

AGAINST MODERATE RATIONALISM*

An increasingly popular position in epistemology is moderate rationalism. As expounded by such writers as George Bealer, Laurence BonJour, and Alvin Plantinga, this form of rationalism differs from the traditional variety found in Descartes or Spinoza principally in admitting the fallibility of intuition, which both forms take to be the ultimate source of *a priori* knowledge and justification.¹ In my view this new form of rationalism is really no better than the old one, and in what follows I shall support this assesment by criticizing the case that a leading proponent of moderate rationalism has recently made for the idea that intuition provides a satisfactory basis for *a priori* knowledge.

Traditional rationalists defended their position by an unsupported appeal to the truth of what they considered intuitively self-evident.² It seemed to them that any rational person who attentively considered such propositions as *Nothing could both have and not have a certain attribute at the same time* or *Nothing could be simultaneously both red and green all over* could simply “see” that they are true. This way of justifying the rationalist’s practice of accepting such propositions as a source of knowledge is not sufficient for moderate rationalists. If some of the propositions that distinguished philosophers have judged to be intuitively obvious are now acknowledged to be false or dubious--for example, Frege’s axiom of abstraction or the principle that every occurrence must have a cause--then the practice of accepting such propositions cannot be approved without further support. Laurence BonJour has acknowledged this point in his recent book, emphasizing that his rationalist practice must be capable of a “dialectical” defense.³ As he sees it, the initial plausibility of intuitive obviousness as a good but fallible basis for thinking something true must be supplemented by a criticism of opposing views concerning *a priori* knowledge and justification and by a rebuttal to objections raised by those who defend such views.

One might wonder how a dialectical defense of this kind can possibly show that moderate rationalism actually delivers *a priori* knowledge. Success in refuting objections to a view merely shows that the view has not been shown to be false or erroneous; and success in opposing alternative views merely shows that the alternatives are false or badly defended. Neither procedure warrants the conclusion that the defended view is actually true or acceptable.

* Forthcoming in *Journal of Philosophical Research*.

In fact, Bonjour does not attempt to prove that we have *a priori* knowledge; this is something he simply assumes. He attacks Quine's arguments against the existence of such knowledge,⁴ and he also argues that if we do not have *a priori* knowledge, we do not have any knowledge, a consequence he regards as "massively implausible from a common sense or intuitive standpoint."⁵ His case for moderate rationalism is essentially an attempt to show how *a priori* knowledge is possible. Put in contemporary terms, his dialectical argument is pretty much an "inference to the best explanation," with *a priori* knowledge (or certain examples of it) the datum to be explained.⁶ His rebuttals to objections and his attacks on alternative views function, in any case, as attempts to accomplish the two crucial tasks that this kind of inference requires. In rebutting the objections he tries to show that intuition (or rational insight, as he prefers to call it) does in fact provide a satisfactory "explanation" for our possession of the relevant knowledge,⁷ and in attacking alternative views he tries to show that it explains it better than they do.

The only significant alternative to moderate rationalism (if Quine's view is rejected) is what Bonjour calls "moderate" empiricism, a view restricting *a priori* knowledge to analytic truths. Bonjour subjects this alternative to thoroughgoing criticism. In what follows I shall argue that Bonjour's dialectical argument fails in both of the tasks I mentioned. Intuition does not adequately explain our knowledge (to the extent that we have it) of the relevant truths, and Bonjour's criticism does not adequately dispose of moderate empiricism, a version of which can account for that knowledge in a satisfactory way.

Rationalist Examples Involving Color

How does intuition account for our knowledge of *a priori* truths? Bonjour offers the following as his initial response to the proposition that nothing red all over can simultaneously be green:

First, I *understand* the proposition in question. This means that I comprehend or grasp the property indicated by the word 'red' and also that indicated by the word 'green', that I have adequate conceptions of redness and greenness.... Similarly, I understand the relation of incompatibility or exclusion that is conveyed by the rest of the words in the verbal formulation of the proposition, together with the way in which this relation is predicated of the two properties by the syntax of the sentence. Second, given this understanding of the ingredients of the proposition, I am able to see or grasp or apprehend in a seemingly direct and immediate way that the claim in question cannot fail to be true--that the natures of redness and greenness are such as to preclude their being jointly realized. It is this direct insight

into the necessity of the claim in question that seems, at least *prima facie*, to justify my accepting it as true.⁸

Just how do the natures of redness and greenness preclude their being jointly realized? Bonjour explains this seven pages later:

It is in the natures of both redness and greenness to exclusively occupy the surface or area that instantiates them, so that once one of these qualities is in place, there is no room for the other; since there is no way for the two qualities to coexist in the same part of the surface or area, a red item can become green only if the green replaces the red” (p. 108).

It seems to me that the passages just quoted disclose two incompatible models of what is going on when one supposedly sees that nothing red can also, at the time, be green. One model, the official one, is perceptual: redness and greenness are somehow presented to consciousness, and one sees their incompatibility directly. The other model is more discursive. Seeing a thing A, one realizes that it is by nature a certain sort of thing--an exclusive occupier, with respect to a certain class of properties, of a surface or area. Seeing another thing B, one realizes that it has a similar nature: it too is an exclusive occupier, with respect to the same class of properties, of a surface or area. Since one sees that A is different from B, one then *concludes* that no surface or area can be both A and B at the same time. How one realizes that A and B are exclusive occupiers in this way is not obvious on this model; but when it is achieved, the conclusion that A and B cannot occupy the same surface at the same time is reached by a patently valid form of reasoning.⁹

Of the two models, I think the second, discursive one is far more realistic than the first one. I also think that the conclusion drawn is not adequately accounted for by the moderate rationalist’s view. Nevertheless, I shall proceed to attack Bonjour’s official view that the relevant incompatibility is intuitively apprehended, leaving till later a discussion of the reasoning that might be brought to bear on such an incompatibility. To prepare the way for my attack on the official view, I want to discuss an example of my own, which makes it easier to see that the relevant incompatibilities cannot be ascertained by an act of mental insight. I shall return to the example involving red and green a little later.

Before presenting my example, I should mention that if colors are conceived of in an empirically realistic way, the idea that we can directly grasp their natures has no plausibility at all. The reason for this is that the nature of physical color is not presented to the senses in the way that common sense might sup-

pose. Its nature is as theoretical as that of H₂O, and there is considerable disagreement about how it is best described.¹⁰ To give the rationalist's position a fair run for its money, I shall therefore assume a more phenomenal conception of color. I shall assume, that is, that our conceptions of colors pertain to the qualitative features that we are aware of when we look at colored objects. What we are thus aware of will not include what is true of the micro-structure of a colored object's surface or what sensory experiences standard observers will have when they see a colored object.

It will be helpful to embed the example I want to discuss in a little story about two people, Mary and Harry, who are not philosophers but who often argue about matters of philosophical interest. They are intelligent and independent-minded, but their color-vocabulary is somewhat modest. On a visit to an arboretum, they observe a shrub with a color they have not seen before. An attendant tells them that the color of the shrub is chartreuse. On the way home they begin talking about it.

Mary says she regards chartreuse as greenish yellow and therefore a kind of yellow. Harry regards it as yellowish green and therefore a kind of green. They argue a little. Harry says they can find out who is right by consulting a dictionary when they get home. When they consult their dictionary, their disagreement remains, because the entry for "chartreuse" says "a clear light green with a yellowish tinge,"¹¹ and Mary says this doesn't describe the color of the bush they saw: that bush was not a *clear* light green. She then adds, "Maybe the color of the bush was not chartreuse; maybe the attendant didn't really know what chartreuse is. But the resolution of our problem should not depend on what anyone else says, anyway. We are concerned with how *we* are going to classify the color we saw, and our classification doesn't have to agree with the classification of anyone else--if anyone else did ever classify that particular color. I still want to say that it is greenish yellow, and since greenish yellow is a shade of yellow, I say it is yellow." Harry is stubborn and wants to say chartreuse is a kind of green, but he also wants to placate Mary; so he says "Let's say it is both green and yellow." He continues: "Things that are greenish yellow or yellowish green have varying amounts of both yellow and green, so in a sense they are both green and yellow all over. Since we (you and I) can't very well measure the relevant amounts of both these colors in chartreuse (how could this be done anyway?), it is useless to continue arguing about whether it is really yellowish green or greenish yellow. Let's just say it is both and be done with it."

Mary doesn't like Harry's proposal. "It is a necessary truth," she says, "that nothing can be both yellow all over and green all over (or yellow and green in the same place) at the same time. I read this in a philosophy book." Tom has read philosophy books, too, but he doesn't like Mary's apparent dogmatism. So he becomes philosophical in his own way. "I am not Humpty-Dumpty," he says, "but in this case we can certainly resolve our disagreement by a stipulation. We are not disagreeing about a matter of fact. We are not concerned to agree with how someone else--some "authority"--describes the color we have seen. Anyone who speaks English will agree that that color contains both green and yellow. As far as we are concerned, we can adjust our color vocabulary to accommodate the new color in four different ways. We can say it is (a) greenish-yellow, a kind of yellow, (b) yellowish-green, a kind of green, (c) both yellow and green all over (for both are admittedly there--all over), or (d) we can say that it is some new color (just call it "the problem color") that is made up of various proportions of yellow and green.

A convinced moderate realist would no doubt be unimpressed with Harry's remarks; certainly BonJour would not concede that they undermine his position. Instead of employing words as generic as "yellow" and "green" or "red" and "green," BonJour could support his case (as he did in one passage) by referring to two different shades of blue.¹² I shall postpone my discussion of possible responses to my example, for Harry's remarks, whatever their dialectical limitations, bring out certain facts that are very important for the basic issue dividing rationalists and empiricists.

Perhaps the most striking fact disclosed by Harry's remarks is that he and Mary can resolve their dispute by stipulating how they are going to use certain color words. There is nothing inherently wrong, I should say, with any of the alternatives Harry mentions. If they stipulate that the color they are concerned about is to count as a shade of yellow, they are adjusting the extension of their adjective "yellow" so that it covers this case. If they give it a new name, they are stipulating that the new name applies to it. And if they decide to call it both yellow and green, so that the bush they saw can be described as both yellow and green in their senses of these words, they are stipulating that they are now using "yellow" and "green" in such a way that their extensions overlap in a certain region of color space, a region containing the color of the bush they saw.

The dispute between Mary and Harry would not have arisen if "yellow" and "green" were not, in their dialect, significantly vague. A vague predicate, as others have observed, is one that is clearly applica-

ble to certain things, clearly inapplicable to certain other things, and neither clearly applicable nor clearly inapplicable to further things.¹³ Owing to their semantic indeterminacy for a range of possible cases, vague words such as “yellow” and “green” should not, in my view, be regarded as corresponding to discrete properties, least of all to simple and unanalyzable ones, as G.E. Moore thought.¹⁴ If a semantic correlate is needed for “yellow,” it is better conceived of as a fuzzy set whose “positive” members include numerous shades of yellow (no doubt thousands of them) of various degrees of brightness and saturation blended with variety of other colors.¹⁵ The set is fuzzy because its membership is not categorically defined: things belong to it in greater or lesser degrees.

If we take the vagueness of generic color words seriously, we must acknowledge that the reality to which we apply a word like “green”—the reality consisting of the items belonging to the relevant fuzzy sets in a significantly positive degree—is enormously complex and cannot be taken in by a single act of mental vision. We must also acknowledge that the relations between such realities cannot possibly be taken in immediately, by an act of rational insight. The grasping and apprehending that moderate realists speak of must therefore be an illusion in cases where the “properties” are as generic as red or green. If we succeed in bringing anything pertinent to mind when we think we are “grasping” greenness, we are perhaps doing what Hume says we always do: we are perhaps bringing to mind some determinate property (in this case some shade of green, such as the color of a shamrock, which “represents” other shades for us). I say “perhaps” here because I am uncertain about what we can actually succeed in bringing to mind in such a case: the grasping and apprehending rationalists speak of is really foreign to me. I am, however, certain--very certain--that no phenomenologically unified, generic color property is ever before anyone’s mind at a particular time. No single greenness is ever presented that can be immediately apprehended as incompatible with yellowness--and no greenness that can be immediately apprehended as incompatible with redness. The fact that reds and greens do not mix, or cannot be compounded, does not support the thesis of moderate rationalism.¹⁶

I noted earlier that BonJour might be willing to rely on examples involving more determinate properties than red, yellow, or green. In a passage I referred to he mentioned the darkish blue colors of two books on his desk; both are darkish blue, he says, but not quite the same shade. Although he has no specific names for these blues, he nevertheless claims to know directly that nothing could have both of them

all over at the same time.¹⁷ One thing missing in his discussion of the example is any mention of what individuates the shades, what determines where in the spectrum one shade ends and another begins. This is an important matter, because without some discussion of it a dispute of the sort Mary and Harry had about their problem color could arise here as well.

Suppose Hilda recognizes a shade C which encompasses both shades of BonJour's darkish blue. Why couldn't some instance of C (one in which the shades overlap) be truly described as having both shades of blue at the same time? This possibility couldn't be dismissed merely on the ground that such an instance of C would be immediately distinguishable from either of the shades, because any candidate for something simultaneously exemplifying two colors is bound to look different from something exemplifying just one of them.¹⁸ If BonJour declares that according to the classification scheme he is tacitly adopting, nothing could count as having one of these shades if it also has another, he would eliminate the possibilities I have mentioned, but he would do so only by means of a stipulation concerning his mode of classification; he would not be appealing merely to the appearance of the shades themselves. Nature is what it is, but it does not delimit extensions or draw distinctions for us. This is something we do, however unconsciously, tacitly, or automatically.

The point I have just made also applies to colors far apart in the spectrum--for determinate colors as disparate as vermilion and Kelly green. It may be physically impossible for a thing to have a mixture of these colors, but there is nothing inconceivable--nothing conceptually impossible--about it. Of course, if one adopts the convention that an object can properly possess a determinate color just when it is indistinguishable in color from a paradigmatic sample of that color, one can immediately infer that if a thing is Kelly green, it cannot simultaneously be vermilion--for things with these determinate colors are patently color-distinguishable. Yet this convention is not something we have to accept; in fact, it is not really acceptable without significant qualification, since (a) we don't possess paradigmatic examples of every discernible color quality, (b) it is unlikely that vermilion and Kelly green are absolutely determinate (they no doubt have instances that are not indistinguishable in color), and (c) we sometimes cannot discern color differences in surfaces that we can know to be nonhomogeneous by inference from further observations (e.g., a foot-wide color-sample can appear homogeneous but differences may be discerned if it is folded so that its edges are brought together). In practice, we do speak of differences in determinate color when we

can discern differences, and we generally deny differences when we cannot discern them. But our practice here is not wholly precise or determinate: If one person sees a difference where no one else does, there is no particular conclusion that has to be drawn. Our practice could be made more determinate, of course; but a change in practice would be a change in tacit conventions, and it would support the conclusion that the color-incompatibilities urged by BonJour do not have the nonconventional status that he supposes.

Immediacy, Concepts, and Propositions

Earlier in the paper I said that BonJour's remarks about seeing that nothing red can also be green disclose two different models of what is going on, a perceptual model and a discursive model. What I have been saying obviously undermines the suitability of the perceptual model. If there is any truth in the idea that one can in some sense "see" that nothing red can also be green, the relevant seeing must be based on a tacit inference. I have just shown how a convention pertaining to the identification of determinate colors--their sameness and difference--can support a conclusion about such colors, but I have done nothing to account for the apparent immediacy of the truths rationalists claim to see or intuit. A return to the example of Mary and Harry may be instructive here.

If Harry accepted Mary's solution to the problem of how to classify their problem color, he would agree that *as they are now resolved to use them*, "yellow" and "green" are mutually exclusive predicates. Anything classifiable by them as yellow will not count as green as they now understand "green," and anything classifiable by them as green will not count as yellow as they now understand "yellow." (As they now understand "green," it will not apply to the color of the bush they saw; and as they now understand "yellow," it will apply to that color.) Having adopted these restrictions on their usage of "green" and "yellow," they will not have to stop and think about what to say if they are asked to identify the color of the bush they saw. They will say right off "That's yellow," and if they are asked "Isn't it green?" they will be disposed to reply, "No; we regard it as yellow." They will not have to do any explicit inferring, at least so long as their decision about how to describe the bush is fresh in their minds. Normal human beings have the ability to do this. Mary and Harry's decision will support the truth of their assertion "As we classify colors, the bush is not green because it's yellow," but they will make the assertion spontaneously: they will not have to reflect on their basis for it. An analogous assertion of "If green, then not red" would be

equally natural and equally justified for anyone who uses “green” and “red” as mutually exclusive predicates and, as a consequence of consciously doing so, conceives of *green* and *red* as incompatible colors.

A moderate rationalist hearing what I have just said might ask, “But aren’t *the concept* of green and *the concept* of red mutually exclusive independently of the way anyone uses words?” I would respond by denying something that this person is assuming--namely, that there is just one concept of green and one concept of red. Even though (as it happens) green does not mix with red, both mix with all sorts of other colors and do so in various degrees of saturation and brightness. There is no reason to suppose that every human being groups reds and greens together in just the same way. If, like the fictional Mary, some actual people decide to classify chartreuse objects as greenish-yellow and therefore yellow, they will use “yellow” in a special way and conceive of yellow as encompassing chartreuse. And if, unlike those people, I conceive of chartreuse as a shade of green--if it is in the extension of “green” as I use that word--then my concept of green will not be the same as theirs. It will coincide with their concept in many if not most cases, but it will not be the same concept.

It is true that if I say “That bush is green,” people who speak English but do not classify greens as I do will report what I say to someone else by saying “Aune said that that bush is green.” But this fact, which Tyler Burge has made so much of in his well-known article,¹⁹ has no significance for the issue at hand here. We customarily report what others say in indirect discourse by giving the words they use (or systematically revised versions of them) even when we know that we and they attach significantly different meanings to those words.²⁰ We are not transmitting a discrete “content”; we are indirectly quoting a speaker’s words, using them in what traditional grammarians called “*oratio obliqua*”.²¹ A person who sincerely believes that Jones *meant* that he didn’t do Y when he said “I didn’t do X” would nevertheless have to agree, if he believed that doing Y is not identical with doing X, that Jones *said* that he didn’t do X. What Jones meant in saying this is obviously important if there is any question whether in so speaking he was telling the truth. When we are concerned with lying and truth-telling, we are concerned with what people meant to convey, not merely with the words they happened to use.

The reader will observe that in describing the inference “If red, then not green” I was referring to sentences rather than propositions. Moderate rationalists can be expected to object to this, for they generally insist that it is propositions, not utterances or sentences, that we can see to be true.²² By “proposi-

tions” they normally mean structures containing properties, for it is by grasping these properties and seeing how they are related to each other that we are supposed to see that the propositions are true.²³ Since I do not believe that vague generic predicates like “red” and “green” express discrete properties that can be grasped by a mental act, I cannot accept propositions so understood. If propositions are understood differently,²⁴ I would not have this difficulty with them, but I would still want to insist that sentences used by speakers (or writers) in particular contexts are primary bearers of truth and primary vehicles of logical (and other) implications.²⁵ It is for such sentences that we have worked out truth theories, and it is sentences that contain the terms and predicates whose referents are established by the naming and classifying practices of appropriate speakers.

It is beyond the scope of an essay like this one to prove that the primary bearers of truth are not the sort of thing rationalists seem to have in mind. (Actually, I doubt that such a thing could be *proved* anyway.) I can, however, support my opinion by a number of considerations that I regard as compelling. I shall introduce some of these considerations now, and some in later sections of this paper.

The first consideration I want to mention is that, if adjectives such as “green” do express or otherwise represent properties, the identity of those properties is fixed by the linguistic behavior of those who use the predicate; the property, independently identifiable, does not give a meaning to the predicate. Again recall Harry and Mary. If they decided to include chartreuse in the extension of “green,” the property expressed by their “green” (supposing that it does express a property) would be different from the property it would express if they had met their problem in a different way. In excluding or including certain shades from a predicate’s extension, we impose a more or less determinate unity on something not inherently possessing it. If the predicate can plausibly be said to represent (or express) a single property, we are, in a sense, creating that property by our classificatory practice: we are singling out something that is not singled out by nature. The fact that a thing can be discerned in nature does not show that it was not (or does not deserve to be counted as) a human creation. The marble constituting the surface of Michelangelo’s *David* was present in the marble slab from which the statue was carved, but Michelangelo nevertheless created that surface by his amazing artistry. Singling out an extension (or partial extension) for “green” is no doubt a more natural, more instinctive, and less creative task than creating Michelangelo’s *David*, but it is nevertheless something only rational beings do: it is not done by nature.

The next consideration identifies a serious weakness in a key argument used to support the rationalist's claim that certain truths can be grasped independently of any language. Bonjour introduced his example of the two shades of blue in support of this very claim. Having identified two shades by reference to two books before him, he appeared to argue that since his color vocabulary (which he describes as "rather meager") contains no names for these specific shades nor any other way of indicating them linguistically, a proposition about them that he can know *a priori* but cannot adequately express in language is that nothing could have both of these colors all over at the same time.²⁶ The trouble with this argument is that Bonjour actually expresses the crucial proposition in language and does so in way that is as adequate for himself as "This pen is mostly white" is adequate for me now. His reader doesn't know the referent of his "these colors" any more my reader knows the referent my "this pen," but each of us knows what the referents of his own words are and each of us understands the sentences in which he has included those words.

In another passage Bonjour quotes with approval A. C. Ewing's claim that a person who is capable of forming visual images might well see the truth of propositions such as the one concerning green and red without having to put them into words.²⁷ To accomplish such a feat the person would no doubt have to have the concept of incompatibility or exclusion (as Bonjour implies in the passage I quoted on page 3 of this paper), and reflection shows that he or she would have to have further concepts as well: concepts of time, space, universality, thinghood, predication, and modality. In other words, to grasp the truth in question, one would have to be capable of thinking such thoughts as "No thing could be both red and green all over (or in the same place) at the same time." I suppose it is conceivable (at least if Wittgenstein was wrong about private languages) that someone could think such a thought without having a conventional language such as English or French, but I am wholly confident that one could not entertain all the propositions we are supposed to be capable of entertaining if one did not have a conceptual system that segments reality in the way that conventional languages do. As far as mental representation is concerned, having such a conceptual system would be tantamount to having a language. Its principal difference would concern its suitability for recording data (making records) and communicating with others.

Other Rationalist Examples

The criticisms I have made thus far are applicable to just a few examples involving color; they are favorite examples for moderate rationalists, but the acceptability of their view does not depend upon them.

Still, they are representative examples.²⁸ To complete my case against moderate rationalism, I shall therefore consider two further sorts of representative examples. One sort is represented by the transitivity of *taller than*; the other is represented by a logical truth (the principle of non-contradiction) and a logical rule of inference (*modus ponens*).

Some of the claims I have made about the red/green examples are also applicable to the one concerning *taller than*. For one thing, the relevant truth or “proposition expressed” is not as language-independent as rationalists suppose. To grasp the relevant truth one must have, if not a language, a conceptual system tantamount to a language, for the relevant necessary truth, appropriately spelled out, is very complex and requires many distinctions. It is not sufficient to grasp *taller than*, *is true of*, and *transitivity*, and then apprehend their interrelations; one must be able to entertain the thought, “For any a, b, c, and t, if a is taller than b at t and b is taller than c at t, then a is taller than c at t,” for things can vary in their tallness over time. But to comprehend what is thought here one needs concepts of quantification, individuality, predication, conjunction (between terms and formulas), conditionality, and, since the principle is said to be apprehended as necessarily true, both truth and necessity.

Another claim applicable to the *taller than* example is that the relevant predicate is actually vague. There is no doubt that Wilt Chamberlain is taller than Yogi Berra and that a dwarf is not taller than a giant, but there are many pairs of objects for which the question “Is A taller than B?” has no more definite an answer than “Is Tom bald?” Can a frog be taller than a tadpole or a wristwatch taller than a ring? Can a mountain be taller than a hill? Frogs, tadpoles, wristwatches, and rings have vertical dimensions, but they are not described as tall or short, and it is not clear that one can be taller than another. As for mountains, they can be tall but not short, and hills can be high or low. Can mountains and hills be compared for tallness? There is no definite answer to this. One can measure the heights of a mountain and a hill and declare that the one with the greatest height is taller than the other, but this way of speaking is not standard, and not clearly right or clearly wrong. The permissible arguments in the schema “x is taller than y” are not sharply demarcated. Yet if *taller than* were a discrete, inspectable property, it should either be possessed by an ordered pair of objects or not possessed by it. We should not have any undetermined cases.

The *red/green* and *taller than* examples have been seized upon by rationalists because they seem to be necessarily true and involve indefinable concepts. If “red,” “green,” and “taller than” cannot be explic-

itly defined, the specimen statements in which they appear cannot, it is thought, be “true by definition” and thus analytic; the empiricists must therefore be wrong about the basis for their truth. Yet our inability to define certain predicates, or find necessary and sufficient conditions for their application to suitable objects, need not be taken as evidence that such predicates connote indefinable properties, or any properties at all. Their application to objects may be justified by reference to properties--as when one justifies the application of “bald” to a man because he may have the property of being utterly hairless. But the property that does the justifying need not, as here, be the property supposedly expressed by the predicate. No such property may exist. We decide to apply “bald” to bare scalps, to withhold it from hirsute ones, and to apply it to intermediate cases only with modifiers such as “slightly,” “nearly,” “almost,” and “kind of.” The sentence “If a man has no hair growing on his scalp, he is bald” may be necessarily true at a certain time and in a certain context, but its truth does not depend on a property of baldness that an attentive mind can grasp and compare with other properties. As I see it, many other necessary sentences are like this.

A number of interesting and important issues can be explored in connection with the examples I have been discussing, but I want to move on to some examples involving logical truths, for they raise problems of a special kind. I explained that Bonjour’s defense of moderate rationalism involved ruling out alternative theories of *a priori* knowledge, moderate empiricism being the standard alternative. Since I share Bonjour’s view that we have some genuine *priori* knowledge, I wish to defend a version of this standard alternative. Bonjour’s criticism of this moderate empiricism was that, in addition to being unable to accommodate the examples concerning color and relative height, it could at best reduce some *a priori* truths to others: it could not accommodate (that is, justify) basic logical truths and basic rules of inference.²⁹ As I see it, Bonjour and other moderate rationalists cannot handle these examples any better than they can handle the examples concerning color and height. They cannot do so, moreover, for fundamentally the same reasons.

Consider the principle, or “law,” of non-contradiction. Formulated in the usual way, “ $\sim(\Phi \ \& \ \sim\Phi)$,” it seems to be a very simple principle, a suitable object of intuitive insight, but the formulation is very misleading.³⁰ The ingredient letter “ Φ ” is schematic; it stands in place of infinitely many formulas of infinitely varying complexity--and this infinite variety is a very inappropriate object of mental vision. When we think about possible members of this infinity, some can be brought to mind that falsify the law unless quali-

fications are made. Suppose we consider two sentences, “A,” and “B,” the first inscribed in a circle and the second in a rectangle. Suppose “A” is “The sentence in the rectangle is true” and “B” is “The sentence in the circle is not true.” By obvious principles of logic and semantics we can easily derive the contradiction, “A & ~A”. I am not saying that we cannot find reasons for disqualifying this case as a genuine counterinstance to the law of non-contradiction; I am saying that the task of finding such reasons shows the law to be, in its extreme generality, a very unsuitable object for a mental act of seeing.³¹ Too many formulas are involved; too many considerations arise; too much cannot be determined prior to the examination of actual cases.³²

For an example of an elementary rule of inference, consider *modus ponens*. Again, the principle seems intuitively obvious; but if it is stated in English, we can find possible counterexamples. Here is one, discovered by Vann McGee.³³ To appreciate it, recall that the 1980 presidential election was won by Ronald Reagan, a Republican, and that Jimmy Carter, a Democrat, was second and Anderson, a Republican running as an Independent, was third. The example concerns this election:

1. If a Republican wins, then if Reagan does not win, Anderson will win.
2. A Republican wins (=does win).
3. Therefore, if Reagan does not win, Anderson will win.

The first and second premises seem obviously true: Reagan won, and he and Anderson were the only Republicans running in the election. But the conclusion seems false. The real race was between Reagan and Carter; Anderson was far behind. So at the time of the election it would be false to say “If Reagan does not win, Anderson will win.”

There is some controversy about whether the argument here is a genuine counterexample to *modus ponens*.³⁴ Its author, Vann Magee, thinks it is a genuine counterexample. At least one writer, Christopher Gauker, thinks it is not a counterexample to *modus ponens* but thinks it can be converted into a counterexample to *modus tollens* by switching lines (2) and (3) and negating them both.³⁵ I think (for reasons I shall mention a little later) that, without some clarification of the English in which the argument is cast, it is impossible to say decisively whether it is or is not an acceptable counterexample. I shall have more to say about the argument later; here I shall merely note that the disagreement about it supports my contention that the validity of *modus ponens* or *modus tollens* is not something that can plausibly be grasped intuitively by

an act of rational insight. As before, too many formulas are involved; too many considerations arise; too much cannot be determined prior to the examination of actual cases.

If the examples that moderate rationalists take to support their view--and I have now considered four of them-- do not actually support it, it is reasonable to conclude that their view has little direct support. On the other hand, if the critical part of Bonjour's argument is successful, moderate rationalists might contend that their view receives some indirect support from the inadequacy of empiricist views, which are the only available alternatives to theirs. To counter this contention I shall therefore offer some criticism of Bonjour's criticism of moderate empiricism, the version implying that we do have some *a priori* knowledge. As I see it, his criticism is entirely unsuccessful.

An Empiricist Alternative to Moderate Rationalism

According to moderate empiricists, the truths they acknowledge to be knowable *a priori* are analytic.³⁶ As I noted, Bonjour's basic criticism of them is that at best³⁷ they can do no more than reduce some *a priori* truths to others: they may be able to show that *Bachelors are unmarried* is reducible to a logical truth by substitution of synonyms for synonyms, but they cannot say anything epistemologically helpful about the *a priori* justification (or truth) of logical truths or logical rules of inference.³⁸ This basic criticism is clearly applicable to the conceptions of analyticity offered by Kant and Frege, which fall under the first conception that Quine discussed in his famous "Two Dogmas"³⁹ article; but it is not applicable to the second and third conceptions that Quine discussed there. The second conception--that analytic truths are "true by virtue of semantical rules"--was not expounded in a helpful way by Quine (who rejected it), but it was fundamental to the thinking of his friend Rudolf Carnap, whom Quine considered "the dominant figure in philosophy from the 1930's onward"⁴⁰ and whose writings Bonjour did not consider. Although (judging by his criticism) Quine seems to have missed the point of what Carnap was sometimes trying to say (I shall try to justify this remark a little later), I think Carnap's writings on analyticity and his analytic practice make his conception reasonably evident to analytically-minded readers sympathetic with the aims of empiricism.⁴¹ Since the treatment I have suggested for the yellow-green and red-green examples accords with Carnap's conception, as the approach I want to take to the principle of noncontradiction, the rule of *modus ponens*, and the *taller than* example also does, I shall elaborate upon it here. I shall ignore the third con-

ception (analytic truths are statements that are “confirmed come what may”), for I do not regard it as acceptable or even initially plausible.

A good way to bring out the plausibility of Carnap’s conception is to consider the logical examples I discussed earlier. I shall begin with the supposed counterexample to *modus ponens*. When I originally presented the example, I expressed the opinion that it is impossible to say decisively whether it is or is not an acceptable counterexample without some clarification of the English in which it is expressed. The logical word *if* featured in it is clear in some respects, but it is not clear in others, for arguments containing it can be expressed in nonequivalent symbols. Suppose we read the argument as having the following logical form:

1. A Republican wins \supset [\sim (Reagan wins) \supset Anderson wins].
2. A Republican wins.
3. Therefore, \sim (Reagan wins) \supset Anderson wins.

Read this way, the argument is clearly not a counterinstance, for the conclusion is plainly true: it is logically equivalent to “Reagan wins \vee Anderson wins,” which is guaranteed to be true if it has a true conjunct--and it does so in this case.

There are, of course, other ways of construing the argument. When I presented it as an ostensible counterexample, I suggested that the conclusion is false because the real race was between Reagan and Carter, Anderson being so far behind as to be effectively out of it. If the conclusion is read with this firmly in mind, it will appear to have a subjunctive force not captured by the horseshoe symbol. Suppose, therefore, that we interpret the *if*’s in the argument as representing the counterfactual conditionality expressed by David Lewis’s symbol “ \rightarrow .”⁴² On this interpretation the argument takes the following form:

- 1*. A Republican wins \rightarrow [\sim (Reagan wins) \rightarrow Anderson wins]
- 2*. A Republican wins.
- 3*. Therefore, \sim (Reagan wins) \rightarrow Anderson wins.

Understood this way, the conclusion is no doubt false, for in the closest world in which Reagan does not win, Carter presumably wins instead of Anderson. Yet the first premise is now false, and it must be true if the argument is to provide a counterexample. The closest accessible world in which a Republican wins in 1980 is the actual world, and in this world it is not true that if Reagan were not to win that election, Ander-

son would. Thus, when the vernacular *if* is replaced by the technical symbol “ \rightarrow ”, the resulting argument also fails to provide an acceptable counterinstance to *modus ponens*.

Not all occurrences of *if* need be replaced by the same technical symbol, of course. Two further arguments could be obtained if one of the following formulas were put in place of 1:

4*. A Republican wins \rightarrow [\sim (Reagan wins) \supset Anderson wins]

5*. A Republican wins \supset [\sim (Reagan wins) \rightarrow Anderson wins]

If 1* were replaced by 4*, the result would not be an instance of *modus ponens*, however; for the consequent of 4* differs from 3*. If 1* were replaced by 5*, we would have an instance of *modus ponens*, but the first premise would not then be true. 5* is logically equivalent to the disjunction of “ \sim (A republican wins)” and 3*, both of which are false. Thus, on these further readings we still do not have an acceptable counterexample.

Other, nonstandard readings of the vernacular *if* are possible, and it is on one such reading that Christopher Gauker defends his counterexample to *modus tollens*. The multiplicity of possible readings of the vernacular argument raises an important question: “Just what is *modus ponens*?” A precise answer seems impossible here. We can say that *modus ponens* is an argument form in which a conclusion *q* is inferred from a premise *p* and conditional premise having *p* as antecedent and *q* as consequent; but since formulas of significantly different logical powers can be described as conditionals, argument forms of significantly different kinds can count as instances of *modus ponens*, some lacking counterinstances and some, for all I know, having them. The vernacular *if* is not so precise in meaning that only a single interpretation is possible for it even in a given context. If we want to single out a definite class of argument forms in speaking about *modus ponens*, we shall have to restrict our reference to the argument forms that can be constructed from the vocabulary of some formal system or group of systems. A “regimented” part of English may count as such a system, the precision (or logical determinacy) of its formulas depending on the way it is regimented.

It should be clear to the reader that the arguments I could confidently declare to be, or not be, counterexamples to *modus ponens* contain logical symbols with precise interpretations. The horseshoe symbol is not a common term whose meaning is determined by the linguistic behavior of ordinary speakers; it is a technical symbol whose logical properties are fixed by special conventions. And it is not by reference to

the dictates of “total science” that we know “ $p \supset q$ ” is logically equivalent to “ $\sim p \vee q$ ” or “ $\sim(p \ \& \ \sim q)$ ”; it is defined as one or the other in standard systems of classical logic. These definitions permit an exact assessment of formulas whose implications are sufficiently parallel to those of certain vernacular statements to be considered the latter’s symbolic transcriptions, but the vernacular statements are far less determinate in what they assert. Consider “Either something is red or everything red is green.” A natural assessment of this statement is that it is a contingent truth, supported by the fact that red things obviously exist. But if it is interpreted as adequately symbolized by the formula “ $\exists xRx \vee \forall x(Rx \supset Gx)$,” it is easily seen to be logically true, since “ $\forall x(Rx \supset Gx)$ ” is true if no x is R .

The rules by means of which we determine the truth or falsity of formulas are, in Carnap’s terminology, “semantical” rules. To apply them in a strict way, we have to know what formulas are acceptable substituends for the schematic letters in assertions such as the principle of noncontradiction. Just as we freely stipulate what our technical symbols are to mean, so we freely stipulate what formulas our systems will contain. Claiming to know whether “ $\sim(\Phi \ \& \ \sim\Phi)$ ” is true for all formulas that can replace “ Φ ” without knowing what those formulas may be is risky and rash;⁴³ but we can be certain that it is true in a logical system L if we stipulate that nothing will count as an acceptable substituend for “ Φ ” in that system if it falsifies the schema. Our semantical rules for “ \sim ” and “ $\&$ ” show us how, given the meaning we attach to those symbols, we should evaluate formulas containing them (“ $\sim\Phi$ ” is true just when “ Φ ” is not true, and “ $\Phi \ \& \ \Psi$ ” is true just when both “ Φ ” and “ Ψ ” are true). If, by these rules, a substituend for “ Φ ” in “ $\sim(\Phi \ \& \ \sim\Phi)$ ” leads to falsity, that substituend is to be regarded as unacceptable or improper. The truth of “ $\sim(\Phi \ \& \ \sim\Phi)$ ” is then ensured by the semantical rules governing the symbols contained within it.

Carnap, in “Quine on Analyticity,” a short paper written in 1952 and not published until after his death, said this: “It must be emphasized that the concept of analyticity has an exact definition only in the case of a language system, namely a system of semantical rules, not in the case of an ordinary language, because in the latter the words have no clearly defined meaning.”⁴⁴ In support of this claim he considered “Everything green is extended,” which rationalists commonly offer as an example of a sentence expressing a synthetic *a priori* truth. Quine had said that he thought his own hesitation to classify this sentence as analytic betrayed an incomplete understanding not of “green” or “extended” but of “analytic.”⁴⁵ Carnap

expressed a very different view, saying that it seemed “completely clear” to him that the difficulty lies in the unclarity of “green”—specifically, in an indecision whether to apply the word to a single space-time point. “Since one scarcely ever speaks of space-time points in everyday life,” he said, “this unclarity about the meaning (or intended application) of ‘green’ plays as small a role as the unclarity about whether the term ‘mouse’ should also be applied to animals which, apart from their greenness, are completely similar to the mice we know, but are as large as cats.” This unclarity is unimportant for the practical purposes of everyday life, but it is vitally important for the philosophical question about the analyticity of “Everything green is extended.”

These words by Carnap express with minor qualifications the attitude I have adopted toward the examples that moderate rationalists commonly offer in support of their views. The words used in formulating the examples belong to everyday language, and as they are normally used, they lack a clearly defined meaning. Typically, they are vague (or semantically determinate) in important respects; their extensions as ordinarily used are fuzzy. Intelligent, clear-headed people can disagree about their application to possible cases, and it is rarely possible for people to agree on precise definitions for them. This is true even of hackneyed examples such as “bachelor.” As Gilbert Harman observed, people are apt to disagree about whether this word is applicable to the Pope, who is male and not married in any ordinary sense, or to a man who has lived with a woman for several years without getting married.⁴⁶

Since most ordinary words are vague, it is implausible to claim that, as routinely used by different speakers in different contexts, they express discrete determinate concepts or properties that philosophers can simply “analyze.” The implausibility of this claim is, however, compatible with the assertion, which I take to be true, that a generally vague word may be used to express something discrete and determinate in certain contexts. When philosophers offer what Carnap called “explications” of philosophically pertinent words, they implicitly specify contexts in which familiar words have clarified or adjusted meanings. The “definitions” some philosophers give are in fact best viewed as proposed explications. We should not suppose, as Matthias Steup does,⁴⁷ that the “knows” in “s knows that p” denotes a unique property that philosophers have been trying to analyze, with conspicuous unsuccess, for three decades; rather, we should suppose that they have been offering what are in effect proposals for how “s knows that p” (and translations thereof) should be understood, or read, in philosophical contexts. Such proposals, when firmly rejected in

the light of “counterinstances,” are not shown to be false or mistaken, though errors or mistakes may have led to them. They are shown to be more or less ill-advised, for they do not accord with the beliefs, assumptions, or expectations that those who speak of knowledge in philosophical contexts generally share.⁴⁸

When a proposed explication of a word, or sense of a word,⁴⁹ is adopted in a certain context, a reading on which a sentence containing that word is analytically true is thereby specified. “A brother is a male sibling” might therefore be taken to be analytically true in the context of a philosophical discussion of analyticity. An explication, as I am using the term here, includes what might be called a partial explication or an adjustment in meaning. If Mary and Harry agree that, as they are going to use “yellow,” their problem color will count as a shade of yellow, they are adjusting the meaning of “yellow” in the context of *their talk about colors*. As a result of this adjustment and their use of the term “the problem color,” the conditional sentence “Anything having the problem color is yellow” is analytic in the context of their talk about color. Similarly, when Carnap in 1936 wished to introduce predicates for dispositions into the context of a technical language having the horseshoe as its sole symbol for conditionality, he was able to “determine” their meaning only incompletely, only “for certain cases,” by means of reduction sentences, a bilateral version of which had the form of “ $Q_1 \supset (Q_3 \equiv Q_2)$,” where “ Q_1 ” and “ Q_2 ” represent old predicates of the language and the new predicate, “ Q_3 ,” is “determined” in meaning only for cases in which the test condition “ Q_1 ” is satisfied.⁵⁰ Applied to the predicate “water soluble,” the reduction sentence would specify a necessary and sufficient condition for its applicability to cases in which something is immersed in water. Its applicability to other cases would remain undetermined in basically the way that the application of “is bald” is undetermined for cases in which a person showing a lot of scalp still has a significant amount of hair.

In his later writings on analyticity Carnap used the term “meaning postulate” for a formula expressing a partial explication of a word or symbol. In his criticism of Carnap in “Two Dogmas...,” Quine seemed to object to this usage, saying:

...given simply a notation, mathematical or otherwise, and indeed as thoroughly understood a notation as you please in point of the translations or truth conditions of its statements, who can say which of its statements rank as postulates? Obviously the question is meaningless--as meaningless as asking which points in Ohio are starting points. Any finite...selection of statements (preferably true ones, perhaps) is as much a set of postulates as any other (p 35).

But Carnap was under no illusions about this matter. Meaning postulates are postulates, he thought, only relative to a purpose, as Quine almost proceeds to say. For the purpose of stipulating how a word or symbol is to be understood (fully or in specified respects) in a certain context, Carnap relates the word or symbol in a particular way to other words or symbols whose meaning is taken to be adequately understood already (for the purpose in question). What is important for him is not the label “meaning postulate” but the stipulation that is recorded by the use of that label.

Carnap explained the point of a possible stipulation in his paper called “Meaning Postulates.”⁵¹ Suppose a person constructing a certain system wishes to use the symbolic predicates “BI” and “R” in a way corresponding to (but not necessarily the same as) the way “black” and “raven” are used in everyday life. Speaking of such a person, Carnap says:

While the meaning of ‘black’ is fairly clear, that of ‘raven’ is rather vague in the everyday language. There is no point for him to make an elaborate study, based either on introspection or on statistical investigation of common usage, in order to find out whether ‘raven’ always or mostly entails ‘black.’ It is rather his task to make up his mind whether he wishes the predicates ‘R’ and ‘BI’ of his system to be used in such a way that the first logically entails the second. If so, he has to add the postulate (P₂) ‘(x)(Rx ⊃ BI x)’ to the system, otherwise not
(p. 225).

If the postulate P₂ is added to the system, the person constructing it has thereby stipulated how, in the context of the system, the predicate “R” is understood in relation to a symbolic predicate corresponding to “black.” If “R” is applicable to a thing x, “BI” must be applicable to it as well.

Gilbert Harman in “Analyticity Regained?” (*op. cit.*, p. 399) says, “...stipulative definitions are assumptions. To give a definition is to say ‘Let’s assume for the time being that the following equivalence holds’.” This is wrong. Assumptions can be false; stipulative definitions cannot. If I decide to use “raven” in accordance with the stipulation (holding for a certain context) that nothing non-black will count as a raven, I will not be proved wrong if something that might be called a raven in the ordinary sense--a bird indiscernible from a raven except for being white--should be observed. It would simply not be a raven in my stipulated sense. Using my special terminology, I might call it a “schraven” and say that ravens and schravens in my sense of the words are pretty clearly subspecies of a distinct kind that might be called

“dravens.” Seeing such a bird might move me to bring my special terminology more into line with common usage and to use “raven” as people ordinarily do. But I would not have made an error in using “raven” as I formerly did.

When Carnap said (in a passage I quoted above) that the concept of analyticity “has an exact definition only in the case of a language system, namely a system of semantical rules,” he was thinking of an exact definition as the kind of definition that would correspond to Tarski’s definition of “is true” for a formal language. In “Meaning Postulates” Carnap showed how he thought such a definition could be constructed for a formula tantamount to “ Φ is analytically true in L.” Although I am in fundamental agreement with Carnap’s approach to analyticity, my own approach differs from his in what I think are relatively minor respects. For one thing, I am not committed to giving the kind of definition that he gave. For me, a much more informal procedure is sufficient. If a sentence S is a logical consequence of stipulations (= formulas expressing what is stipulated) or meaning postulates accepted by a person or persons P for a context C, then S is analytically true for P in relation to C. For another thing, I am prepared to speak of analytic truth without explicit reference to stipulations or meaning postulates.

In view of what I have just said, I must answer two important questions if I am to offer a complete exposition of my conception of analyticity: How can I justifiably describe some sentence as analytic without explicitly referring to stipulations or meaning postulates? And how can I justifiably describe logical truths as analytic? The subjects of both questions are complicated and require extensive discussion, but I must limit myself here to some general remarks.

I begin with the first question. Although I am most confident in relating analytic truths to explicit stipulations, certain implications of most words are so firmly embedded in the language (or in particular dialects of the language) to which they belong that any stipulation involving those words that does not preserve those implications would universally be acknowledged by alert and attentive speakers of the language (or the dialects of it) to be odd, puzzling, or paradoxical. If this condition is satisfied by a word or symbol, a sentence of the language clearly and unambiguously expressing the entrenched implications can reasonably be regarded as analytically true for those alert and attentive speakers. This explanation is vague in important respects, and in consequence the predicate “analytically true” (as I am using it) is also vague in corresponding respects. This consequence is not damning because many if not most useful words are vague,

and almost every vague word has some firmly embedded implications. One is given by “A thoroughly hairless person is bald.”

The clearest examples of logically true sentences are formulas containing logical symbols with precise interpretations; my remarks about “the” law of noncontradiction are sufficient to indicate how I think the truth of such statements can be justified or ascertained by reference to semantical rules, which provide the relevant interpretations. The clearest examples of valid rules of logical inference also contain logical symbols with precise interpretations; my remarks about *modus ponens* are sufficient to indicate how I think the validity (the truth-preserving property) of such rules can be justified or ascertained by reference to semantical rules, which provide the relevant interpretations of the symbols they contain. Less clear but nevertheless relatively unquestionable examples of logical truths and logical rules of inference can be identified by the considerations mentioned in the last paragraph. If substitution instances of the schema “If Φ and Ψ , then Φ ” are restricted to formulas lacking vague predicates (see footnote 43) and problematic self-reference (either of the indirect kind I described or of a more direct kind accomplished by demonstratives), they can no doubt be regarded as logically true in spite of the informal language in which they are expressed. Such instances are only relatively unquestionable because surprises are always possible in a natural language. An example is an instance that prompted the qualification I just made about demonstratives in connection with “If Φ and Ψ , then Φ .” The example is, “If this clause is an antecedent and this sentence is written in English, then this clause is an antecedent.” If an occurrence of “this clause” refers to the clause in which it appears, the example arguably falsifies an unqualified “If Φ and Ψ , then Φ .”⁵²

There is a final point to consider, one prompted by a criticism that is commonly made of conventionalist treatments of logic. Bonjour expresses the criticism in his attack on moderate empiricism. He does so as follows:

...once it is agreed that...the set of conventions is finite and the set of *a priori* justifiable statements infinite, there must be logical relations between the conventions and the further statements that determine the *a priori* status of the latter and that must themselves be justified *a priori* if the account is to work. But the *a priori* status of these logical relations...cannot itself be accounted for by those same conventions, on pain of obvious circularity (p. 57).

This argument is highly compressed and therefore difficult to set forth in detail, but it contains a serious flaw that deserves to be pointed out.⁵³

To see the difficulty with the argument, assume that we are using a standard version of the sentential calculus, the rules for which are understood as freely adopted (but carefully-considered)⁵⁴ conventions. If the rules include conditional proof and disjunction introduction, we can by these means alone obtain an infinite set of logical truths. By conditional proof we first derive “ $p \supset p$ ” and then derive a potentially infinite sequence of disjunctive formulas by adding disjuncts to this initial formula, running through the vocabulary of the language in a systematic way. Thus we may derive “ $((p \supset p) \vee q)$ ”, “ $((p \supset p) \vee r)$ ”, “ $((p \supset p) \vee (q \ \& \ r))$ ”, and so on. Since we obtain these formulas by constructing proofs for them in accordance with the rules of the system, we are entitled to accept them as theorems, that is, as derived logical truths. Where is the circularity in this?

It is true that we do not obtain formulas of the form “ Φ is true,” “ Φ is L-true” or “ Φ is true in all models of L.” To prove formulas of this kind, we have to ascend to a metalanguage and use the logic assumed there. Some of the theorems of this higher-order logic will look a lot like the theorems of our object language, but they will not be the same theorems, for their schematic letters (their variables in some systems) will not have the same substituends. So we shall not, strictly speaking, be involved in any circularity if we ascend to the metalanguage.⁵⁵ But this ascent to the metalanguage should not really be necessary, anyway. If we are prepared to accept the theorems of a metalanguage (assuming that we are happy with the logic used there), we should be prepared to accept the theorems of the object-language, for we have constructed rigorous proofs for those theorems. We obtain those theorems by *applying--by using--*the rules of our object language. We do not deduce the theorems from the rules. The difficulty that BonJour speaks of does not arise for us. It does not arise at the object language, and it does not arise at the meta level either.

When BonJour completed his criticism of moderate empiricism, he said he found it “shocking” that a view “as ill-defined, poorly defended, and [*sic*] in this way ultimately incoherent as moderate empiricism should have been held so long, so confidently, and often so complacently by so many philosophers” (p. 61). I cannot deny that many moderate empiricists have had a very poor understanding of the credentials of the epistemological position they were assuming, but I am certain that the same was not true of men like Car-

nap: he at least had a very clear idea of what he was doing. I hope what I have said here increases the attractiveness of moderate empiricism. I think it is far preferable to the new rationalism.⁵⁶

Bruce Aune

University of Massachusetts, Amherst

Endnotes

¹ Laurence Bonjour seems to have introduced the term “moderate rationalism” in his book, *In Defense of Pure Reason* (Cambridge: Cambridge University Press, 1998), *passim*. See also George Bealer, “The Incoherence of Empiricism,” *Proceedings of the Aristotelian Society Supplementary Volume*, 1992., and Alvin Plantinga, *Warrant and Proper Function* (New York: Oxford University Press, 1993), ch. 6.

² See Descartes, *Rules for the Direction of the Mind*, 425; translated in John Cottingham *et al*, *The Philosophical Writings of Descartes*, Vol. 1 (Cambridge, UK: Cambridge U.P., 1985), p. 48.

³ See Bonjour, p. 99.

⁴ See Bonjour, ch. 3.

⁵ Bonjour, p. 98.

⁶ Bonjour, pp. 100-106.

⁷ I put “explains” in quotes because an erroneous account (or theory) cannot actually explain anything. When we ask of incompatible theories “Which provides the best explanation of phenomenon *E*?” we must be understood as asking something more qualified, such as “Which theory, if true, would best explain *E*?”

⁸ Bonjour, p. 101.

⁹ The reasoning would appear to feature a premise whose truth is arguably an analytic consequence of the meaning of “exclusive occupier, with respect to *C*, of a surface or area”—namely, “For all *x, y, z, t*, if *x* and *y* are exclusive occupiers, with respect to *C*, of a surface or area *z* at time *t* & $x \neq y \supset \sim(x \text{ and } y \text{ occupy } z \text{ at } t)$).

¹⁰ See the anthology by Alex Byrne and David R. Hilbert, *Readings on Color, volume 1: The Philosophy of Color* (MIT Press, 1997).

¹¹ This is what is given in *The Random House Dictionary of the English Language, College Edition* (New York: Random House, 1968).

¹² See Bonjour, pp. 57f.

¹³ See David Lewis, “Many, But Almost One,” in Lewis, *Papers in Metaphysics and Epistemology* (Cambridge, UK: Cambridge Univ. Press, 1999), pp. 164-182.

¹⁴ See G. E. Moore, *Principia Ethica* (Cambridge: Cambridge Univ. Press, 1903), p. 7.

¹⁵ For an informal discussion of fuzzy set theory and arguments for the view that the semantics of basic color terms is best represented in the formalism of fuzzy set theory, see Paul Kay and Chad K. McDaniel, “The Linguistic Significance of the Meanings of Basic Color Terms,” in Alex Byrne and David R. Hilbert, *Readings on Color, Volume 2: The Science of Color* (MIT Press, 1997), pp. 399-442. In this article Kay and McDaniel also argue that however different the languages of normal human beings may be, they nevertheless segment color space in a limited and systematic variety of alternative ways.

¹⁶ The fact that red does not mix with green as yellow and blue do is explained by the physics of light and the physiology of human retinas. For information this, see the essay by Kay and McDaniel cited in fn. 13

¹⁷ BonJour, pp. 57f.

¹⁸ BonJour responded to C. L. Hardin’s claim that purple is both red and blue by saying that he could be accused of lying if he told someone that a thing he knows to be red is purple. (See BonJour, p.101n.) Such an accusation would be justifiable, however, only if it is assumed that the speaker knew that his hearer thought of purple as different from red (not a shade of it or a mixture containing it) and made no effort to enlighten the hearer about how he viewed it.

¹⁹ See Tyler Burge, “Individualism and the Mental,” *Midwest Studies in Philosophy*, vol. IV: *Studies in Metaphysics* (Minneapolis: University of Minnesota Press, 1979), pp. 73-122.

²⁰ If Mary told you “I will see you here tomorrow,” you would naturally relate what she said on the following day in a different location by saying “She said that she would see me there today.” Only one of her actual words is preserved in the that-clause.

²¹ See, e.g., H. W. Smyth, *Greek Grammar* (Cambridge, MA: Harvard University Press, 1920), where Smyth says “In an indirect quotation the words or thoughts are given at second hand with certain modifications to indicate that the words or thoughts are reported” (p. 584).

²² One of the ironies of contemporary philosophy is that Roderick Chisholm, the most influential defender of rationalism in the middle years of the twentieth century, ended up denying that propositions and states of affairs exist. See L. E. Hahn, ed., *The Philosophy of Roderick M. Chisholm* (Chicago & LaSalle, IL: Open Court, 1997), p. 358.

²³ BonJour says, “If having a thought whose content is, for example, the claim that nothing can be red and green all over at the same time involves being in a mental state that instantiates a complex universal of which the universal’s redness and greenness are literal constituents [as he believes it does], then at least much of the mystery surrounding my access to those universals and my ability to intuitively apprehend the relation of incompatibility between them is removed” (pp. 184f). I cannot resist observing here that if BonJour is right about the incompatibility of his two shades of blue, two shades of green would be equally incompatible, and since greenness cannot be instantiated without some shade of it being instantiated, the awareness he speaks of would evidently involve the awareness of one shade and the exclusion of others. It is hard to see, then, how the greenness grasped on one occasion could include the greenness of other shades.

²⁴If propositions are conceived of as functions from possible worlds to truth values, or as sets of possible worlds, as David Lewis suggests in his book *Counterfactuals* (Cambridge, MA: Harvard, 1973), p.46, they are not incompatible with the position I defend here and in the remainder of this paper. The “propositional” function expressed on a standard reading of the sentence “Bachelors are unmarried males” would assign the value *T* to a world *w* just when every bachelor in *w* is an unmarried male. Since a necessary condition of being a bachelor in the relevant sense is that a thing be an unmarried male human being, the propositional function pertinent here assigns *T* to every possible world on a purely conceptual basis.

²⁵ A given sentence can have different truth values in different contexts for reasons other than its possession of indexical expressions. As David Lewis observes, different standards of precision may be assumed in different contexts, so that in some contexts “France is hexagonal” will be regarded as true although it is regarded as false in others. He has made a similar claim about “know.” See David Lewis, “Index, Context, and Content,” in Lewis, *Papers in Philosophical Logic* (Cambridge: Cambridge University Press, 1998), pp. 24f. See also fn.48 below.

²⁶ BonJour, p. 58.

²⁷ BonJour, p. 57.

²⁸ An anonymous reviewer of this paper has claimed that the examples I have considered are not representative for traditional rationalists, since the propositions whose truth they claim to intuit involve clear and distinct ideas rather than vague ones. Descartes did hold this view, but (a) it is evidently not held by to-

day's moderate rationalists and (b) many if not most truths that most philosophers would concede to be knowable *a priori* involve patently vague concepts. The positive case I make for our knowledge of *a priori* truth is based on the idea that *most of our concepts are vague*, and it accommodates truths involving precise concepts as well.

²⁹ BonJour claims that the conceptions of analyticity espoused by moderate empiricists are either obfuscating or merely reductive (see his chapter 2). I am assuming here that a moderately empiricist conception is not obfuscating "at best." I make the same assumption below, when I discuss his criticism of moderate empiricism.

³⁰ BonJour (p. 33) formulates it with a quantifier, "for any proposition P, not both P and not P". Unless the quantifier is read substitutionally, however, the formulation does not make sense; for the inner formula will then lack a verb, like "not both Tom and Mary."

³¹ I first developed this point in Bruce Aune, "Determinate Meaning and Analytic Truth," in Guy Debrock and Menno Hulswit, eds., *Living Doubt: Essays Concerning the Epistemology of Charles Sanders Peirce* (Dordrecht: Kluwer, 1994), p. 55-65.

³² See Graham Priest, "The Logic of Paradox," *Journal of Philosophical Logic*, 8 (1979), 219-241, and, for criticism, David Lewis, "Logic for Equivocators," *Noûs*, 16 (1982), 431-441.

³³ Vann McGee, "A Counterexample to Modus Ponens," *Journal of Philosophy*, 82 (1985), 462-471.

³⁴ Since this paper was written, Bernard Katz has also argued that the argument does not provide a successful counterinstance to *modus ponens*; see his "On a Supposed Counterexample to Modus Ponens," *Journal of Philosophy*, XCVI (1999), 404-415. Katz's criticism is not the same as mine, however.

³⁵ See Christopher Gauker, *Thinking Out Loud* (Princeton: Princeton University Press, 1994), p.p. 141f.

³⁶ I put the point this way to accommodate the moderate version of moderate empiricism that I endorse. Instead of claiming--rashly, in my opinion--that anything knowable *a priori* is definitely analytic, I should prefer to make this weaker claim, which is compatible with an agnostic attitude toward mathematical truth. My knowledge of mathematics is not extensive, and I hesitate to make sweeping claims about the status of mathematical assertions. If a particular mathematician cannot provide a better justification for an axiom than a moderate rationalist such as BonJour might provide, I would not want to acknowledge that a theorem he or she infers with the aid of that axiom represents genuine knowledge of some mathematical object. I

would remain skeptical.

³⁷ See n. 23 above.

³⁸ See BonJour, pp. 34f.

³⁹ W.V.O. Quine, "Two Dogmas of Empiricism," in Quine, *From a Logical Point of View*, (Cambridge, MA: Harvard University Press, 1953), pp. 20-46.

⁴⁰ Quine said this in his "Homage to Rudolf Carnap," reprinted in Richard Creath, ed., *Dear Carnap, Dear Van* (Berkeley: Univ. of Cal. Press, 1990), p. 463. His homage was a eulogy, not a critical assessment, but there is no doubt that he regarded Carnap as the dominant figure among logical empiricists.

⁴¹ Even though Quine and Carnap were friends and admired each other's work, they sometimes misunderstood one another and argued at cross-purposes. This is obvious from their interchange on analyticity in *The Philosophy of Rudolf Carnap*, ed. Paul A. Schilpp (LaSalle: Open Court, 1963). I can't explain how this misunderstanding was possible; but it happened. For a perceptive discussion of the Carnap-Quine dispute on analyticity, see Alexander George, "On Washing the Fur Without Wetting It," *Mind*, CIX (2000), 1-24.

⁴² See David Lewis, *Counterfactuals*, *passim*.

⁴³ Another class of formulas that will yield plausible exceptions to the law are those containing vague predicates. If P is a predicate that is applicable to some objects, not applicable to some objects, and (owing to its vagueness) neither applicable nor inapplicable to certain others, formulas containing it can plausibly be assigned one of three values: T, F and I, the latter abbreviating "indeterminate." According to Lukasiewicz's rule for evaluating formulas possessing a third value, if 'Pa' (= 'Tom is bald') has the value I, the formula ' $\sim(Pa \ \& \ \sim Pa)$ ' will also have that value; it will not have the value T. (See William and Martha Kneale, *The Development of Logic* [Oxford: Clarendon Press, 1962], p. 571.) If we assess the validity of the formula ' $\sim(Pa \ \& \ \sim Pa)$ ' by van Fraassen's method of supervaluations, we may assign it the value *super-true*, for it is true for all ways of making the predicate 'P' precise. But this doesn't make the formula true as it stands; its vagueness (its semantic indeterminacy) prevents that. See Bas C. van Fraassen, "Singular Terms, Truth-value Gaps, and Free Logic," *Journal of Philosophy*, LXIII (1966), 481-495.

⁴⁴ This paper is included in *Dear Carnap, Dear Van*, pp. 427-432; the sentences I quote in this paragraph are on pages 427 and 428 respectively.

⁴⁵ See “Two Dogmas...,” p. 32.

⁴⁶ Gilbert Harman, “Analyticity Regained?,” *Nous* (1996), p. 399.

⁴⁷ Matthias Steup says that words denote concepts and that concepts are properties that have instances, implications, and meanings. See his *Contemporary Epistemology* (Upper Saddle River, NJ: Prentice Hall, 1996), pp. 21f, 25, 26.

⁴⁸ I explain the failure of epistemologists to obtain a trouble-free “analysis” of *knowing that p* to three factors: one, a conviction that there is just one relevant concept to analyze; two, a belief that we have a lot of knowledge in ordinary life when our evidence is imperfect; and three, a traditional idea that is often suppressed but keeps popping up, like Austin’s frog at the bottom of the beer mug: the idea that if your basis for thinking that *p* leaves open the possibility that you are wrong, you don’t know that *p*. David Lewis’ treatment of “*s* knows that *p*,” which drops the conviction about a single concept in favor of stricter or less strict readings of the formula depending on the context, is a big step forward in the subject. See Lewis, “Elusive Knowledge,” in *Papers in Metaphysics and Epistemology*, pp. 4187-445.

⁴⁹ In “Two Concepts of Probability,” *Philosophy and Phenomenological Research*, 5 (1945) Carnap introduced his technical notion of an explication and, after doing so, offered explications of what he called “two concepts” of probability. I would describe him as offering explications for two senses of the word “probability.” This essay is reprinted in H. Feigl and W.S. Sellars, eds., *Readings in Philosophical Analysis* (New York: Appleton-Century-Crofts, 1949), pp. 330-348.

⁵⁰ See Rudolf Carnap, “Testability and Meaning,” *Philosophy of Science*, 3 (1936) and 4 (1937); reprinted, with omissions, in Feigl and Sellars, pp. 47-92.

⁵¹ Included as an appendix or “supplement” in Carnap’s *Meaning and Necessity*, 2nd ed. (Chicago: University of Chicago Press, 1956), pp. 22-229.

⁵² This example was suggested by a related example invented by Bas C. van Fraassen, I would not agree that the example is a genuine counterexample to “If Φ and Ψ , then Φ .” but I do agree that the matter is arguable.

⁵³ BonJour attributes the argument to Quine’s “Truth by Convention,” but the argument Quine actually gives in that paper is different from the one BonJour offers. I discuss Quine’s argument in my book, *Knowledge of the External World* (London: Routledge, 1991), pp. 153f.

⁵⁴It is vital to realize that freely adopted conventions need not be--in fact, hardly ever are --adopted capriciously or irrationally. Adopting a convention or rule is a practical act that is normally done for specific purposes. In the case of orthodox logical conventions, one purpose relevant to their adoption is that they do not permit the derivation of contradictions. Experience shows that, unless certain restrictions are made, contradictions will be derivable, and relevant restrictions are therefore made. A purpose relevant to the adoption of “ \supset ” as the symbol for conditionality in *Principia Mathematica* was that, as Russell and Whitehead say, the relation it represents has the essential property required of a conditional: what a true proposition is conditional on is true. (See *Principia Mathematica* to *56 (Cambridge: CUP, 1962), p. 94, where Russell and Whitehead use “implies” where I use “is conditional on.”) It is possible to prove without actual circularity that certain sets of logical conventions (rules, definitions) will have various desirable properties, such as consistency and completeness--and this is another reason for choosing some conventions rather than others. The fact that conventions are normally adopted in view of their supposed or, in many cases, even known consequences does not disqualify them from being freely chosen. Not every logician has standard or orthodox purposes; one exception is Graham Priest; see footnote 32 above and “What Is So Bad About Contradictions?”, *Journal of Philosophy*, XCV, Number 8 (1998), 410-426.

⁵⁵ The epistemic status of our metalinguistic logic is basically the same as that of our lower level, object language logic: in this context we have turtles all the way up.

⁵⁶I wish to thank my colleagues at UMass (particularly Phil Bricker) for helpful criticism or useful suggestions; I also wish to thank Joseph LaPorte, Steve Herman, Jay Garfield, and an anonymous referee for written comments that improved the final product.