1981). "A Confutation of Convergent Realism," *Philosophy of Science*, 48: 19–49.

Miller, Richard W. (1995). "The Advancement of Realism," *Philosophy and Phenomenological Research*, 55: 637–645.

Millikan, Ruth (1999). "Historical Kinds and the Special Sciences," *Philosophical Studies*, 95: 45–65.

Roland, Jeffrey W. (2007). "Maddy and Mathematics: Naturalism or Not," *The British Journal for the Philosophy of Science*, 58: 423–450.