Preprint
Forthcoming in *Philosophy and Phenomenological Research*

Naturalism, Reduction and Normativity: Pressing from Below

JOHN F. POST
Vanderbilt University

David Papineau's model of scientific reduction, contrary to his intent, appears to enable a naturalist realist account of the primitive normativity involved in a biological adaptation's being "for" this or that (say the eye's being for seeing). By disabling the crucial anti-naturalist arguments against any such reduction, his model would support a cognitivist semantics for normative claims like "The heart is for pumping blood, and defective if it doesn't." No *moral* claim would follow, certainly. Nonetheless, by thus "pressing from below" we may learn something about moral normativity. For instance, suppose non-cognitivists like Mackie are right that the semantics of normative claims should be "unified": if the semantics of moral claims is non-cognitivist, so too is that of all normative claims. Then, assuming that a naturalist reduction does yield a sound cognitivist account of the primitive normativity, it would follow that our semantics of moral claims is cognitivist as well.

I. Introduction

In "The Status of Teleosemantics, Or How to Stop Worrying About Swampman" David Papineau (2001) defends teleosemantics against objections advanced by David Braddon-Mitchell and Frank Jackson (1997). Papineau's defense succeeds, I believe, and yet it poses a problem for what he says about normativity:

> Wherever the normativity of content comes from, it can't be from biology, since biology deals in facts, not prescriptions. . . . It has always mystified me why anybody should think that biology helps with normativity (Papineau 2001, 280).

The problem is this. Papineau argues that Braddon-Mitchell and Jackson "fail to take their reductionist moral to heart" when it comes to selectional content (Papineau 2001, 279). But it seems Papineau fails to take his own reductionist moral to heart when it comes to normativity. When his reductionist moral is applied to the normativity he too sees in the selectional property of a biological adaptation's being for this or that — say the heart's being for pumping blood — it looks as though the normativity can come from biology after all, by way of being reducible to biological facts.

Normativity reducible to facts? Some philosophers might see this as a reduction to absurdity of Papineau's reductionism. Others, however, might see it as a welcome consequence, however unintended. His reductionist method, if taken to heart, would appear to disarm the crucial arguments against naturalist realism about the normativity in question. This would greatly improve the prospects of a positive account, in naturalist-realist terms, of an important

kind of normativity some think is objectively in the world — the primitive normativity involved in a biological adaptation's being for, or designed to do this or that. Such normativity is the target of this paper, though the discussion will have implications for moral normativity as well. Not that moral normativity can somehow be inferred from the normativity involved in selectional properties; far from it. Nonetheless, by "pressing from below" on this primitive normativity, we may learn something about the higher-level moral normativity, as we shall see.

The crucial arguments against naturalist normative realism include Hume's-Law arguments (no norm can be inferred from any facts); open-question arguments (no norm can be reduced to them either); Mackie-style queerness arguments (Mackie 1977); arguments that no naturalistic reduction can preserve the critical gap between what something actually does or is disposed to do and what it ought to do; and arguments that so-called objective normative properties would be unacceptably epiphenomenal, playing no significant causal role, and no predictive role either (e.g., Harman 1986 and 1998, and Papineau 2001, 287). Opponents of naturalist normative realism rely on these arguments, or their variants, in rejecting any such realism. Versions of naturalist normative realism which disarm the crucial arguments and their variants would thus be well positioned to provide a sound positive account of the targeted normativity in realist terms. Among such versions, I argue, is one enabled by what Papineau endorses as "the standard model of scientific reduction." This model, and models relevantly like it, would disarm the crucial arguments against naturalist normative realism, often by rendering them irrelevant, or so I hope to show in the case of the targeted normativity.

First, some preliminaries, starting with a question about "pressing from below." Why should meta-ethicists bother with the lower-level, primitive kind of normativity involved in a biological adaptation's being for something? One reason to bother is that most if not all non-cognitivists about normativity — notably Mackie (1977), 15, 53-58 — assume that the semantics of normative claims should be "unified," in the sense that, as David Copp puts it, "If our semantics of moral claims is non-cognitivist, then . . . the semantics of all normative claims must be non-cognitivist."[1] It follows that if a unified semantics is right, and if a Papineau-style model of reduction yields a successful cognitivist account of the low-level selectional normativity, then our semantics of the high-level moral claims must be cognitivist too. (For clarity, it should be emphasized that I nowhere depend on the unified-semantics thesis. Indeed I suspect it is false. Esthetic normativity, for example, may well have a non-cognitivist semantics, so that esthetic claims are neither true nor false, whereas claims to the effect that a biological adaptation is normatively for this or that may be true or false. My target is non-cognitivists and others who believe that if the semantics of moral claims is non-cognitivist, the semantics of all normative claims must be non-cognitivist. Note also that the thesis would not automatically be threatened by different kinds of normativity having different kinds of semantics. So long as the different kinds of semantics all imply that the various normative claims are, say, non-cognitivist, the thesis could still be true.)

[1] Copp (1995), 18, though not himself a non-cognitivist. What lies behind this unification thesis, I suspect, is the belief that key anti-cognitivist arguments are as effective in the case of non-moral normativity as in the moral.

Another reason to bother with this primitive normativity is that it is simpler and more tractable than moral normativity. Typically the debate between realists and non-realists about normativity is conducted in terms of the latter. This amounts to starting with what is presumably the most complex and sophisticated kind of normativity there is — the kind meant to apply to the behaviors and moral principles of creatures endowed with reason, self-aware rule-following and moral agency, awareness of the moral agency of others, and so on. Instead of beginning at the bottom, so to speak, with the simplest and most tractable cases, philosophers typically begin at the top — a significant difference between Richard Boyd's naturalist normative realism and mine; while he presses from above, I press from below. By beginning at the top, as also does Mackie, philosophers risk projecting the lessons drawn from struggles with the semantic status of moral normativity onto the status of non-moral normativity, thereby overlooking, among other things, the possibility that what we learn about the latter might tell us something about the former, perhaps even that the semantic status of moral norms must be cognitivist after all (if with Mackie one assumes a unified semantics, though not only if). Of course there are comparable risks in pressing from below, but most philosophers are far more alive to such risks than to the risks in pressing from above.

A second preliminary is that, due to limits on length, I make no attempt here to defend a theory of meaning and reference adequate to support certain presuppositions about meaning and reference. Suffice it say that a theory like Ruth Millikan's, perhaps among others, would be adequate for present purposes.[2] Her account accords especially well with rejecting, as I do, (i) a priori or other epistemically privileged access either to the meaning of a term, or to its reference, or to the relevant "central" properties of the affair to which the term refers (if it refers); and (ii) the widespread assumption that the mind's contents alone, or a society of minds' contents, determine the criteria for a term's success or failure in referring, so that "What's inside determines how things must be outside for the reference to be successful."[3]

To be more specific, one way in which Millikan's work bears on the intended application of Papineau's model of reduction is this. Some philosophers might object that the relevant folk terms have an *indexical* semantics, say a semantics according to which the folk term 'water' tells us that water is whatever it is around here that plays the folk role. The objection conflicts with Millikan's strongly realist account, according to which the semantics of the relevant folk terms is not indexical, nor is the reference of the folk term 'water' fixed by, say, what Braddon-Mitchell and Jackson 2002, 372ff, call the term's "association with the folk roles associated with . . . water." Instead, what fixes the reference is the term's biological function or purpose.

Furthermore, the objection conflicts with her own account of indexicals, according to which "to interpret an indexical, one must have prior knowledge of, one must know

---

[2] Cf. Millikan (1984), (1993), (2001).

[3] Millikan (2001), 131. As regards the objection that it would be question-begging to use her account here because it assumes the targeted normativity, see footnote 15, below.

independently and ahead of time, what item bears the indexical's adaptation relation to the indexical token" (Millikan 1993, 270-271, emphasis supplied; cf Millikan 2001, 131ff). Hence an indexical could not be what tells us (if anything does) that water is whatever it is around here that plays the folk role, since we must already know independently what extant item bears the relevant relation to the indexical token.

Or suppose someone challenges the intended application of Papineau's model by suggesting that the application could not explain why water is not a disjunctive kind, on the grounds that the folk semantics specifies that the sample whose empirical nature is relevant is that which locally actually plays the role. This too conflicts with both Millikan's semantics for such terms and her account of indexicals, which raise serious questions about the very relevance of twin water. According to her account, XYZ would have to be physically possible, whereas it appears not to be, hence is beside the point of the intended kind of reduction.

A third preliminary concerns whether Papineau's model of reduction really is, as he says, "the standard model of scientific reduction." I think it is, at least near enough for present purposes, though limits on length preclude arguing the point here. In order not to beg the question I will speak instead of Papineau's model of scientific reduction, or Papineau reduction for short.

A fourth is this. When contemporary naturalists consider whether a certain phenomenon is real and, if so, how best to account for it, usually they begin by looking to the relevant science — not for the last word, but for some effective first words. Yet the same naturalists mostly balk at doing so when it comes to normative matters, and deny the applicability of scientific reductive method to them. This includes Papineau himself, to judge from the displayed quotation we started with. Nonetheless, the strategy of this paper is to apply his model of scientific reduction to the targeted primitive normativity, then see what follows, in order to argue that applying the model disarms the crucial arguments against naturalist normative realism at least in this case.

An immediate complication is that applying the model to normativity may strike some as question-begging. Those who oppose to the very idea of applying any such model or method to normativity may object that doing so begs the question whether it does apply.[4] But the objection neglects a key role played by trial-balloon assumptions in our reasoning about what there is and its nature. Suppose I believe that Mill's methods of inductive reasoning apply to all physical phenomena whatever, and I unqualifiedly reject, as question-begging, anyone's assuming instead that some conflicting method or model applies to this or that phenomenon. The trouble with such an unqualified stance is that it excludes out of hand any trial model or method that, contrary to Mill's methods, (i) does not construe every phenomenon as a system whose components are causally separable and structurally invariant over operating conditions, but (ii) construes some phenomena as, say, dissipative systems in dynamic equilibrium, or functional complexes with

[4] As an anonymous referee has objected, at length.

feedback interdependencies, or single field structures obeying an integro-differential equation.[5] To exclude such phenomena out of hand conflicts with the possibility that there are more things in heaven and earth than are dreamt of in Mill's methods, as we now know there are.

The unqualified stance appears to go wrong in neglecting the distinction between (i) flatly assuming that a certain model or method applies to the phenomenon in question, and (ii) assuming for the sake of argument that it does, in order to see what follows. As Barbara Herman remarks, "one way to argue for a metaphysical view is to see how much that matters follows from it" (Herman 2003, 2). Furthermore, if what follows is born out by relevant testing or critique — including but not limited to empirical observation and experiment — the model's applicability may gain significant support. If it is not born out, the model may be in trouble (as in fact Mill's methods are when it comes to a dissipative system in dynamic equilibrium, or the like). This means that the trial model is at significant epistemic risk when assumed for the sake of argument. It is not just assumed, take it or leave it, but held liable to refutation, or to substantial revision, or at least to anomalies that demand explanation. To assume for the sake of argument that the trial model or method applies is not question-begging.[6]

The fifth and final preliminary concerns the following objection.[7] Even if the crucial arguments against naturalist normative realism were all disarmed, it would not follow that such realism should be accepted. In addition to surviving arguments against it, the theory should enjoy some positive support as well; even if there remain no effective arguments for *rejecting* it, there should be some good reason for *accepting* it. The objection presupposes that disarming the arguments against a theory never yields positive support, whereas I think it often does. But rather than argue the point here, it suffices to emphasize that my present aim is not to show that we should accept the theory, but to argue for a conditional**:** if we adhere to a method relevantly like Papineau's, and to a theory of meaning and reference relevantly like Millikan's, then the crucial objections to any such naturalist realism can be disarmed. Should positive support accrue as a result, so much the better. Meanwhile it remains true that, as Herman says, "one way to argue for a metaphysical view is to see how much that matters follows from it."

So much for preliminaries. Now, what sort of account of the targeted normativity might be enabled by a model of scientific reduction relevantly like Papineau's? According to one such trial account or hypothesis, and put baldly, a biological adaptation *A*'s normative property of

---

[5] Hooker (1987), 288, uses the example of Mill's methods to illustrate the theory-dependence of method (how a method's applicability depends on which theory of the targeted phenomenon is true). Cf. Robert Brandon (1990), 144-149, on the empirical presuppositions of applicability of the Principle of Natural Selection, and Alan Chalmers (2003) on the theory-dependence of the use of instruments in science.

[6] See further Section III and Subsection (1) of Section IV, below, in connection with Hume's-Law arguments; and Brown (1993) on how "a theory-laden observation *can* test a theory."

[7] Urged at length by an anonymous reader.

being for *E* is the selectional matter of *E*'s being the effect in virtue of which, ancestrally, having *A* had an adaptive advantage over its alternatives. Equivalently, *A*'s normative property of being for *E* is a matter of *E*'s being the effect of *A*'s past instances in virtue of which *A* was selected for.[8]

Already one can hear the usual chorus of protest against any such account, from Hume's-Law objections to charges of epiphenomenalism. *E*'s being the effect of *A*'s past instances in virtue of which *A* was selected for, it will be said, is a purely factual affair, as one would expect of anything coming from biology. And no prescription can come from any facts. In particular, as Papineau says in the displayed quotation we started with, "Wherever the normativity . . . comes from, it can't be from biology, since biology deals in facts, not prescriptions." End of story.

Well, not quite the end, at least not if we take Papineau's own reductionist moral to heart. Sections II and III review the relevant features of Papineau's model of scientific reduction. Section IV explains how these features disarm the crucial objections to reducibility in the case of the primitive normativity involved in an adaptation's being for this or that, thereby greatly improving the prospects of a robust realist account of such normativity in naturalist terms. The final Section V hazards some further conclusions. One is that teleosemantics could explain the normativity of content after all, contrary both to Papineau and to Braddon-Mitchell and Jackson (not that no other account could do so). Another is that the model would disarm crucial arguments against naturalist normative realism in the case of *other* kinds of normativity — including *moral* normativity — to the extent that such arguments have the same form and presuppositions as the arguments against naturalist normative realism in the present case (some may not). Still another conclusion is that if the naturalist account offered here should prove sound, so that the semantics for this primitive kind of normativity is cognitivist, then those who hold to a "unified semantics" for normative claims, as does Mackie, would have to be cognitivists about moral claims as well.

II. Reduction

Papineau uses the example of the reduction of water to H_2O to illustrate his model of scientific reduction. For present purposes, however, his example needs to be replaced with one in which, as in the case of normativity, there has been substantial doubt as to the real existence of the supposed phenomenon to be reduced. Otherwise we risk being misled by cases like that of water, about whose real existence there has been no such doubt.

To this end, consider "ball lightning," reported by ordinary folk since antiquity as fantastic, glowing, floating balls of colored light, often accompanied by a hissing sound and

[8] Cf. Sober (1993), 84: "characteristic c is an adaptation for doing task t in a population if and only if members of the population now have c because, ancestrally, there was selection for having c and c conferred a fitness advantage because it performed task t." I am indebted to Derek Turner for reminding me of this passage. See also Sterelny and Griffiths (1999).

distinct odor. As in the case of normativity, its real existence has been doubted, indeed dismissed as old wives tales — by physicists, for example, whose theories of electromagnetism ruled it impossible. Nonetheless, let's suspend judgment and adopt the trial hypothesis that ball lightning is a phenomenon objectively in the world best approached not by way of conceptual analysis or the like, but by way of a Papineau model of scientific reduction.

This is fundamentally how plasma physicists have actually proceeded. Their implicit strategy has been to "equate" ball lightning, provisionally, with a high-density plasma — "equate" them in the sense of supposing they are either (i) identical in standard conditions (say of temperature and pressure), or perhaps (ii) only equivalent in such conditions, an equivalence expressible by a relevantly modalized biconditional that serves as a bridge principle (as does the identity claim). Either way, equating them, if successful, would support the claim that the folk role of 'ball lightning' is realized by a high-density plasma.[9]

In order to apply a Papineau model of reduction to ball lightning, we need only replace 'water' throughout his account with 'ball lightning'. Thus what he would call the folk role of 'ball lightning' is "the set of descriptions which pre-theoretical intuition uses to pick out instances" of ball lightning. According to the plasma theory of ball lightning, the folk role is fulfilled, or realized, by a high-density plasma. The point of the theory is not to eliminate either ball lightning or the folk concept of ball lightning or our pre-theoretical intuitions about it. Rather, the point is to find the "theoretically interesting states" that in fact fill or realize the folk role. What is "theoretically interesting" about them is that classifying ball lightning as a high-density-plasma state introduces "properties that are causally efficacious." Specifically, by so classifying ball lightning we relate it to "the basic laws governing physical causation." In this way, we "become better able to understand the behavior" of ball lightning's manifest everyday properties within the framework of "a powerful, unifying, explanatory theory" which "tells us about the underlying nature" of ball lightning (Papineau 2001, 282, 285, 287).

The aim of the reduction is not analysis — not to capture or to conform to received concepts or usage or meaning, or to conform to our pre-theoretical intuitions about ball lightning. Instead, the aim is to find how folk roles are filled in the *actual* physically possible world (PPW). Among other things, this means that even though merely logically or conceptually possible scenarios — say certain versions of Twin Earth, Swampman, and so on — "are all right for teasing out the structure of everyday thinking," they "have no bearing on how those folk roles are filled in the actual world." So too for our pre-theoretical intuitions about wildly counterfactual scenarios. It follows that such scenarios — Papineau calls them "merely

[9] There are other physical accounts of ball lightning, including accounts in terms of a chemiluminescent process, an air vortex containing luminous gases, microwave radiation within a plasma shell, or an atmospheric maser. The jury is still out on what is the best account. Cf. Stenhoff (1999)

possible" cases — are irrelevant to the reduction.[10]

As Papineau says of 'water', it might seem that 'ball lightning', being a natural kind term, has the sort of semantics Saul Kripke proposes for such terms. Hence it might seem that 'ball lightning' is a rigid designator, "referring in all contexts, including modal contexts, to the actual stuff which plays the [ball-lightning] role in this world" (where the modal contexts include those that are, or are in, logically or conceptually possible worlds that are not physically possible). Papineau's response is basically that for purposes of scientific theoretical reduction, what counts is the role-realization claim — here the claim that in the *actual* PPW the "ball-lightning" role is realized by a high-density plasma. So far as the plasma theory of ball lightning is concerned, it does not matter whether the semantics of the folk term 'ball lightning' is one of (i) rigid designation, (ii) flaccid designation (the term "refers in any context to the stuff which would play the [ball-lightning] role under the suppositions constituting that context"), or (iii) role designation ('ball lightning' refers to the role, not the realizers). The choice among these three options "seems to collapse into a matter of local sociolinguistics. . . . Nothing important to [plasma physics] hangs on this choice." The significant question is "the relation between everyday *roles* . . . and *realisers*, the theoretically interesting states which fill these roles in the actual world." Any objection is beside the point which presupposes that the folk term being reduced — or the reduction or the role-realization claim — requires a semantics of, say, rigid-designation (Papineau 2001, 282, 284-286; cf. my remarks under the second preliminary in Section I, on Millikan's semantics for the relevant terms).

Clearly, the plasma theory does not claim that ball lightning and a high-density plasma are logically or conceptually identical. Rather, in Papineau's terms, the theory claims that the "ball lightning" role is realized in the actual PPW by a high-density plasma. I would add that given only the evidence for the theory, and given what the theory may therefore legitimately claim, what happens in merely logically or conceptually possible worlds is a "don't-care" — for the good reason that it is the testing of the theory which tells us the range of conditions in which it can be shown to hold, thereby circumscribing what count as the relevant standard conditions in the actual world. The theory may claim identity, but only a relatively weak contingent identity — identity in the relevant proper subset of the PPW's, namely those in which the relevant actual-world standard conditions obtain.[11] Attempting to counter-example a plasma theory of ball

[10] Papineau (2001), 282-284. Among those who think that giving an account of biological "function," or what an adaptation is for, is a job for the conceptual analyst are Bigelow and Pargetter (1987), 188; Boorse (1976), 74; Ernest Nagel (1977), 284; Searle 1995, 16-19; and Wright (1976), 97. I interpret Papineau as meaning by "merely possible" cases those that are, or are in, logically or conceptually possible worlds that are not physically possible. Papineau aside, this is how I understand the relevant sort of reduction. Thus there are non-actual PPW's in which the reduction might for all we know be false, making the reduction sensitive to non-actual cases even though it is about an actual-world kind.

[11] Cf. Post (1995), 88-90. Such identity claims, like realization claims, do have some modal force (contrary, perhaps, to Papineau's account). The most plausible modality, and not only from Papineau's point of view, is truth in the relevant proper subsets of the PPW's. Again cf. Brandon (1990), 144-149, on the empirical presuppositions

lightning by conjuring "merely possible" worlds in which the identity claim fails — or the realization claim — would be irrelevant.

Summing up, a Papineau model of scientific reduction involves assuming a bridge hypothesis for the sake of argument, one which equates a folk phenomenon (a phenomenon under a folk description) with a "theoretically interesting" affair (in the sense of claiming they are either identical in standard conditions or perhaps only equivalent in them). This trial hypothesis is at least a biconditional, the modality of which is truth in the PPW's in which the relevant actual-world standard conditions obtain. The biconditional not only connects the folk role with its hypothesized realizer, it supports the role-realization claim. The aim is not conceptual analysis but to understand how folk roles are realized in the actual PPW; "merely possible" cases are irrelevant. Nor is a rigid-designator semantics required; indeed, I would add that insofar as a rigid-designator semantics involves claims of identity in all logically possible worlds, it is at odds with the irrelevance of such "merely possible" worlds. The reduction may claim identity, but only a relatively weak contingent identity — identity in the relevant proper subset of the physically possible worlds. By so classifying the folk phenomenon as a plasma-theoretic affair, we relate it to the basic laws governing physical causation and improve our understanding of the phenomenon's manifest everyday properties within the framework of a powerful, unifying explanatory theory that tells us about its underlying nature.

These, I take it, are the features of a Papineau model of scientific reduction that are relevant here. They will play a key role when we apply the model to the primitive normativity involved in an adaptation's being for this or that. But first a bit more about method.

III. More about Method

Given only a description of a high-density plasma purely in the vocabulary of plasma physics, one cannot infer that it is ball lightning or that it has this or that higher-level or folk property — say odoriferous, hissing, or red. Just as, in conformity to Hume's Law, one cannot infer an ought given only a purely descriptive is, so too one cannot infer ball lightning or an odor or hiss or color given only a description purely in the plasma-theoretic vocabulary (or, for that matter, folk water or its properties given only a description purely in the vocabulary of quantum chemistry). But to conclude from this, in line with what many do in the case of normativity, that the theorist has not successfully reduced the would-be objective ball-lightning and its properties to high-density plasma properties would miss the point. Plasma theorists are not out to infer ball lightning or its higher-level properties given only a description in the plasma-theoretic vocabulary. Rather, they provisionally propose — hypothesize, posit, try on for size, assume for the sake of argument — a bridge principle to the effect that key ball-lightning properties equate with certain plasma properties in standard conditions. Like all such bridge principles, this one

---

of applicability of the Principle of Natural Selection. Such presuppositions amount to actual-world standard conditions for the applicability of PNS.

includes vocabulary from both the physical level and the higher level.[12]

Thus it would seem to follow that Hume's-Law objections are irrelevant. Yet some philosophers would deny that it does.[13] After all, consider the plasma-theorist's bridge equivalence claim, which is expressed by a relevantly modalized biconditional, say by

BL. *x* is ball lightning if and only if *x* is a high-density plasma of kind k,

where BL is meant to be true in the PPW's in which the relevant standard conditions obtain. Once given this biconditional, plasma theorists can infer that *x* is ball lightning from a purely plasma-theoretic description of *x* as a high-density plasma of kind k. Without some such bridge principle, one cannot infer that something is ball lightning — or that it has such folk properties as being odoriferous, hissing or red — given only a description purely in the plasma-theoretic vocabulary. Hence it begins to look as though one who doubts plasma-theoretic accounts of ball lightning could indeed run a Hume's-Law argument against them. After all, such accounts must eventually endorse an inference, *underwritten by BL*, from a description in the plasma-theoretic vocabulary to the folk-vocabulary conclusion that *x* is ball lightning — and has such folk properties as being odoriferous, hissing, red. Since plasma-theorists justify this inference by assuming the biconditional BL, surely they beg the question against anti-plasma theorists, who reject any principle like BL which enables inferences to folk ball lightning from purely plasma-theoretic descriptions. It does no good to replace the inference with a conditional.

One problem with the objection is this. If some such Hume's-Law objection works against the plasma theory of ball lightning, why not also against the H_2O theory of water, or indeed against any scientific reduction of a higher-level phenomenon to a lower? After all, in every such reduction some bridge principle is assumed in order to enable the relevant inferences, a principle that is itself a biconditional or supports one. So it looks as though something has gone wrong with Hume's-Law objections in this context.

The culprit appears to be the assumption that using a biconditional like BL to underwrite the inference of its left-hand limb from its right must always be question-begging. The idea seems to be that because BL amounts to the theory at issue, or at least is a crucial component of it, using BL as a premise in an argument in support of the theory amounts to assuming at least a crucial component of the theory, which is question-begging. But consider. Well before the

[12] Jackson (2001), 656-659, appears to overlook, or at least to underestimate, the role of such bridge principles in (i) providing an a posteriori passage from one kind or level of phenomena to another, and in (ii) thereby making conceptual analysis unnecessary. Jackson thinks that given an account of the nature of a gas which is complete in purely [statistical-]mechanical terms, "There is nothing else relevant to be learnt about gases" — as if there were nothing relevant to be learnt about gases, or ball lightning, from the a posteriori bridge principles that successfully connect their key folk properties with their purely physical properties, or indicate how a folk role is realized by this or that physical state.

[13] As does an anonymous reader, who presses the following objection at length.

general theory of relativity (GTR) was regarded as at all well confirmed, GTR itself was assumed as a premise in an extended argument in its own support. This came of GTR's being used as a premise in a sub-argument — of the extended argument — for the conclusion L that light rays from a distant star are bent to a certain specific degree as they pass by the sun. L was subsequently confirmed by independent observational evidence *E*, and L thus confirmed was then taken as providing support for GTR. The basic structure of this extended argument is that theory *T*, conjoined with auxiliaries *A*, entails *C*; *C* is supported by independent evidence *E*; hence *C* provides support for *T*.[14]

The same basic structure applies when plasma theorists of ball lightning use the biconditional BL as a premise in a sub-argument for the conclusion that in standard conditions, a certain kind of high-density plasma will have key folk properties of ball lightning, so that if this conclusion is born out by independent experiment, it provides support for the plasma-theory of ball lightning. Again a trial thesis serves as a premise in an extended argument in support of the self-same thesis. This sort of "circularity" appears to be what Robert Adams (1999) has in mind in passages that meet with Boyd's approval for allowing "that the defining natures of moral terms sought by a realistic moral theory need not be non-circular, since they are not offered as introducing definitions but as accounts of the natures of things presumed to exist."[15]

In such cases, which could easily be multiplied, it is not question-begging to use the theory or the biconditional in question to underwrite an inference in the above sort of extended argument in support of the self-same theory or biconditional, in particular an inference from one limb of the biconditional to the other. The main reason why it is not question-begging is that in the relevant contexts the theory and the biconditional are advanced in such a way as to be at significant epistemic risk. They are held to be liable to refutation, or to substantial revision, or at least to anomalies that demand explanation, when the inferred conclusion is not born out by independent observation or other relevant testing or critique.[16] It thus appears that without violating Hume's Law naturalists can use the same sort of extended-argument structure on behalf of a realist account of normativity. Or rather they can unless it is shown that there is something special about normativity which renders this sort of argument structure inadmissible.

Open-question arguments appear to be designed in part to show just that. According to

---

[14] Post and Turner (2000) reply to foundationalist objections that this sort of extended argument is inadmissibly circular. See also Brown (1993) and Brown (1994). In actual practice the fine structure of the extended argument is of course more complex, but in ways that have no substantive effect on the present discussion.

[15] Boyd (2003b), 27. For the same reason, using Millikan's theory of meaning and reference, even though it presupposes the targeted normativity, is not question-begging. Her theory and the one developed here would form a package "not offered as introducing definitions but as accounts of the natures of things presumed to exist" — including meaning, reference and what these and other (ultimately) selectional affairs are for.

[16] Again cf. Brown (1993) on how "a theory-laden observation *can* test a theory," and Chalmers (2003) on the theory-dependence of the use of instruments in science.

Moore's open-question argument, we can easily imagine ourselves both recognizing that some factual condition *C* obtains (say that an act or policy *x* would conduce to the greatest happiness of the greatest number) and nonetheless asking meaningfully — or, as he says, "with significance" — whether *x* has normative property *N* (say whether *x* is good). Since this is an open question, it "shows clearly that we have two different notions before our minds"; therefore, *N* and *C* cannot be identical or equivalent, let alone their predicates synonymous (Moore 1980, 15ff). Hence their would-be identity or equivalence is of no use even as a trial hypothesis.

One problem with this line is that we can likewise imagine ourselves being told that something is a high-density plasma of kind k and nonetheless asking "with significance" whether it is ball lightning. We do indeed have two notions before our minds: the folk notion of ball lightning and the purely physical notion of a high-density plasma. Yet to conclude from this that ball lightning cannot be equated with, or is not realized by, a high-density plasma would miss the point of the plasma theorist's bridge principle BL. The aim is not to capture the meaning of the folk term, so that the biconditional would express a meaning equivalence and/or could be used to support a claim of logical or conceptual equivalence or identity. Rather, the aim is to capture how folk roles are realized in actual-physically-possible-world standard conditions; "merely possible" cases are irrelevant.

This is not to say that Moore would himself conclude that ball lightning cannot be reduced to a high-density plasma. Rather, he would owe us some justification for the double standard — that such theorizing is OK for ball lightning but not for normativity. Until this double standard is justified, anti-naturalists would not have shown by these means that there is something special about normativity which renders the above extended-argument structure inadmissible, and along with it Papineau-style reduction.

Granted, Moore's is hardly the latest in open-question arguments. But the new, improved open-question arguments likewise fail against a Papineau model of scientific reduction to the extent that they too appeal to "merely possible" cases (that is, cases that are, or are in, logically or conceptually possible worlds that are not physically possible). What I take to be the leading such argument — Horgan and Timmons (1992b), endorsed by Hare (1995) — is aimed at reductions that presuppose a rigid-designator semantics for the reduced term, whereas a Papineau model (among others) presupposes no such thing, indeed is at odds with any account that presupposes the relevance of merely logically or conceptually possible worlds as counter-examples. Granted, there is at least one open-question argument, namely Adams's, which may not to appeal to the "merely possible" (it is hard to tell). But because the "critical stance" on which his argument turns is meant to be characteristic of *ethical* thinking, the argument would appear to have no bite in the case of the *non*-ethical normativity involved in a biological adaptation's being for this or that.[17]

Summing up, it begins to look as though relevant features of the reductive method

[17] See Adams (1999), 77-78, and Adams (2003), 127.

Papineau endorses could free naturalist normative realists from the oppression of Hume's-Law objections, open-question arguments, and counter-exampling by weird worlds or "merely possible" cases. This would improve the prospects of a naturalist realism about the normativity involved in the targeted selectional properties. In order to see whether such an account might indeed be constructed, and to deal with objections, we need to have before us an explicit, concrete reduction of the targeted kind of normativity. To this we now turn.

IV. Reducing the Normativity

According to Brandon's rigorous account of adaptation, it makes sense to think of an adaptation as for something.[18] Of course a given trait might not be for anything, in which case a what-for question is out of order. But as Brandon says, "Whenever we hypothesize that some trait is an adaptation, it makes sense to inquire about its function," what it's for. I would add only that the assertion that an adaptation is *for* this or that appears to be a normative assertion, as is at least implicit in Brandon's account (and made explicit by him in conversation). To say that the heart is for pumping blood is ordinarily to imply that even when a given heart cannot possibly pump blood, nonetheless pumping blood is what it is for, what it is supposed to bring about, what it should or "ought" to do. We call a heart "bad" or "defective" when it cannot do what it is for, distinguishing between what it actually does or is disposed to do and what it should or "ought" to do. These are among the characteristic features of the folk role of something's being for this or that.

Note that this normativity, this "ought," is a *thin* kind of "ought" or "prescription" — the kind involved in the folk role of something's being for this or that. When the folk say that the thermostat is for keeping the house at constant temperature, they mean that's what it's supposed to do, what it ought to do; if it doesn't, it's defective. Likewise, when they say that your heart is for pumping blood, they mean that's what it's supposed to do, what it ought to do; if it doesn't, it's defective. This *thin* normativity is obviously not *moral* normativity; thermostats and hearts are not morally obligated to do what they're for. This thin normativity is what Mackie has in mind when he discusses "functional words" — words that refer to what something is for — in the course of arguing that, as in the case of moral normativity, even the normativity involved in functional attributions "always imports some reference to something like interests or wants," hence cannot be objective (Mackie 1977, 15, 53-58). This distinction between "thin" and "moral" normativity will prove crucial.

Now, what theoretically interesting objective affair — an affair that among other things does not involve "something like interests or wants" — might realize the folk role of 'adaptation *A* is for *E* '? An obvious candidate is the causal/mechanical selectional affair of *E*'s being the physical effect in virtue of which *A* was selected for. So let's try assuming, for the sake of argument, the following bridge principle. Where *A* is an adaptation,

DFOR. *A* is directly *for E* (normative sense) if and only if *E* is the effect of *A*'s past

[18] Brandon (1990); all quoted Brandon passages below are from pp. 139, 165, 185-89 unless otherwise noted.

instances in virtue of which *A* was selected for.[19]

The qualification 'directly' is necessary because many adaptions, in addition to being directly for something in the way characterized by DFOR, can also be for other things in ways not characterizable by DFOR. The imprinting mechanism in a newly-hatched chick is an adaptation directly for imprinting Junior on its mother. But the mechanism is thereby also for imprinting Junior on the here-now particular individual that is Junior's mom — call her Henna. Since the here-now Henna appears nowhere in the evolutionary history, imprinting on *Henna* can't be the effect of the mechanism's past instances in virtue of which it was selected for. It follows by DFOR, as it should, that Junior's imprinting mechanism is not *directly* for imprinting on Henna.[20]

Of course there are philosophically more urgent questions about DFOR, certainly the following six, each of which is raised by one or another of the crucial arguments against any naturalist normative realism.

(1)

Can a theory that presupposes DFOR conform to Hume's Law? It helps to begin with Darwin's theory of adaptation, which tacitly relies on bridge principle

DA. Biological trait *A* is an adaptation if and only if, ancestrally, *A* had some effect in virtue of which *A* was selected for.

In a bit more detail, this is to say that trait *A* is an adaptation if and only if, ancestrally, *A* had some effect *E* (often enough) such that there was selection for having *A* and *A* conferred a fitness advantage because it had effect *E*.[21]

Next, insofar as the folk notion of an adaptation is the notion of a modification made to suit a purpose, the notion is normative. Hence *one way to construe Darwin's achievement is as having shown, among other things, how this folk normative role is realized in the biological world by natural selective processes*, not by a knowing designer. To this day, evolutionary biologists maintain DA (or its equivalents) in support of what amounts to the claim that the folk role of 'adaptation' is realized in the actual world by a certain kind of natural selective process. Neither DA by itself nor *A*'s being an adaptation is meant to be inferred given only the fact that

---

[19] Of course there can be more than one effect in virtue of which a trait A was selected for, but talking as if there were just one simplifies the exposition, as do some other idealizations that likewise make no substantive difference to the argument. Note too that A's past instances need not always have had effect E, or even very often, just often enough for there to have been selection for A.

[20] Nonetheless, the mechanism is for doing so in a derivative sense explicated by Post (2001), §§6.1-6.2, drawing on Millikan (1984).

[21] Cf. notes 8 and 19, above.

E is the effect of *A*'s past instances in virtue of which *A* was selected for. DA is a bridge-principle assumed for the sake of argument and meant to be evaluated in light of what follows from it. A Hume's-Law objection would be irrelevant.

Now suppose we follow evolutionary biologists in calling the effect *E* in virtue of which *A* was selected for *A*'s function, or what *A* is for. Suppose further that we heed the above distinction between being directly for something and being indirectly for it. Under these suppositions (conjoined with DA), biological trait *A* is an adaptation if and only if *A* is for some *E*, namely the effect in virtue of which *A* was selected for. Furthermore, and again under these suppositions, it follows (i) that DA implies

> DAfor. Trait *A* is an adaptation directly for *E* if and only if, ancestrally, *E* is the effect of *A*'s past instances in virtue of which *A* was selected for,

and (ii) that DAfor implies DA, hence (iii) that DAfor is equivalent to DA. Since DFOR is equivalent to DAfor, DFOR is likewise equivalent to Darwin's DA, under the foregoing suppositions; if running a Hume's-Law argument against Darwin is irrelevant, so too is it irrelevant against DFOR. Moreover, as seen in Section III, to assume such bridge principles for the sake of argument is not question-begging, nor do open-question arguments appear to show that there is something special about normativity which renders the relevant extended-argument form inadmissible (see further Subsection (3) below). Thus in view of its inferential provenance, DFOR conforms to Hume's Law in basically the same way as bridge principles DA and BL do in their respective theories — Darwin's theory of adaptation and the plasma theory of ball lightning.

Unfortunately this conclusion appears to conflict with an objection Papineau presses in another paper:

> Whatever norms are, I take it that they must involve some kind of *prescription*, some kind of implication about what *ought* to be done. This simply isn't true of biological facts (Papineau 1999, 21n5).

Yes, of course, the biological facts — including such facts as that *E* is a certain fitness-enhancing effect of *A*'s past instances — simply do not imply, by themselves, any prescription for *A*, or anything about what *A* ought to do. But it is equally true, as seen, that the plasma-theoretic facts — those expressed in the vocabulary purely of plasma physics — imply nothing by themselves about the folk role of 'ball-lightning' or ball lightning's folk properties. In both cases some bridge principle is required, which, when conjoined with the lower-level facts, effects the implication. We allow this strategy in the latter case, why not in the former? What would justify the double standard according to which inference effected by bridge principles assumed for the sake of argument is OK in the non-normative cases but not the normative?

Papineau might defend the double standard by pressing an argument he advances immediately after the one quoted just above:

> these biological "norms" aren't norms in any prescriptive sense. It is a vulgar, and indeed dangerous, error to infer, from the premise that X has been biologically designed to Y, that in some sense X *ought* to Y. My knuckles have arguably been biologically designed to hit people with, but it doesn't in any sense follow that I ought so to use them (Papineau 1999, 21n5).

Again, yes, of course, given only the premise that X has been biologically designed to Y, one cannot infer that in some sense X *ought* to Y. The premise must be conjoined with some suitable bridge principle, say with DFOR. So this part of Papineau's further argument goes the way of his previous one, leaving us still to deal with the implied double standard.

Consider, then, the second part of Papineau's further argument, the bit about knuckles: given only the premise that knuckles were biologically designed to hit people with, "it doesn't in any sense follow that I ought so to use them."[22] Right, absolutely; it doesn't even follow that I am so much as permitted so to use them. But DFOR implies nothing to the contrary. Rather, in conjunction with the premise that *E* is the effect of *A*'s past instances in virtue of which *A* was selected for, DFOR enables us to infer that *A* is normatively for *E*, hence that *in some sense A* is supposed to, or should, or "ought" to effect *E*.

Which sense? This "ought," as noted before, is the *thin* kind of "ought" or "prescription" involved in the folk role of something's being for this or that. DFOR is about this thin normativity, not *moral* normativity. In view of the gap between the two, DFOR does not imply, or even support, an inference from one to the other — say from (i) *A*'s having been designed by selection to *E*, so that *A* is for or is supposed or "ought" to *E*, to (ii) the containing organism's having a moral obligation or permission to use *A* to that end. Such an inference would equivocate between 'ought' in the thin sense and 'ought' in some moral sense. Anyone who supposes that according to DFOR, my knuckles' having been designed to hit with would imply that I ought, or even that I am permitted so to use them, commits this equivocation.

Papineau elsewhere argues that moral utterances do not express beliefs in the first place (Papineau 1993, 198-199). But even if his 1993 argument works in the case of *moral* normativity, Papineau would have to show that it works as well for the *thin* kind of normativity DFOR is about. That this can be shown seems unlikely, since his 1993 argument (i) relies on a Hume's-Law objection ("Hume long ago observed that you can't infer an 'ought' from an 'is'" (Papineau 1993, 199)), and (ii) assumes that the aim of reduction is to preserve meaning ("reductionist readings . . . seem clearly not to do justice to the intended meaning of the moral terms" (Papineau 1993, 198)). Both (i) and (ii) are irrelevant, according to Papineau's model of reduction, as seen in Section II. To restore their relevance, one would have to show that although the model applies to non-normative folk roles, it does not apply even to the thin normativity DFOR is about — another version of the double standard at issue.

---

[22] Harman has made the same argument, using a different example, in comments on an earlier version of this paper at an APA Symposium.

(2)

What is DFOR's modality? Again it helps to begin with Darwin's theory of adaptation and its bridge principle DA. Evolutionary biologists maintain principles like DA in support of what amounts to the claim that the folk role of 'adaptation' is realized by a certain natural selective process in the relevant actual-world standard conditions.[23] The modality of such principles is truth in the physically possible worlds in which the relevant standard conditions obtain, namely certain actual-world conditions that enable adaptation by natural selection. Evolutionary biologists would not be amused by counter-examples in Swamp-Man or Twin-Earth worlds (or such imaginary places in the actual world).

In order to counter-example DA, one must show that the scenario or world *W* in which DA is supposed to be false is not only (i) physically possible but (ii) such that the relevant standard conditions obtain in *W*. That some philosopher's imagined world *W* satisfies these two conditions cannot be inferred from (a) the fact (when it is one) that *W* is conceivable, and not even from (b) the fact (when it is one) that it is conceivable that *W* is both physically possible and such that the relevant standard conditions obtain in *W*. [24] The inference from (a) or (b) to (i) or (ii) is a non-sequitur. The method of counter-exampling by invoking weird worlds or "merely possible" cases has no force against Darwin's bridge principle DA, however useful such worlds or cases may be for teasing out the structure of everyday thinking.

As with DA, so with DFOR. Insofar as DFOR is advanced in the same spirit as DA, one would expect DFOR's modality to be the same — truth in the relevant proper subset of the PPW's; what happens in "merely possible" worlds is a don't-care. Furthermore, recall that DFOR is equivalent to DA (under the suppositions that the effect *E* in virtue of which *A* was selected for is what *A* is for, and that we heed the distinction between being directly and being indirectly for something). In view of this inferential provenance of the equivalence, DFOR's modality is the same as DA's. So again what happens in merely logically or conceptually possible worlds is irrelevant to DFOR. Nor is a rigid designator semantics in order, being at odds with the aims of a Papineau model of reduction (as seen at the end of Section II). Any argument against DFOR which presupposes a rigid-designator semantics for the reduced term, as does, say, the new, improved open-question argument advanced by Horgan and Timmons (1992b), is irrelevant.

(3)

What else might be said of open-question arguments? Once more it helps to begin with Darwin. His aim was not to propose necessary and sufficient conditions of the meaning of the folk term 'adaptation', and not to conform to our pre-theoretical intuitions about either the usage of the term or the nature of the phenomenon it refers to (if it refers). We folk can easily imagine

---

[23] Again, see Brandon (1990), 144-149, on the presuppositions of the applicability of the Principle of Natural Selection, which amount to actual-world standard conditions for its applicability.

[24] Contrary to Moser (1992), 73-74.

ourselves both recognizing that some trait *A* is the result of a biological selectional process and nonetheless asking "with significance" whether *A* is an adaptation — especially if we have strong intuitions to the effect that adaptation (analytically) requires a knowing designer.[25] And this does show that we have two different notions before our minds — a folk notion and a biological selectional notion. But it does not follow that Darwin's theory of adaptation should be rejected.

So too in the case of DFOR. We folk can easily imagine ourselves recognizing that *E* is the effect in virtue of which *A* was selected for, and nonetheless asking "with significance" whether *A* is normatively for *E*; we do indeed have two different notions before our minds. But from this fact it does not follow that *A*'s being for *E* is not realized by *E*'s being a certain effect of *A*'s past instances. Nor does it follow that the folk role of *A*'s being for *E* has been eliminated — eliminated because, say, being selectionally for *E* does not entail a knowing designer whereas the folk notion of *A*'s being for *E* does. To suppose that either conclusion follows would miss the point of the reduction. Open-question arguments appear to be as irrelevant here as in the case of theories like the plasma theory of ball lightning, the H_2O theory of water — and Darwin's theory of adaptation, even when an adaptation is conceived normatively as a modification made to suit a purpose.

So too are open-question arguments irrelevant insofar as they presuppose that moral judgments are intrinsically motivating, hence in that sense are "prescriptive." Even if such internalism should prove true of moral normative judgments, it is not true of the non-moral normative judgment that biological adaptation *A* is for *E*. When I judge that my heart ought to pump blood, my heart is not thereby motivated to palpitate, nor am I. If Thomas Nagel (1970, 8) and Darwall, Gibbard and Railton (1997, 4) are right that this internalism is what underlies the seeming persuasiveness of open-question arguments, then such arguments again have no force in the case of an adaptation's being normatively for this or that.

(4)

Does DFOR preserve the critical gap between what something actually does or is disposed to do and what it is supposed to do? Specifically, where *x* is an instance, or token, of adaptation type *A*, does DFOR entail that *x*'s having the normative property of being for *E* neither equates with nor follows from what *x* actually does or is disposed to do? By DFOR, and where *x* is a token of *A*,

> *x* has the normative property of being for *E* if and only if *E* is the effect of *A*'s past tokens in virtue of which *A* was selected for.

The right-hand limb of this biconditional says nothing about the actual behavior or dispositions of token *x*. Instead, the right-hand limb is about an adaptation type *A*, *A*'s *past* tokens, and their effect *E* in virtue of which *A* was selected for. It follows that even when a token *x* of *A* does not

[25] As apparently does Searle (1995), 16-19.

have effect *E*, or is not even so much as disposed to have effect *E*, it remains the case that *x*'s having the normative property is determined by *E* 's being the effect of *A*'s past tokens in virtue of which *A* was selected for. Such past affairs cannot be affected by what *x* actually now does or is disposed to do. Hence *x*'s having the normative property of being for *E* neither equates with, nor follows from, nor is affected by what *x* actually does or is disposed to do. The gap is preserved.

(5)

What of Mackie's argument from queerness (AQ)? According to AQ, so-called objective or real normativity would be a queer sort of thing, because its relation to what is objectively the case would be quite mysterious. No allegedly objective normative property *N* of an item *x* is inferable from non-normative affairs, as Hume taught us. Nor can *N* be reduced to such affairs, as Moore's open-question argument taught us. Furthermore, talk of supervenience of a specific normative property *N* on certain specific facts is itself in need of naturalistically acceptable explanation.[26] Since no other relation has been spelled out that works, the allegedly objective normativity must be queer indeed. Mackie concludes,

> How much simpler and more comprehensible the situation would be if we could replace the [normative] quality with some sort of subjective response which could be causally related to the detection of the natural features on which the supposed quality is said to be consequential (Mackie 1977, 41).

What we think of as objective normativity is just our subjective valuation projected onto the value-neutral real world, a process Mackie calls objectification, likening it to what Hume calls the mind's "propensity to spread itself on external objects."[27]

Clearly, AQ contains a couple of sub-arguments that amount respectively to a Hume's-Law objection and an open-question argument, which are rendered irrelevant by a Papineau model of reduction, as seen in Subsections (1) and (3). By so much, then, is AQ likewise rendered irrelevant. That leaves the rest of AQ: the challenge to naturalists to explain (i) the needed non-mysterious relation between the would-be objective normative properties and the relevant natural affairs, and (ii) why such affairs subvene the would-be objective fact that some given specific individual has this or that specific normative property.

As regards (i), recall that DFOR is a biconditional. So there is at least a relation of equivalence between the would-be objective normative properties and the relevant natural

---

[26] Thus for present purposes I follow Horgan and Timmons (1992a) in interpreting Mackie's AQ.

[27] Mackie (1977), 42. Another part of Mackie's argument — that something's being morally good, and/or the judgment that it is, is intrinsically or necessarily motivating and therefore must be ontologically queer from the point of view of naturalism — is irrelevant so long as we are discussing non-moral normativity. Likewise irrelevant is a related internalism, according to which "accepting the judgment that something is morally good must, by itself, provide some reason for choosing or preferring it," as Boyd (2003a), 532, puts it, but only to reject it.

affairs, and its modality, as seen in Subsection (2), is the relatively weak modality of truth in every PPW in which the relevant actual-world standard conditions obtain. Furthermore, equivalence in a set *S* of worlds implies determination in that set; if *p* is equivalent to *q* in every world *W* in *S*, then *p* determines *q* in every *W* in *S*, and vice versa (for, given the equivalence in *W*, the truth value of *p* in *W* fixes the truth value of *q* in *W*, and vice versa). When one adds that *q*, say, has the relevant sort of explanatory and compositional priority over *p*, it follows that in such worlds, *p* supervenes on *q*. Likewise, in light of the equivalence DFOR together with the explanatory and compositional priority of its right-hand limb over its left, the normative matter of *A*'s being for *E* is — in the PPW's in which the relevant standard conditions obtain — determined by, and supervenes on, *E*'s being the effect of *A*'s past instances in virtue of which *A* was selected for. All this is compatible with, and indeed supports, the claim that the latter realizes the former.[28]

As regards (ii), the matter of *why* a relation of supervenience obtains between the given specific normative property of being for *E* and the relevant specific natural affairs can be explained as follows. Focused supervenience obtains between them because the former is equivalent to the latter in the relevant PPW's, and the latter has the relevant sort of explanatory and compositional priority over the former. So too for why the relations of focused determination and realization obtain between them. It follows further that there is not just one non-mysterious relation between the objective normativity and the relevant specific natural features, but at least four: equivalence (in the relevant worlds), focused supervenience and focused determination (again in the relevant worlds), and realization (also in the relevant worlds). Contrary to queerness arguments, and thanks to DFOR deployed in accord with a Papineau model of reduction, the normativity of a biological adaptation's being for this or that is no more queer than ball lightning — or adaptation by natural selection.

(6)

Would the so-called objective normative properties be unacceptably epiphenomenal, playing no significant causal or explanatory role, and no predictive role either? Papineau seems to think so:

> selectional classifications . . . clearly don't introduce causally efficacious properties. To classify something as a biological heart . . . implies that its ancestral . . . effects led to its preservation in the species, but says nothing about the physical make-up which enables it to produce those effects (Papineau 2001, 287).

[28] The foregoing relation of "focused" supervenience or determination is just the global supervenience/determination relation with the range of its quantification over worlds restricted to the PPW's in which the actual-world conditions obtain that are relevant in a specific case. Focused supervenience and focused determination, like global, are nonreductive, in the sense that they do not entail, though they are compatible with, property-property equivalence and identity in the relevant worlds; not all a thing's higher-level properties need be identical or even equivalent to, or realized in, certain (compounds) of its own base properties (whether intrinsic or relational). Cf. Post (1995), 89-93.

Clearly, by DFOR, classifying an adaption *A* as being normatively for *E* is likewise a selectional classification. It would follow, according to Papineau, that such normative classification does not introduce causally efficacious properties. And yet it does. Granted, such classification tells us nothing by itself about *A*'s purely physical make-up. But it does tell us that if a token *x* of an adaptation *A* — say my own heart — is operating as designed by selection and in design conditions, then *x* is able to produce circulation of my blood, and we can predict that it will do so (more on "design" below). Furthermore, such classification tells us that whatever *x*'s physical make-up, if *x*'s physical make-up is operating as designed and in design conditions, then this make-up enables *x* to produce circulation of my blood.

In general, under the assumption that a token *x* of an adaptation *A* is operating as designed and in design conditions, *x*'s having the normative property of being for *E* implies that *x* has the causally efficacious property of being able to produce *E*, and that *x* will do so. Under the assumption, which is frequently warranted, we can not only infer that *x* is so enabled, but explain *why* it is. In this conditional manner, selectionally typing *A* as being normatively for *E* introduces a property that is causally efficacious — causally efficacious when *A* is operating as designed and in design conditions.[29]

As regards the notion of "design," Papineau is right that these "designeds" cannot be explained by everyday thinking. But as he himself goes on to say, surprisingly,

> It is specifically here that selectional typing adds theoretical power to everyday thought. It tells us about the underlying nature of [*A*'s] design, and thereby directs us to the past selectional processes which fixed [*A*'s] real purposes (Papineau 2001, 288).

Thus Papineau himself appears to believe that the folk role of "design" is realized in certain past selectional processes. Furthermore, because these processes are physical processes, selectionally typing *A* as designed to produce *E*, hence as normatively for *E*, not only introduces properties that are causally efficacious (in the conditional way described above). It does so by relating them, as Papineau would require, to "the basic laws governing physical causation," thereby enabling us to understand such properties within the framework of "a powerful, unifying explanatory theory" which "tells us about the underlying nature" of *A*'s being normatively for *E*. In this way, selectionally typing *A* as normatively for *E* "adds theoretical power to everyday thought" (Papineau 2001, 282, 285, 287).

V. Concluding Remarks

The arguments of Subsections (1)-(6) are not meant by themselves to show that the theory built around DFOR, when deployed in accord with a method relevantly like Papineau's and a theory of meaning and reference relevantly like Millikan's, provides a fully adequate

---

[29] Cf. Post (2001), §5.6 on Harman, with regard to DFOR's empirical adequacy compared to that of competing bridge principles or theories. For comments on earlier versions of this paper I am greatly indebted to Allen Coates, Fred Dretske, Gilbert Harman, Stephen Schiffer, Brian Skyrms, Ken Taylor, Derek Turner, and the referees.

naturalist realism about the targeted normativity. As emphasized in Section I, my aim is not to argue that we should accept the theory, but to argue that if we adhere to the relevant method and theory of meaning and reference, then the crucial objections to naturalist realism about the targeted normativity are disarmed, thereby greatly improving its prospects. To show that the theory should indeed be accepted would require substantially more argument than is possible here. Meanwhile, we may conclude at least the following.

In order to make good on their claim that no norm can come from biological facts, naturalists like Papineau would need to explain why reductionist models relevantly like his should apply to non-normative folk roles but not to the folk-role primitive normativity involved in a biological adaptation's being for this or that. They would need to justify this double standard. Yet the prospects of doing so appear dim. The crucial objections to any naturalist-realist attempt to reduce norms to facts — from Hume's-Law objections to charges of epiphenomenalism — are disarmed by the very same models, largely by rendering the objections irrelevant, as explained in (1)-(6). To be sure, someone might revise one or more of these objections, or develop an entirely new one, in such a way as to justify the double standard. But until then the problem remains: it appears one cannot in good conscience apply the model to non-normative folk roles while refusing to do so when it comes to the primitive normative matter of a biological adaptation *A*'s being for *E*.

It follows, as a special case, that one cannot in good conscience apply a model of reduction relevantly like Papineau's to non-normative folk roles and yet reject it for the normativity involved in selectional content. Until the double standard is successfully defended, if it can be, it looks as though the best bet for teleosemanticists who want to explain the normativity of selectional content is, ironically, to apply a Papineau model to such content along the lines developed in Sections III-IV (again, see footnote 15).

Are there any implications in all this for *moral* normativity? No direct implications, surely. As repeatedly emphasized, the normativity involved in an adaptation *A*'s being for *E* is a thin kind of normativity, not moral, so that a moral ought or permission cannot be inferred from *A*'s being for this or that. Even if knuckles had been selectionally designed to hit people with, in no sense would it follow that morally I ought, or am even so much as permitted, so to use them.

Nonetheless, there may be some indirect implications for moral normativity. Here are two possibilities. First, it will not have escaped notice that the objections to the realist DFOR considered in Subsections (1)-(6) have long been used against naturalist moral realism. To the extent that the objections in the moral case have the same form and presuppositions as those in the non-moral, they too would be disarmed by applying a Papineau model of reduction. This of course assumes that it is possible to apply the model to the moral case without inferring moral norms from selectional. I think it is not only possible but promises a robust naturalist moral realism, though that is another story, far too long to attempt here. But even if applying the model to the moral case should prove a dead end — say by virtue of the semantics of folk moral terms being relevantly different from the semantics of 'adaptation *A* is for *E* ' — Papineau and like-minded others would still need to explain why the model is supposed to apply to non-

normative folk roles but not to the primitive folk-role normativity involved in a biological adaptation's being for this or that. Otherwise their anti-reductionism about such normativity would have to be substantially qualified if not retracted.

Second, Mackie and most if not all other non-cognitivists about normativity believe that the semantics of normative claims should be unified, in the sense that, as we saw Copp put it, "If our semantics of moral claims is non-cognitivist, then . . . the semantics of all normative claims must be non-cognitivist." If this sort of unified semantics is assumed, and supposing that a Papineau model applied to DFOR yields a successful cognitivist account of an adaptation's being normatively for *E,* it would follow that our semantics of moral claims must be cognitivist as well.

REFERENCES

Adams, Robert (1999). *Finite and Infinite Goods* (New York and Oxford: Oxford University Press).

Adams, Robert (2003). "Anti-Consequentialism and the Transcendence of the Good," *Philosophy and Phenomenological Research, 67,* 114-132.

Bigelow, J. and Pargetter, R. (1987). "Functions," *Journal of Philosophy 84*, 181-196.

Boorse, Christopher (1976). "Wright on functions," *Philosophical Review 85*, 70-86.

Boyd, Richard (2003a). "Finite Beings, Finite Goods: The Semantics, Metaphysics and Ethics of Naturalist Consequentialism, Part I," *Philosophy and Phenomenological Research, 66,* 505-553.

Boyd, Richard (2003b). "Finite Beings, Finite Goods: The Semantics, Metaphysics and Ethics of Naturalist Consequentialism, Part II," *Philosophy and Phenomenological Research, 67*, 24-47.

Braddon-Mitchell, D. and Jackson, F. (1997). "The Teleological Theory of Content," *Australasian Journal of Philosophy 75*, 474-489.

Braddon-Mitchell, D. and Jackson, F. (2002). "A Pyrrhic Victory for Teleonomy," *Australasian Journal of Philosophy 80*, 372-377).

Brandon, Robert N. (1990). *Adaptation and Environment* (Princeton: Princeton University Press).

Brown, Harold I. (1993). "A Theory-laden Observation *Can* Test the Theory," *British Journal for the Philosophy of Science, 44*, 555-559.

Brown, Harold I. (1994). "Circular Justifications," *Proceedings of the Biennial Meetings of the Philosophy of Science Association, 1*, 406-414.

Chalmers, Alan (2003). "The Theory-Dependence of the Use of Instruments in Science,"

Philosophy of Science, *70*, 1493-509.

Copp, David (1995). *Morality, Normativity, and Society* (New York and Oxford: Oxford University Press).

Darwall, Stephen; Gibbard, Allan; and Railton, Peter (1997). "Toward Fin de siècle Ethics: Some Trends," in Darwall, Gibbard and Railton, eds., *Moral Discourse and Practice: Some Philosophical Approaches*, ed. by (Oxford: Oxford University Press), 3-47.

Hare, Richard M. (1995). "A New Kind of Ethical Naturalism?" in P. A. French, T. E. Uehling and H. K. Wettstein, eds., *Midwest Studies in Philosophy, XX* (Notre Dame: University of Notre Dame Press), 340-356.

Harman, Gilbert (1986). "Moral Explanations of Natural Facts — Can Moral Claims Be Tested Against Moral Reality?" *Southern Journal of Philosophy*, *24*, *Supplement*, 57–68.

Harman, Gilbert (1998). "Responses to Critics," *Philosophy and Phenomenological Research, 58*, 207-213.

Herman, Barbara (2003). "A Doctrine of Ends," delivered at the Chapel Hill Philosophy Colloquium, October 10, 2003.

Hooker, C.A. (1987). *A Realistic Theory of Science* (Albany: State University of New York Press).

Horgan, Terry and Timmons, Mark (1992a). "Troubles on Moral Twin Earth: Moral Queerness Revisited," *Synthese, 92*, 212-260.

Horgan, Terry and Timmons, Mark (1992b) "Troubles for New Wave Moral Semantics: The "Open Question Argument" Revived," *Philosophical Papers*, *21*, 153-175.

Jackson, Frank (2001). "Responses," *Philosophy and Phenomenological Research*, *62*, 653-664.

Mackie, John L. (1977). *Ethics: Inventing Right and Wrong* (New York: Penguin Books).

Millikan, Ruth Garrett (1984). *Language, Thought, and Other Biological Categories* (Cambridge, Mass.: MIT Press).

Millikan, Ruth Garrett (1993). *White Queen Psychology and Other Essays for Alice* (Cambridge, Mass.: MIT Press).

Millikan, Ruth Garrett (2001). "Cutting Philosophy of Language Down to Size," in *Philosophy at the New Millennium*, ed. Anthony O'Hear (Cambridge: Cambridge University Press), 125-140.

Moore, G. E. (1980). *Principia Ethica* (Cambridge: Cambridge University Press).

Moser, Paul K. (1992). "Physicalism and Global Supervenience," *The Southern Journal of Philosophy*, *30*, 71-82.

Nagel, Ernest (1977). "Teleology Revisited: The Dewy Lectures 1977" *Journal of Philosophy*

74, 261-301.

Nagel, Thomas (1970). *The Possibility of Altruism* (Princeton: Princeton University Press).

Papineau, David (1993). *Philosophical Naturalism* (Cambridge: Blackwell).

Papineau, David (1999). "Normativity and Judgement," *Aristotelean Society, 73, Supplement*, 17-43.

Papineau, David (2001). "The Status of Teleosemantics, or How to Stop Worrying About Swampman," *Australasian Journal of Philosophy, 79*, 279–289.

Post, John F. (1995). "'Global' Supervenient Determination: Too Permissive?" in *Essays on Supervenience*, ed. Elias Savellos and Ümit Yalçin (Cambridge: Cambridge University Press), 73-100.

Post, John F. (2001). "From Is to Ought: Another Way," a revised and expanded version of "How to Get an Ought from a Biological Is" delivered at the Symposium on Evolution and Normativity, APA, April 7, 2000. Available at www.vanderbilt.edu/~postjf/fritoaw7.htm

Post, John F. and Turner, Derek (2000). "Sic Transitivity: Reply to McGrew and McGrew," *The Journal of Philosophical Research, 25*, 67-82.

Searle, John (1995). *The Construction of Social Reality* (New York: The Free Press).

Sober, Elliott (1993). *Philosophy of Biology* (Boulder, CO: Westview Press).

Stenhoff, Mark (1999). *Ball lightning : An Unsolved Problem in Atmospheric Physics* (New York: Kluwer Academic).

Sterelny, K. and Griffiths, P. (1999). *Sex and Death: An Introduction to Philosophy of Biology* (Chicago: University of Chicago Press).

Wright, Larry (1976). *Teleological Explanation* (Berkeley: University of California Press).