Philosophical Perspectives, 21, Philosophy of Mind, 2007

INTENTIONAL INEXISTENCE AND PHENOMENAL INTENTIONALITY

Uriah Kriegel
The University of Arizona, The University of Sydney

Introduction

How come we can represent Bigfoot even though Bigfoot does not exist, given that representing something involves bearing a relation to it and we cannot bear relations to what does not exist? This is the problem of intentional inexistence. This paper develops a two-step solution to this problem, involving (first) an adverbial account of conscious representation, or phenomenal intentionality, and (second) the thesis that all representation derives from conscious representation (all intentionality derives from phenomenal intentionality). The solution is correspondingly two-part: we can consciously represent Bigfoot because consciously representing Bigfoot does not involve bearing a relation to Bigfoot, but rather instantiating a certain non-relational ("adverbial") property of representing Bigfoot-wise; and we can non-consciously represent Bigfoot because non-consciously representing Bigfoot does not involve bearing a relation to Bigfoot, but rather bearing a relation to conscious representations of Bigfoot.

1. The Problem of Intentional Inexistence

The problem of intentional inexistence is one of the perennial problems of philosophy. It is also one of its simplest: it arises out of the trivial fact that we can think of what does not exist.

How is this a problem? Here as elsewhere, we can appreciate the problem by considering an inconsistent triad of initially highly plausible propositions:

1. One can think of non-existents.
2. One cannot bear a relation to non-existents.
3. Thinking of something involves (constitutively) bearing a relation to it.

The first is the trivial claim with which we opened: you can think of Bigfoot, even though Bigfoot does not exist. The second claim is equally straightforward:

one can kick a parrot or sit in a tub, but only if there is a parrot or a tub; nobody has ever kicked a dragon or sat on a magic rug. The third claim appears likewise self-evident: when one thinks of one's mother, one enters a relation to one's mother—the thinking-of relation.

Similar inconsistent triads show up for activities other than thinking of something. There are the activities of seeing something, hearing something, and generally perceiving something;[1] considering something, contemplating something, and generally thinking about something; wanting something, trying to do something, and generally willing something. Beyond our activities, there are also inked and mouthed linguistic expressions, pictures and traffic signs, and myriad other entities for which a similar inconsistent triad can be constructed. What all these have in common is that they involve intentionality, or *representation*: they represent something.[2,3] Thus, seeing a table involves representing the table; wanting an ice cream involves representing the ice cream; depicting the Mona Lisa involves representing her.

The general problem, then, concerns representation. We may formulate it in terms of the following triad:[4]

(a) One can represent non-existents.
(b) One cannot bear a relation to non-existents.
(c) Representing something involves (constitutively) bearing a relation to it.

This is what I will refer to as the *problem of intentional inexistence*. The problem is that none of (a)–(c) can be comfortably rejected.[5]

It might be thought that the problem of intentional inexistence is not at bottom a problem about relations to non-existents. Consider the fact that you can think of a parrot without thinking either of a blue-eyed or of a brown-eyed parrot, but you cannot kick a parrot without kicking either a blue-eyed or a brown-eyed parrot; that you can want a tub without either wanting a small tub or wanting a big tub, but you cannot sit in a tub without either sitting in a small tub or sitting in a big tub; etc. This might suggest to some that representation is special not only in apparently relating us to non-existent entities but also in apparently relating us to indeterminate entities.[6]

However, one plausible diagnosis of the problem presented by representation of indeterminate entities is that it arises because there *are* no indeterminate entities: there are no tubs that are neither small nor big. On this diagnosis, the present problem is derivative of the problem of intentional inexistence as already formulated. The reason it is problematic that we can represent indeterminate entities is that indeterminate entities are non-existent; if there existed indeterminate entities, there would be no inconsistent triad associated with the possibility of representing them.

Nonetheless, a solution to the problem of intentional inexistence should apply equally to the problem presented by the representation of indeterminates. This problem arises because what is represented is sometimes indeterminate

even though what exists is always determinate. Interestingly, there may also be the converse problem, closely associated with Quine's (1960) "inscrutability of reference." This problem arises when what is represented is determinate but what exists is indeterminate. When you think of a rabbit, it seems obvious that a rabbit is determinately what you represent; yet some philosophers may hold that it is indeterminate whether what exists is (i) a rabbit, (ii) a synchronic fusion of rabbit parts, (iii) a diachronic fusion of rabbit stages, (iv) a diachronic fusion of synchronic fusions of rabbit part stages, or perhaps (v) some other metaphysical freak.[7] Again, it is a desideratum of a solution to the problem of intentional inexistence that it solves this problem for those who believe in such ontological indeterminacy. That is, philosophers who believe in ontological indeterminacy should be able to coherently hold that representation is nonetheless determinate, and it is a desideratum of a solution to the problem of intentional inexistence that it show them how to do so.[8]

As with other perennial problems of philosophy, attempts to solve the problem, in this case by justifying the rejection of one among (a)–(c), often strike the reader as perhaps clever or even ingenuous but in the last count not really believable. The purpose of this article is modest: to present a solution to the problem that might strike the reader as believable, that is, to sketch a possible solution and argue that it is at least viable and worth pursuing.

The solution in question is a version of denying (c). I will suggest that representing something does not involve bearing a relation to it. But the reason I will claim this is different for two different kinds of representation, conscious and non-conscious. For *conscious* representations, it is plausible that they represent in virtue of a non-relational phenomenal character they exhibit, and therefore without standing in a relation to *anything*. For *non-conscious* representations, it is plausible that they represent in virtue of bearing a relation to conscious representations, and therefore in virtue of bearing a relation to *things other than what they represent*. Both kinds of representation do not require standing in a relation to what is represented.

The article proceeds as follows. In §2, I will consider the possibility of rejecting (a) or (b), and dismiss it as unviable. In §3, I will offer a first pass at how one might attempt to reject (c). In §4, however, we will see that the approach outlined in §3 faces several problems. In §5, a modified approach will be sketched that may overcome those problems. The modified approach involves two central theses. A full defense of these theses will not be possible here, but a preliminary defense will be sketched in §§6–7.

2. Some Aborted Attempts at Solving the Problem

Our triad is genuinely inconsistent. So the problem can be solved only by rejecting one of (a)–(c). In this section, I argue quickly against the viability of rejecting either (a) or (b); quickly, because what I have to say there is not

particularly new, and the focus of this paper is on the positive proposal to be developed later.

Rejecting (a) would amount to claiming that we cannot represent non-existents. There are several ways this might play out, but they all have in common that, when we seem to ourselves to be thinking of dragons, we are not really. Either we are not thinking at all, or we are thinking of entities that are perhaps associated with dragons but are not themselves dragons.

The notion that when we seem to be thinking of dragons, we are not thinking at all—are not having a thought—might be supported with some strong kind of semantic externalism. Perhaps the idea is that dragon thoughts would have to be directly referential (perhaps because they are natural kind thoughts), and since they are reference-less, they are content-less, and therefore not genuine thoughts at all.

To my mind, it is not ultimately believable that apparent dragon-thinkers are not thinkers at all. The proposal implies that there is a gap between trying to have a thought and having one—that sometimes we try to have a thought, and seem to ourselves to have one, but actually do not. Furthermore, even if the present approach solved the problem for the activity of thinking of something, the problem will resurface for the activity of *seeming to oneself to be thinking of something* (in the relevant sense of "seeming," where an intentional mental state is actually attributed). That is, we can devise a new inconsistent triad: one can seem-to-think of non-existents; one cannot bear relations to non-existent; yet seeming-to-think of something involves (constitutively) a relation to it.[9]

Marginally more plausible is the idea that, when we seem to ourselves to be thinking of dragons, we are indeed thinking of something, just not dragons. The claim is typically that even when we seem to represent non-existents, what we are really representing are existent *abstracta* (in one version of the idea) or existent *mental concreta* (in another version). Thus, when you think of Bigfoot, you are thinking of something that exists, but not the concrete flesh-and-blood monster who lives in the forest. Rather, you are thinking either of an abstract entity that "lives" altogether outside space-time or of a mental object that "lives" only in your mind.

Again, this approach faces a number of obvious problems, not least of which the counter-intuitiveness of the notion that, when we seem to ourselves to be thinking of dragons, we are not thinking of what we seem to ourselves to be thinking of.[10] (In fact, the approach may imply a sort of second-order error theory about intentional states.) There are, in addition, familiar intuitive, ontological, epistemological, and phenomenological difficulties here: (i) intuitively, Bigfoot seems to be a non-mental concretum, though one that does not exist, rather than an existing abstractum or mental concretum; (ii) ontologically, commitment to abstracta and mental concreta is a liability that we should not have to incur merely to account for the facts of representation; (iii) epistemologically, the notion that what we are in direct representational contact with are abstracta or mental concreta throws a veil of appearances over the realm

of external concreta, producing a corrosive skepticism about our knowledge thereof; (iv) phenomenologically, the entities we are aware of when we think of dragons and parrots present themselves to us, from the first-person perspective, as external concreta, not as abstracta or mental concreta. In any case, the same problem arises here that we have seen before: even if when we seem to think of dragons we are really thinking of something else, a new inconsistent triad could be constructed for seeming-to-think of something.

A more recent version of the view under consideration is that, when we seem to ourselves to think of dragons, we are thinking of *possibilia*. Against the background of one ontology of possibilia, this suggestion collapses onto the abstracta view. But against the background of another ontology (Lewis's realism), the suggestion is that we are thinking of non-mental concreta that do not exist in the actual world, though they do exist *simpliciter*. This concrete-possibilia view overcomes most problems associated with the abstracta and mental-concreta views. However, it does not overcome, and probably exacerbates, the ontological problem, inasmuch as it commits us to the existence of non-actual concreta. It is also inconsistent with a general principle about representation, namely, that representation is always grounded in causal contact with the represented. On some views, causal contact is constitutive of representation. The requirement I have in mind is much weaker: that causal contact is a precondition for the possibility of representation (even if it is not constitutive of the representation itself). Since there is no trans-world causation, representation of non-actual concreta is inconsistent with this principle.

Let us move, then, to the option of rejecting (b). To reject (b) would be to maintain that we can bear relations to non-existents. This option is certainly counter-intuitive: just as we do not think that a monadic property can be instantiated in the absence of a particular that instantiates it, so we should not think that a relation could be instantiated in the absence of particulars that instantiate it.

It might be thought that although this is true of most relations, it is not true precisely of intentional relations: perceiving, thinking, remembering, doubting, depicting, and the like. However, the variety of intentional relations should not obscure the fact that they are all varieties of a single (alleged) relation, intentionality, and that presumably they do not require the existence of the relevant relata precisely because they are such. They are all determinates of a single determinable relation,[11] and it is the determinable relation that is claimed to be the only exception to the rule that the instantiation of a relation requires the existence of the relata. At bottom, then, the suggestion under consideration is a sort of *intentionality exceptionalism* that is unappealing in the same way any metaphysical exceptionalism would be: it is odd to maintain that a certain metaphysical truth holds of all entities of type *T* except one.

To relieve worries about intentionality exceptionalism, the proponent of rejecting (b) should adduce a non-intentional relation that may intuitively be thought to be borne to non-existents. So, for example, one might suggest

being-smaller-than: in the sentence "lions are smaller than dragons," lions are claimed, truly no less, to bear the (non-intentional) relation of being-smaller-than to non-existents. But the sentence seems better thought of as ellipsis for something like "if there were dragons, lions would be smaller than them."[12] This is more clearly evident in cases involving only non-existents, such as "unicorns are smaller than dragons," which seems to be ellipsis for "if there were unicorns and dragons, the former would be smaller than the latter." These sentences are counterfactuals, then, so they do not claim that a certain relation is instantiated by non-existents; only that it *would* be instantiated *if* they existed. The problem of intentional inexistence arises, however, because a relation seems to *be* instantiated (in the actual world), since the representing of something does occur (in the actual world), even though the relevant relatum does not exist (in the actual world). Thus, "you are thinking of dragons" is not ellipsis for "if there were dragons, you would be thinking of them." No: the statement is that you *are* thinking of dragons.

I conclude that rejecting either (a) or (b) is not a viable option. This suggests that if a viable solution to the problem of intentional inexistence is to be found, it is probably in the rejection of (c).

It is noteworthy that (c), the claim that representing something involves constitutively bearing a relation to it, is presupposed by the mainstream research program on naturalizing intentionality as it has been pursued since the late seventies. This mainstream is characterized by two principles. The first is that the overarching goal of research in this area ought to be the naturalization of intentionality. The second is that the way to achieve naturalization is to identify the natural (broadly causal) relation that holds between x and y when, and only when, x represents y. A centerpiece of this outlook is thus the idea that there is a "representation relation" that constitutes intentionality, so that naturalizing intentionality would require naturalizing the representation relation.

Since (c) is presupposed by the relevant research program, the viability of that program depends on the plausibility of solving the problem of intentional inexistence by rejecting either (a) or (b). But as the foregoing discussion suggests, the prospects for doing so with any plausibility are unpromising. To that extent, the mainstream research program on naturalizing intentionality may be fatally flawed.[13]

3. Adverbialism about Intentionality

To reject (c) is to hold that thinking of something, and more generally representing something, does not constitutively involve standing in a relation to it. Perhaps representation *often* involves a relation to the represented, but it never does so *constitutively*. That is, it is never the case that a representation represents *in virtue* of bearing a relation to the represented, so that the destruction of the relation would destroy the representation.[14]

There is of course a well-known precedent in modern philosophy to a non-relational treatment of an apparently relational, indeed apparently intentional, phenomenon. This is the adverbial account of sense perception we find in Ducasse (1942) and Chisholm (1957). That account differed from the present proposal in both scope and motivation: its *scope* was limited to sense perception, which is a form of representation, but only one form; its *motivation* was to undermine the sense-datum theory of sense perception, not to solve the problem of intentional inexistence.[15] But the adverbial account of sense perception did involve arguing that some apparently relational intentional phenomenon is really non-relational. Thus the option of rejecting (c) could be cast as a generalization of the adverbial account from the case of sense perception to all representation; that is, as an adverbial account of intentionality in general.

In the adverbial account of sense perception, it is shown that although the surface grammar of "you perceive red" casts perception as a relation between two concrete particulars (you and a red), there is a way to paraphrase such straightforward perceptual ascriptions into sentences with a completely different grammar, suggesting a completely different metaphysic. It can be paraphrased into "you perceive redly," which suggests a state of affairs consisting in a single particular instantiating a property and that property instantiating another property.[16] Just as in "Jim walks slowly," the adverb "slowly" denotes an intrinsic modification of Jim's walking, not a relation between Jim and slowness, so in "Jim perceives redly," the adverb denotes an intrinsic modification of Jim's perceiving, not a relation between Jim and redness.[17]

It is important to realize that offering an adverbial paraphrase does not by itself settle any metaphysical questions. It only shows that there are two intertranslatable ways of speaking of sense perception, one that casts it as a relation between a perceiver and a perceived and one that casts it as a non-relational property of the perceiver. To choose the latter over the former we must provide independent arguments for a non-relational conception of sense perception. The purpose of the paraphrase itself is only to show that despite the naturalness of the relational way of speaking, there is also a consistent non-relational way of doing so. Presumably this is supposed to render the non-relational conception of sense perception *intelligible*. The fact that we can speak of perception in non-relational language is supposed to show that the idea is not nonsensical. It shows that while our ordinary way of speaking seems to commit us to a relational understanding of perception, such commitment is not inevitable, since there are equivalent ways of speaking that do not carry that commitment. This removes a general obstacle to the very viability of a non-relational conception of sense perception. Once automatic commitment to the relational conception is relieved and the non-relational conception is shown to be intelligible and viable, other arguments can be marshaled to show that the latter is actually preferable.[18]

We might pursue the same strategy for intentionality in general. An adverbial paraphrase of intentional ascriptions would be provided to establish the

intelligibility of a non-relational conception of intentionality. Once the intelligibility is established, it would be argued that such a non-relational conception of intentionality is preferable to a relational one on independent grounds. These grounds might simply be that only the non-relational conception is consistent with the fact that we can represent non-existents but cannot bear relations to non-existents (i.e., that a non-relational conception is the only way to solve the problem of intentional inexistence).

The adverbial account of intentionality is fairly straightforward. Although the surface grammar of "you are thinking of Bigfoot" casts the thinking as a relation between you and Bigfoot, the sentence can be paraphrased into "you are thinking Bigfootly," or perhaps more naturally, "you are thinking Bigfoot-wise." The latter casts the thinking as a non-relational property of yours.[19]

A few clarifications. First, the variable *x* in "*x* represents Bigfoot-wise" ranges over the domain of representors. Thus, it may take as value tokens of the word "Bigfoot," paintings of Bigfoot, or activities of thinking of Bigfoot. Confusingly, it can take as value both thinkers (perceivers, desirers, etc.) and thoughts (perceptions, desires, etc.), since we can say both of thinkers and of thoughts that they represent. Plausibly, however, the more fundamental phenomenon here is that of a thought representing. The thinker's representing Bigfoot is plausibly construed as the thinker's *having* a thought and the thought *representing* Bigfoot. For an adverbialist, then, to say that a thinker represents Bigfoot is to say that the thinker has a thought and the thought represents Bigfoot-wise.

Second, it is worth emphasizing that adverbialism is only one apparatus for recasting an apparently relational phenomenon as non-relational. There may be others. For example, we may use the device of hyphenation, paraphrasing "you perceive red" into "you perceive-red" and "you are thinking of Bigfoot" into "you are thinking-of-Bigfoot." Here "perceive-red" and "are thinking-of-Bigfoot" are monadic predicates, ostensibly picking out non-relational properties.[20] Note that the grammar of "you perceive redly" is different from that of "you perceive-red," inasmuch as "perceive redly" is syntactically structured, whereas "perceive-red" is not. In the former, "redly" ostensibly picks out a property of the property picked out by "perceive," whereas in the latter "perceive-red" is a single unbreakable (syntactically unstructured) predicate: "red" is no more a syntactic part of "perceive-red" than "apple" is part of "pineapple." Adverbialization and hyphenation might each have its own advantages (e.g.: the former captures compositionality better, the latter avoids ostensibly denoting second-order properties). For present purposes, the relative advantages and disadvantages of various apparati for a non-relational account of intentionality can be treated as an in-house matter to be settled later.[21]

Third, a non-relational account of intentionality does not imply that representors do not bear relations in virtue of representing. What it implies is rather the converse: that representors do not represent in virtue of bearing relations. Thus it may well be that when a representor represents an existent, it enters a

relation to it, and enters that relation in virtue of representing it. This relation we may well call "the representation relation." Thus, thinking of one's mother may involve bearing a representation relation to her—*in addition* to bearing the non-relational property of representing mother-wise.[22] What is crucial is that it is not in virtue of bearing the representation relation to one's mother that one's mother-thought has intentionality. Rather, the thought has intentionality purely in virtue of instantiating the non-relational property of representing mother-wise. On this view, there exist both a non-relational representation property and a representation relation; it is just that only the former constitutes intentionality.[23]

To summarize this part of the discussion: in Ducasse and Chisholm's account of sense perception, the adverbial move was deployed relatively narrowly, to circumvent the need to posit mental concreta to which perceptual experiences are related. In the present context, the adverbial move is deployed with the more sweeping purpose of solving the problem of intentional inexistence. Your thought of Bigfoot does not involve constitutively a relation to Bigfoot, on the present account, but rather the instantiation of a non-relational property of representing Bigfoot-wise. This is why your thought can represent a non-existent even though it cannot bear a relation to a non-existent. Thus an adverbial account of intentionality solves the problem of intentional inexistence.

The account also accommodates the possibility of representing indeterminates. The reason one can think of a parrot without thinking either of a blue-eyed parrot or of a brown-eyed parrot is that thinking of a parrot is a matter of having a thought that represents parrot-wise, and a thought can represent parrot-wise without representing either blue-eyed-parrot-wise or brown-eyed-parrot-wise. Granted, one cannot bear any relation to a parrot without bearing that relation to either a blue-eyed or a brown-eyed parrot. But this presents no problem if representing a parrot is not (constitutively) a matter of bearing a relation to it.

Conversely, the reason one can think of a rabbit even when the relevant existent is indeterminate as between a rabbit and a fusion of rabbit parts is that thinking of a rabbit is a matter of having a thought that represents rabbit-wise, and a thought can represent rabbit-wise even if there is no fact of the matter as to whether the relevant existent is a rabbit or a rabbit-like fusion. Granted, if the relevant part of the world is indeterminate as between a rabbit and a rabbit-like fusion, one cannot determinately bear a relation to one and not the other, since one can only bear a relation to the relevant part of the world. But again this presents no problem if representing a rabbit is not (constitutively) a matter of standing in a relation to the rabbit part of the world.

4. The Limits of Adverbialism

The adverbial account of intentionality does face a number of serious difficulties, however. I will mention what I take to be the three main ones.

The best-known objection to the adverbial account of sense perception is due to Jackson (1977).[24] According to Jackson, the adverbial paraphrase cannot distinguish between "you perceive a white rectangle and a purple circle" from "you perceive a white circle and a purple rectangle." Both will have to be paraphrased into "you perceive whitely and purply and rectangularly and circularly."[25] We *could* paraphrase the former into "you perceive white-rectangularly and purple-circularly" and the latter into "you perceive white-circularly and purple-rectangularly" (which *are* different), but such paraphrase would lose the compositionality of representation. This compositionality manifests itself, e.g., in the fact that we can infer "you perceive a rectangle" from "you perceive a purple rectangle." Since "purple-rectangularly" is a simple, unstructured expression (of which "rectangularly" is only a morphological, not syntactic, part), however, we *cannot* infer "you perceive rectangularly" from "you perceive purple-rectangularly."[26]

Another objection is that semantic externalism shows that at least *some* forms of intentionality are certainly relational, since the individuation of the relevant kinds of representation is determined partly by the nature of what is represented. In particular, some linguistic representations—most notably, natural kind terms and proper names—must involve relations to their referents, as the Twin Earth thought-experiment (Putnam 1975) allegedly demonstrates. Presumably, the same would apply to natural kind and singular *thoughts*.

We will return to these objections in the next section. I now want to focus on a third objection, one that is less technical but probably goes deeper to the motivation for avoiding adverbial accounts. This is the simple fact that, although we can *say* things like "this painting represents Bigfoot-wise," it is not at all clear what that means. The problem is that the property of representing Bigfoot-wise is quite mysterious.[27] Consider an inked token of the word "Bigfoot." We could spend considerable time examining the non-relational properties of the word, but I predict that we would never find a property that could credibly be taken to constitute representing Bigfoot-wise. A token word's only remotely relevant non-relational properties seem to be its graphic or phonetic properties. Surely the property of representing Bigfoot-wise is not a graphic or phonetic property of "Bigfoot." The word "Bigfoot" could mean New York City in a language yet to be invented, and tokens of it in that language would be representationally different, but graphically and phonetically (hence non-relationally) indistinguishable, from its English tokens.

What are the adverbialist's options for responding to this objection? One option is to try to identify a clever non-relational property of "Bigfoot" that might constitute representing Bigfoot-wise. But that seems hopeless. A more promising option is to point out that, from the fact that *linguistic* representations lack non-relational properties that might constitute representing *x*-wise, it does not follow that *all* representations lack such properties. Perhaps it is the case that the adverbial move is suitable for some kinds of intentionality but not others.

This option is problematic on two scores, however. First, it remains to be seen what non-relational property can capture representing *x*-wise in any non-linguistic intentionality. Secondly and more importantly, this move seems to amount to giving up on a general solution to the problem of intentional inexistence. Even if the adverbial move solves the problem of intentional inexistence for those forms of intentionality that are amenable to adverbial treatment, there would be other forms of intentionality for which the problem would persist.[28]

5. Phenomenal Intentionality

We have explored the option of rejecting (c) by claiming that representing something does not constitutively involve bearing a relation to what is represented, because it is a non-relational property of the representor. However, this is not the only way to reject (c). Another way would be to claim that representing something does not involve bearing a relation to what is represented because it is grounded in relations to things other than what is represented. There is also a third way of rejecting (c), which is a mix of the other two, claiming that some representing is non-relational and some involves relations to things other than what is represented. Although the notion that all representation is grounded in relations to things other than what is represented is extremely implausible on the face of it, in this section I will suggest that the mixed view is actually plausible.[29]

The sort of mixed view I have in mind involves two key moves. The first is the thesis that there is a kind of intentionality that is amenable to adverbial treatment. The second is the thesis that all other intentionality is derivative from this adverbial intentionality. The upshot is the claim that the only kind of basic, non-derivative intentionality is one that is amenable to an adverbial treatment.

More specifically, I want to suggest that a restricted adverbialism might work for the intentionality of phenomenally conscious mental states, or more accurately the intentionality phenomenally conscious states have *in virtue of being* phenomenally conscious states. For this form of intentionality, independently dubbed *phenomenal intentionality* by Loar (2002) and Horgan and Tienson (2002), there *is* a non-relational property that is a good candidate for constituting the property of representing Bigfoot-wise, namely, the representation's phenomenal character.[30] A preliminary defense of this view will be offered in §6.

In addition, I want to claim that it is quite plausible that all intentionality derives from phenomenal intentionality (all representation derives from conscious representation). The idea is that there is a distinction to be made between derivative and non-derivative intentionality, and that only phenomenal intentionality is non-derivative. On this view, conscious representations are the only representations that represent in and of themselves, not because they are suitably related to other representations. Non-conscious representations, by contrast, represent only insofar as they are suitably related to conscious representations

(namely, by whatever relation underlies the "derivation" of derivative intentionality from non-derivative intentionality). A preliminary defense of this view will be offered in §7.

With these two theses in place, a general solution to the problem of intentional inexistence is possible after all. The solution has two parts, one for conscious representations and one for non-conscious ones. First, a conscious representation of Bigfoot does not represent Bigfoot in virtue of bearing a relation to Bigfoot, but in virtue of instantiating some non-relational property of representing Bigfoot-wise. Second, a non-conscious representation of Bigfoot does not represent Bigfoot in virtue of bearing a relation to Bigfoot, but in virtue of bearing a relation to conscious representations of Bigfoot.[31] More generally, non-conscious representations derive their intentionality from conscious representations, not from relations to intentional objects, while conscious representations have an altogether non-relational intentionality.[32]

To make the case for this general solution, I would have to defend the two theses upon which it is founded: (i) that phenomenal intentionality is non-relational, and (ii) that all non-phenomenal intentionality derives from phenomenal intentionality. Clearly, neither task can be seriously undertaken here. Fortunately, however, these theses have been defended elsewhere in the literature: the first in (e.g.) Horgan and Tienson 2002, Horgan et al. 2004, and Loar 2002; the second in (e.g.) Kriegel 2003a, Loar 2002, McGinn 1988, and Searle 1991, 1992.[33] My goal here will be merely to show that both theses are quite plausible. This would suffice to show that the solution to the problem of intentional inexistence they afford is viable and believable—the target I have set for myself in §1.[34]

At the background here is a tension between two opposing outlooks on intentionality. The first is the mainstream research program on naturalizing intentionality mentioned in §2, where the goal is to identify a natural relation that holds between two things when and only when one represents the other. The second is a more heterodox outlook that has been receiving growing attention in recent years. On this heterodox outlook, phenomenal intentionality is in some sense special and basic,[35] in a way that places it at the heart of the theory of intentionality and makes an understanding of consciousness a precondition for understanding intentionality.[36]

I will not attempt to develop here a full defense of this outlook. I am merely concerned to point out what (as far as I can tell) has not been observed to-date, namely, that this picture delivers a solution to the problem of intentional inexistence. This is particularly significant given that, as we saw in §2, approaches to intentionality that follow the mainstream research program on naturalizing intentionality—by naturalizing "the representation relation"—are ill-positioned to deliver a viable solution to this problem. This suggests that the issue of intentional inexistence offers a deep reason to prefer the heterodox outlook that focuses on phenomenal intentionality *at the expense of* the mainstream picture that focuses on the representation relation.

(None of this is to say that the heterodox outlook is inherently anti-naturalistic. If the heterodox outlook is correct, the implication is merely that, if intentionality is to be naturalized, it is through the naturalization of the non-relational property of representing *x*-wise, not through the naturalization of a representation relation. That is, what the heterodox outlook implies for naturalism is not an abandonment of the project of naturalizing intentionality, but rather a redirecting of it.)

6. Adverbialism about Phenomenal Intentionality

Among the things that represent are phenomenally conscious mental states.[37] Thus, when you conjure up a visual image of Bigfoot, you enter a phenomenally conscious state that represents Bigfoot. The thesis of this section is that such states have a non-relational intentionality. If so, there exists a kind of intentionality that is amenable to adverbial treatment.

One way to think of the main point of this section is as follows. In §4, we noted that when we examine the word "Bigfoot," we cannot find among its non-relational properties any property that could be credibly taken to constitute the property of representing Bigfoot-wise. According to this section's thesis, when we examine a phenomenal visualization of Bigfoot, we *do* find a non-relational property that could be plausibly construed as the property of representing Bigfoot-wise. This is simply the visualization's *phenomenal character*.[38] More generally, while for non-phenomenal representations there is not a non-relational property that might credibly be taken to constitute representing what they do, for phenomenal representations there is, namely, their phenomenal character.

To show that the phenomenal character of some mental states constitutes a non-relational intentionality, we would have to show that these mental states have a phenomenal property that is at once intentional and non-relational. To a first approximation, the thesis is that some mental states have a phenomenal property and an intentional property, such that (i) the relevant intentional property is constituted (realized?) by the relevant phenomenal property and (ii) the relevant phenomenal property is non-relational. More formally:

(PDT) For some mental states $M_1, \ldots, M_n$, there is a phenomenal property P and an intentional property I, such that for each M_i, (i) M_i is P, (ii) M_i is I, (iii) M_i's being P constitutes M_i's being I, and (iv) P is a non-relational property of M_i.[39]

The argument I will present for this claim is basically this: 1) some phenomenally conscious states are intentional, and intentional in virtue of being phenomenal; 2) these states' phenomenal character is a non-relational property of theirs; therefore, 3) some phenomenally conscious states are intentional, and intentional in virtue of a non-relational property of theirs.[40]

The first premise starts from the observation that some phenomenally conscious states are intentional, but goes beyond it. The claim is that such conscious states have an intentional content which they carry purely in virtue of their phenomenal character. They may have intentional properties unrelated to their phenomenal character, but they *also* have some intentional properties that are instantiated in virtue of, indeed are constituted by, their phenomenal properties. When a mental state M has an intentional content that is constituted by its phenomenal character, we may say that M exhibits a *phenomenal intentionality*, or is phenomenally intentional.[41] The claim before us, then, is that some mental states are phenomenally intentional, or that phenomenal intentionality is actually instantiated.[42] Call this the *phenomenal intentionality thesis*; it consists in (PDT) minus condition (iv). I now present two considerations in favor of the phenomenal intentionality thesis.

The first thing to note about the phenomenal character of many conscious states is that, in virtue of having this phenomenal character, the states in question are *assessable for accuracy* (Siewert 1998). Thus, a visual experience of a box of raspberries can be more or less accurate. But what would endow the experience with its level of accuracy if not the degree to which the world is the way the experience's phenomenal character, in some sense, "says" it is? Thus the experience has accuracy conditions that effectively constitute a representational content: under the condition that things are indeed the way they seem, the experience is accurate; under the condition that they are not, it is not. These accuracy conditions may be indistinguishable across phenomenally indistinguishable states. Thus there is a kind of intentionality that phenomenal duplicates share: duplicate the phenomenology, and you have duplicated the intentionality.

That some phenomenal character is inherently intentional, and constitutes an intentional content, is brought out clearly in the following thought experiment.[43] Suppose your brain is hooked up to a machine—call it "the inverter"—that can rewire the visual channels in your brain in such a way that when the operator presses the right button, your color qualia are inverted.[44] Suppose further that, while hooked to the inverter, you are looking at pictures of red apples passing on a monitor for five seconds each. In some cases, after three seconds the operator alters the picture on the monitor into a picture of a green apple. In other cases, after three seconds she presses the button on the inverter, thus inverting your color qualia. The point is this: from the first-person perspective, you will not be able to tell apart those trials. You will not be able to tell whether the operator changed the picture on the monitor or inverted your qualia. Whether it is the world that changed or your brain, your experience of the change is the same—it is an experience as of the world changing. This suggests that your experience is inherently directed at the world, that is, is inherently intentional.

This fact about visual experience has been appreciated quite often over the past couple of decades. It is basically the point often referred to in the relevant literature as the "transparency of experience" (Harman 1990).[45] The idea is that

whenever we try to introspect the qualities of our conscious experiences, we manage only to become aware of the properties of what these are experiences *of*. This suggests that the phenomenal character of our conscious experiences is intentional.[46]

(It may be thought odd to cite the transparency of experience in defense of phenomenal intentionality. The transparency claim is often adduced on behalf of the attempt to naturalize phenomenal character by combining an intentionalist account of it with traditional naturalistic accounts of intentionality in terms of the representation relation. By contrast, the notion of phenomenal intentionality is often touted in the context of rejecting this program. However, regardless of the wider motivations commonly at play in these discussions, the transparency of experience—the fact that phenomenology appears, from the first-person perspective, to be inherently intentional—is surely evidence for the existence of phenomenally constituted intentionality.)

The second premise in the argument for PDT is the claim that the phenomenal character of conscious states is a non-relational property of these states.[47] Call this thesis *phenomenal internalism*. In abstraction from any theoretical considerations, this claim is very intuitive: it just seems that what it is like for a person to undergo a certain experience, though perhaps *causally* dependent upon external affairs, is not *constitutively* dependent on them. This intuition is brought out vividly when we consider Cartesian demon and brain-in-vat scenarios (Horgan et al. 2004). It seems that a brain-in-vat duplicate of you should have the same phenomenal experiences as you. If so, non-relational duplicates are also phenomenal duplicates. To abstract from immaterial complications to do with vats and their operation, we may consider an otherwise disembodied brain, or for that matter a soul, floating through otherwise empty space but (by sheer luck) undergoing conscious experiences that are phenomenally indistinguishable from yours. Since there is nothing for this "space soul" to bear relations to, it would seem that the phenomenal character of its experiences is a non-relational property.[48]

One might attempt a Putnam-style externalist treatment of envatted brains, arguing that in virtue of standing in representational relations to different entities, an envatted brain is in fact in a different intentional state from its non-envatted duplicate (Putnam 1981). However, since our concern here is with *phenomenal* intentionality, this would have to mean that the envatted brain is in a different *phenomenal* state as well, which is false *ex hypothesi*. It might be claimed that what this shows is that in fact the hypothesis we are entertaining is impossible, but there is no reason to accept such a claim.[49] Worse, to apply the Putnamian treatment to the space soul would be to claim that the space soul has no intentional states, since its internal states have no systematic causes at all. But that borders on the absurd: right now your space-soul duplicate is having an experience as of reading a paper about intentional inexistence and phenomenal intentionality, and in a moment it might pause to visualize a camel; clearly, these internal states are about a paper and a camel.[50]

There is thus a good case for construing phenomenal character both as non-relational and as intentional, that is, for both phenomenal internalism and the phenomenal intentionality thesis. Some philosophers have used the latter against the former. According to these "phenomenal externalists," the phenomenal character of a conscious state is a relational property of it (Dretske 1996; Lycan 2001). The reasoning is this: phenomenal character is intentional, and all intentionality is relational, so phenomenal character must be relational. This reasoning mirrors perfectly ours: phenomenal character is intentional, and it is non-relational, so phenomenal intentionality is non-relational. To decide between these two package deals, we must determine which is the more plausible premise: that all intentionality is relational (as per the externalist package) or that phenomenal character is non-relational (as per the internalist one). In the present context, the externalist's position seems to beg the question. But even if this is just an artifact of this particular paper's dialectic, there may in fact be no non-question-begging way to decide between the two package deals. What this means is that there is no non-question-begging objection here to our holding both phenomenal internalism and the phenomenal intentionality thesis.

If we accept that phenomenal character is both sometimes intentional and always non-relational, it follows that there is a kind of non-relational intentionality. The picture we get is one where many conscious states involve something like *intrinsic phenomenal directedness*: some sort of phenomenally constituted non-relational feature of being-directed-at-something. In an adverbial vein, we might say that a visual experience of a rabbit has the non-relational phenomenal property of being rabbit-ward-esque (i.e., intrinsically directed at a rabbit), and a visualization of Bigfoot has the non-relational phenomenal property of being Bigfoot-ward-esque (intrinsically directed at Bigfoot).

To be sure, there is something perplexing about the notion of intrinsic phenomenal directedness. Is not saying that phenomenal experience presents us with the external world precisely saying that it is inherently relational? The short answer is No: to say that phenomenal experience presents us with the external world is to say that it is inherently *directed* at the external world, not that it is inherently *related* to the external world. The former would entail the latter only if directedness at the external world involved a relation to it. The claim made here is that there is a kind of phenomenal directedness that does not involve a relation to the external world. At the same time, there is admittedly something somewhat mystifying about the idea of non-relational directedness.

How to understand this notion of intrinsic phenomenal directedness is an important question. Loar (2002) and Horgan and Tienson (2002) offer their own glosses, which may not be entirely satisfactory.[51] To my mind, the best way to wrap one's mind around the notion of phenomenal directedness is by triangulation from the various theoretical constraints on it. That is to say, the best answer to the question "what is phenomenal directedness?" is: it is the kind of property instantiated by a space soul having camel experiences that are transparent to introspection! That such a property exists is not something

we are asserting here lightly. It is a substantive claim: the claim made in PDT, which we may call the *phenomenal directedness thesis*.

If there is such a property as intrinsic phenomenal directedness, then our objections to adverbialism about intentionality from §4 are unconvincing as objections to adverbialism about phenomenal intentionality. Clearly, the objection that representations of Bigfoot have no non-relational properties that are good candidates for constituting anything like representing Bigfoot-wise is no longer troubling. For phenomenally conscious representations of Bigfoot have a perfect candidate for constituting the property of representing Bigfoot-wise, namely, the property of being Bigfoot-ward-esque. That is to say, phenomenal directedness is a perfect candidate for constituting non-relational, adverbialized intentionality.[52]

As for the externalist objection, although semantic externalism is widely popular when applied to (some types of) linguistic representation, as well as belief and the propositional attitudes, it is much less popular as applied to phenomenal experiences. Thus phenomenal externalism is still a non-starter for most philosophers, who typically maintain that the intentional contents of phenomenal water and twin-water experiences are the same. Even among those who accept that phenomenally conscious states have an externally determined "wide" content, it is often thought that such states must also have, in addition, a "narrow" content that is invariant across twins.[53] As long as conscious states are conceded to have such a narrow content, there is no externalist objection to an adverbial account of phenomenal intentionality.

It might be insisted that, even if conscious representations of natural kinds can be construed adverbially, *singular* conscious representations (i.e., conscious representations of concrete particulars) cannot. When you consciously think of your mother and your Twin on Twin-Earth thinks of his or her mother, it may seem that you are in a different representational state, but the phenomenology of your thoughts is the same.

There are two ways to respond to this objection. One is to deny that you and your Twin have the same phenomenology; the other is to deny that you and your Twin are in a different representational state. I prefer the latter response.[54] There are two moves involved in this. The first is to construe the singular component of singular representations—what makes them about one specific concrete particular and not about a qualitatively indistinguishable particular—as a hidden indexical. The second is to adopt a broadly internalist treatment of indexicals, say in the style of Searle's (1983) token-reflexive account or the two-dimensional semantics of Jackson (1998) and Chalmers (2002, 2006).[55]

Consider finally Jackson's objection, and the idea that adverbial representation would have to lack compositional structure. There is a sense in which this is accurate: if adverbialism is true, we cannot infer from the fact that you consciously represent purple-rectangle-wise that you consciously represent rectangle-wise. This is the sense in which we cannot (validly) go through the following inferential process: you consciously represent purple-rectangle-wise,

therefore you consciously represent purple-wise and rectangle-wise, therefore you consciously represent rectangle-wise. However, this may not be the only kind of valid inferential process in the area. If, as is presumably the case, consciously representing rectangle-wise is a determinable of which consciously representing purple-rectangle-wise is a determinate, then we *can* (validly) go through something like the following inferential process: you consciously represent purple-rectangle-wise, consciously representing purple-rectangle-wise is consciously representing rectangle-wise, therefore you consciously represent rectangle-wise. Compare: we can infer from the fact that this is a strawberry that it is a berry. We do not do so by reasoning that since it is a strawberry, it must be a straw and a berry, and therefore a berry. Rather, assuming that being a berry is a determinable of which being a strawberry is a determinate, we reason that since it is a strawberry, it is a berry.[56]

Thus all three objections to adverbialism about intentionality are unconvincing as objections to adverbialism about *phenomenal* intentionality.[57] Phenomenal intentionality allows inferences from representational determinates to representational determinables, does not allow Twin-Earth-style scenarios, and has a natural candidate for constituting the property of representing *x*-wise. Furthermore, as we saw in the first half of this section, there is an independent argument for adverbialism about phenomenal intentionality (and more strongly for the thesis that there exists a non-relational intentionality), namely, the argument from phenomenal internalism and the phenomenal intentionality thesis. I conclude that adverbialism, although implausible as applied to all representations, is quite plausible as applied to conscious ones.[58]

7. Non-Phenomenal Intentionality

That an adverbial treatment of phenomenal intentionality is viable helps us with the problem of intentional inexistence only if one of the following obtains: (i) there is no intentionality other than phenomenal intentionality or (ii) all non-phenomenal intentionality derives from phenomenal intentionality. Obviously, (i) would help us because it would make an adverbial account of phenomenal intentionality entail an adverbial account of all intentionality. Although (i) has been held by some philosophers (Georgalis 2006, Strawson 2004), I will not pursue it here. Instead, I will pursue the more lenient (ii), which helps us with the problem of intentional inexistence because it makes entities endowed with non-phenomenal intentionality represent not in virtue of standing in a relation to potentially non-existent entities, but rather in virtue of standing in a relation to entities endowed with phenomenal intentionality.

Option (ii) presupposes a distinction between *derivative* intentionality and *non-derivative* intentionality. The distinction enjoys relatively wide adherence in discussions of representation. It surfaced originally in the context of claims to the effect that the intentionality of (natural) language is derivative from the

intentionality of the mental, whereas the intentionality of the mental is in some sense intrinsic to it and does not derive from any other source.[59] The thesis of this section is that in reality only *phenomenal* intentionality is non-derivative, and all other intentionality—including non-phenomenal mental intentionality—derives from phenomenal intentionality.

For present purposes, we may construe the thesis that only phenomenally conscious states have non-derivative intentionality as follows:

(NPI) There is a relation R, such that for any non-phenomenal representation N, there are phenomenal representations $P_1, \ldots, P_n$, such that N is the representation it is, and a representation at all, in virtue of bearing R to $P_1, \ldots, P_n$.

I have explored the case for NPI at length in "Is Intentionality Dependent upon Consciousness?" (Kriegel 2003a), and concluded in its favor. Here I will recapitulate only in the barest outlines the main consideration that I took in that paper to favor NPI. (The case presented here might seem underdeveloped; but for its full development I refer the reader to §3 of Kriegel 2003a.) I will then offer a second, newer argument in favor of NPI.

The older argument can be appreciated in three steps. First, it seems plausible that something counts as a genuine representation only if it is at least *possible* for it to represent what it does *to someone*. We may put this more precisely by saying that *x* is a representation of *y* only if there is some *z*, such that, possibly, *x* represents *y* to *z*.[60] Secondly, if this is so, then it would seem that for something to qualify as a representation in and of itself, and not merely derivatively, it would have to be able to represent to someone in and of itself, and not merely derivatively. Since it is a necessary feature of representations that they can in principle represent to someone, it is a necessary feature of non-derivative representations that (in principle) they can non-derivatively represent to someone. Finally, it is highly plausible that only conscious representations represent to someone in and of themselves.

This final claim is based on the observation that there is an important difference between conscious and non-conscious representations (Georgalis 2006, Kriegel 2003a, McGinn 1988). If you perceive a rabbit consciously, your perceptual experience represents the rabbit *to you*; if you perceive a rabbit subliminally, or blindsightedly, there is a straightforward sense in which your perception represents the rabbit, but not *to you*.[61] More generally, conscious representations seem to always envelop within them both a representation-of component and a representation-to component, whereas non-conscious representations exhibit universally only the representation-of component.[62]

The newer argument for NPI is based on the methodological idea that, in studying intentionality, we should adopt as our starting point the least committal approach to intentionality, and then make the minimal departure from it that would enable a workable account of intentionality. The claim is that, when

we do this, we end up with an account of intentionality that is committed to NPI.

As "the least committal approach to intentionality," I have in mind a kind of *interpretational semantics*. Interpretational semantics works with the assumption that things ought to be treated as exhibiting intentionality just if so treating them is explanatorily advantageous. As Dennett (1987) puts it, we ought to take an *intentional stance* toward a set of phenomena, and treat them as intentional, when and only when there is some distinctive explanatory payoff in doing so. The thought is that there is no deep fact of the matter as to whether anything in the world actually exhibits intentionality; it is just an instrumentally useful way of talking. Dennett himself conceived of his interpretational semantics as involving treating intentionality as a *useful fiction*. But the "fiction" part of this is not inherent to interpretational semantics as such. It is possible to take interpretational semantics to imply merely that intentionality is a useful posit and remain silent on the matter of fictitiousness. It is this version of interpretational semantics that I think of as "the least committal approach" to intentionality. It is least committal in that it does not commit to anything in the world being mind-independently endowed with anything like intentionality. On this view, something has intentional content just in case it usefully admits of interpretation, where "interpretation" is the act of ascribing content (to "interpret" a stop sign is to ascribe intentional content to it).

According to this sort of interpretational semantics, *all* intentionality is derivative (Dennett 1990). For on this picture, all representations derive their intentional content from acts of interpretation that target them. Thus, the word "Bigfoot" represents what it does, and represents at all, in virtue of being interpreted to stand for Bigfoot. It thus derives its representational content from another representation, one whose content is "'Bigfoot' stands for Bigfoot." In general, all representations represent what they do, and at all, in virtue of bearing the right relation (the relation underlying useful interpretability) to other representations (interpretative ones). None represent in virtue of bearing a relation to what is represented, nor in virtue of their intrinsic character.

An immediate difficulty with interpretational semantics, so construed, is the fact that interpretations are themselves representational and contentful, and in fact could not confer intentionality on other representations if they were not. Given that they do not have content in and of themselves, however, they must await conferral of content by yet further interpretations—and we are quickly off on an infinite regress. The interpretationist's hope must be that the wider system of representations will function in such a way that intentionality will arise out of the web of mutual interpretations, but it is unclear how this might happen.

In a way, the interpretationist faces here an acute version of the problem faced by coherentists in epistemology, who must take epistemic justification to somehow emerge from the mutual relations of epistemic support among the various beliefs that comprise one's belief system. In epistemology, the foundationalist fix is to suppose that some beliefs (the "foundational" ones) are

either self-justifying or not in need of justification. The parallel move in semantics would ground the web of interpretations in a select group of representations that are either self-interpreting or not in need of interpretation. To say that they are self-interpreting is to say that they confer content upon themselves; to say that they are not in need of interpretation is to say that they do not need to have content conferred upon them, presumably because they have content in and of themselves. These elite representations, which are either self-interpreting or interpretation-exempt, would have their content non-derivatively, in the sense that their intentionality would not be grounded in the intentionality of other representations.

The claim I would like to make now is that phenomenally conscious representations are the only natural candidates to serve as the "foundational" representations. This is so on both accounts of non-derivative content, the self-interpretation account and the interpretation-exemption account.

Suppose non-derivative content is to be found only in self-interpreting representations. Given that interpreting something is a matter of representing it to have a certain representational content, self-interpreting representations would have to be representations that represented themselves to have a certain representational content. Thus, a self-interpreting representation of *x* would be a representation with a dual content: a first-order content "*x*" and a second-order content "this very representation represents *x*." The question, then, is which representations are most plausibly construed as self-representing. The most plausible answer is that these are phenomenally conscious representations.[63] There may be here and there individual non-conscious representations that happen to represent themselves (e.g., "this is a sentence"), and some may even represent themselves to represent something ("this is a sentence about a sentence"). But if we seek a well-demarcated set of systematically self-representing representations, surely our best bet is the set of conscious representations.

Alternatively, suppose that non-derivative content is to be found in a select group of representations that are altogether exempt from the need to be interpreted. These representations have whatever content they do in and of themselves and independently of any interpretation (either by other representations or by themselves). If, as I have argued in the previous section, there are representations endowed with intrinsic phenomenal directedness, then it would be most natural to suppose that these are the interpretation-exempt representations. Surely the need for interpretation would be obviated by the exhibiting of intrinsic phenomenal directedness if by anything at all. Since only phenomenally conscious representations are endowed with intrinsic phenomenal directedness, it follows that only conscious representations are interpretation-exempt. So on this account too, only conscious representations have their content non-derivatively.

We have seen that on both the self-interpretation and interpretation-exemption accounts of non-derivative content, conscious representations are the best candidates for having non-derivative content. We have also seen that the need for such self-interpreting or interpretation-exempt representations arises

from the inherent limitations of global interpretationism, which is otherwise the least committal approach to intentionality. The point is this: if we start out with global interpretationism, and make the minimum modification necessary to make it work, we end up with the thesis that only conscious representations are non-derivatively such. Thus the minimally committal *but workable* picture of intentionality is one that designates conscious representations as a select group of representations with non-derivative content, and takes all other representational content to derives from conscious representations via something like the interpretational mechanism of the intentional stance. This is a version of NPI.

8. Conclusion: A Viable Solution to the Problem of Intentional Inexistence

I conclude that both claims needed for the solution to the problem of intentional inexistence sketched in §5 are highly plausible. The solution is this. For all kinds of representation, the inconsistent triad presented in §1 is avoided by rejecting (c), the proposition that representing something involves constitutively bearing a relation to it. But the reason (c) is rejected is different for conscious and for non-conscious representations. For conscious representations, (c) is rejected on the grounds that one can consciously represent something without being related to it, namely by instantiating the relevant non-relational phenomenal property. For non-conscious representations, (c) is rejected on the grounds that one can non-consciously represent something without being related to it, namely by bearing the right relation to conscious representations of it. Thus, a conscious representation of Bigfoot represents Bigfoot not in virtue of bearing a relation to Bigfoot, but in virtue of instantiating the non-relational property of representing Bigfoot-wise; and a non-conscious representation of Bigfoot represents Bigfoot not in virtue of bearing a relation to Bigfoot, but in virtue of bearing a relation to representations that represent Bigfoot-wise.

This solution depends, of course, on the adoption of a somewhat heterodox outlook on intentionality, according to which (i) all intentionality derives from phenomenal intentionality and (ii) phenomenal intentionality is non-relational or "adverbial." I have offered a preliminary defense of both planks of this outlook, and noted (in §5) that fuller defenses can be found elsewhere in the literature. For my part, I find this outlook independently plausible, on broadly phenomenological grounds. It is therefore particularly satisfying to observe that it also delivers a solution to the problem of intentional inexistence.

For readers who do *not* find this outlook immediately attractive, the present paper can be offered as a challenge. The challenge is to show how the problem of intentional inexistence can be solved *without* adopting this heterodox outlook. Given that the rejection of (a) and (b) seems hopeless, and that it is unclear how to reject (c) on other grounds, the challenge ought to be taken seriously. Thus the present paper can also be taken as an indirect argument for the heterodox outlook on intentionality adopted here.

My primary goal here was not, however, merely to erect a challenge for opponents of the heterodox outlook. Rather, it was the more positive goal of outlining a viable solution to the problem of intentional inexistence. In offering preliminary defenses of each plank of the heterodox outlook (in §§6–7), I was concerned to convince the reader that both are plausible, or at least that neither is blatantly implausible. For it is already remarkable that the two claims are quite plausible, and that together they deliver a solution to one of the perennial problems of philosophy.[64]

Notes

1. There is a usage of "perceive" as a success verb. In that usage, there can be no perceiving of something that does not exist. But there can still be an apparent perceiving of something that does not exist, so the problem would still arise—it would just have to be reformulated. In this paper, I use "perceives" in a non-success way for the sake of ease of discussion.
2. In the literature, one often finds the idea of explaining intentionality in terms of representation treated as a substantive claim. There is certainly a sense of "representation" in which it would be so, but here I will use "representation" in a sense in which it is not a substantive claim that intentionality involves, and is even exhausted by, representation. I do so mainly for esthetic reasons: while "representation" has a natural verb form in "represents," "intentionality" does not (the verb "intends" in the relevant sense is sometimes used in this context, but is ugly).
3. There may be a question as to whether "represents" is a success verb. If it is, then what is common to all of the above is not necessarily that they represent, but that they *attempt* to represent. For present purposes, we may treat this as a matter of bookkeeping.
4. Perhaps more formally, though somewhat awkwardly, we might write: (a) Possibly, there is an x, such that x represents a non-existent; (b) Necessarily, there is no x, such that x bears a relation to a non-existent; (c) Necessarily, for any x, if x represents something, then x bears a relation to it. This formalization is awkward, because any formalization would be problematic, inasmuch as the adoption of certain formalisms is liable to beg all the important questions in this area. The formalization I use here is somewhat awkward, but has the advantage of not begging the questions we are interested in here.
5. We have focused so far on *imaginative* representations of non-existents (e.g., thinking of a dragon), but the problem arises also in the context of *hallucinatory* representations of non-existents. Thus, if you hallucinate a floating lemon in front of you, you represent something that does not exist. In the imaginative case, there are no dragons, so what is represented does not exist. In the hallucinatory case, there do exist lemons, just not the token lemon you represent. Thus imaginative representation of non-existents involves representation types every token of which represent non-existents, whereas hallucinatory representations of non-existents are token representations that belong to representation types some tokens of which may represent existents. (If what one hallucinates is Bigfoot, then one's

hallucination will not belong to a type some tokens of which are accurate representations. But if what one hallucinates is a lemon, then one's hallucination will represent a non-existent lemon, but will belong to a representation type some tokens of which are (potentially accurate) representations of existing lemons.) In this paper, we will concentrate on the imaginative variety of representation of non-existents, just because doing so will allow us to disregard problems in the philosophy of perception that do not go to the core of the problem of intentional inexistence.

6. Thus we could construct the following inconsistent triad: one can represent indeterminates; one cannot bear a relation to indeterminates; representing something involves bearing a relation to it.
7. Note well: this is the converse of the problem of representing indeterminates only if it is indeed indeterminate which of (i)–(v) exist(s). Clearly, most metaphysicians today hold that it is determinate which of (i)–(v) exist(s)—even when their views on the matter are otherwise counter-intuitive. Nonetheless, it a viable position, one that is apparently shared by many philosophers outside the metaphysics community, that there is no fact of the matter as to which of (i)–(v) exist(s). Note also that this is not quite Quine's own problem of the inscrutability of reference, though it is inspired by similar considerations.
8. Another pair of related problems concerns the way sameness and difference of what is represented and sameness and difference of what exists come apart. Consider Hesperus and Phosphorus. If we want to individuate representations in terms of what is represented, then given that the representations of Hesperus and Phosphorus are different representations, we must say that what is represented in those representations is different. Yet what exists is clearly one and the same. The converse case is presented by water and twin-water, where what is represented is in some good sense the same, yet what exists is certainly different. We will not have occasion to treat of this pair of problems here. In any case, the problem of intentional inexistence goes deeper, in that it does not presuppose any assumptions about the individuation of representations.
9. Moreover, there are certain distinctions among "apparent thoughts" that still make sense, and which must be accounted for, regardless of whether these are not genuine thoughts. Thus, thinking of different kinds of non-existent seems to involve some difference in activity, but for the proponent of the current solution they are the same activity—the activity of failing to have a thought. (Try to convince a science-fiction fan that there is no difference between thinking of wizards and thinking of goblins and you will be accused, justifiably for once, with not knowing what you're talking about.) Or at least it is not clear how it can be argued that these are different kinds of activity.
10. One might attempt to go externalist with respect to seeming-to-oneself-to-think in such a way that the alleged external constituent of the content of an apparent dragon thought would also become a constituent of the content of the seeming. There are ways to extend Burge's (1988) solution to the problem of externalism and self-knowledge that may suggest this. But whatever the merits of Burge's own position, the extension is surely highly implausible.
11. Thus, the relations of perceiving and remembering are just perceptual intentionality and mnemonic intentionality; the point being that intentionality is really the only relation that is being excepted from the rule here.

12. Or perhaps: "if there were dragons, and they were more or less the way they are characterized in the relevant fictions, lions would be smaller than them."
13. To reject (c) is not necessarily to give up on naturalizing intentionality. It is just to give up on doing so in the mainstream way. To embrace (c) is to claim that representation is a non-relational property. It is perfectly possible to maintain that representation is a natural(izable) non-relational property. So naturalism is not averse to that approach to representation. It just requires a reconceptualization of how we are to go about naturalizing representation.
14. If the above critiques of the other options are right, this view is forced on us by the trivial facts that we can think of what does not exist but cannot bear relations to what does not exist. It follows that we can think of things to which we cannot bear relations, and therefore that it is possible to think of something without bearing a relation to it.
15. Sense-datum theory claimed that sense perception involves constitutively a relation to mental concreta. The idea was that if it could be shown that sense perception did not involve a relation at all, then it would follow that it did not involve a relation to mental concreta, and therefore that the sense-datum theory was false.
16. In the former construction, "perceive" is used as a transitive verb and "red" is its grammatical object. In the latter, "perceive" is used as an intransitive verb, and "redly" is an adverb that *modifies* it. Thus the sentence "you perceive red" suggests a relation between what is picked out by the subject term and what is picked out by the object term, a relation picked out by the verb. The sentence "you perceive redly" suggests that the verb picks out a property of what is picked out by the subject term and the adverb picks out a property of the property picked out by the verb (hence a second-order property of what is picked out by the grammatical subject).
17. This corresponds to the fact that the former features two names and a transitive verb, whereas the latter features one name, one intransitive verb, and one adverb. In general, adverbs and other grammatical modifiers pick out intrinsic modifications of the properties picked out by what they modify. In other words, if we were to read off from the grammar of "you perceive red" the ontological structure of the state of affairs described by that sentence, it would be that there is a perception relation that holds between you and a red; whereas for "you perceive redly," the ontological structure seems to be that you instantiate the non-relational property of perceiving, and the perceiving instantiates the non-relational (intrinsic-modification) property of being in some sense red-like.
18. In Chisholm's case, these would be mainly the epistemological considerations regarding the threat of skepticism about our knowledge of the external world.
19. More generally, "*x* represents Bigfoot" is to be paraphrased into "*x* represents Bigfoot-wise." Again, if we were to read ontology off of grammar, this would suggest a state of affairs involving a particular *x* instantiating the non-relational property of representing, which in turn instantiates a non-relational property of occurring Bigfoot-wise.
20. Goodman (1976) in fact uses this apparatus for something like the same purpose the adverbial apparatus is used here. More specifically, he uses it to develop a non-relational way of classifying—hence, presumably, type-identifying—

representational items of some kinds (pictures mainly). Goodman does not endorse, however, the extra ("reductive") move explored in the present section, of treating this manner of type-identification as the only viable one (nor the move of extending the non-relational classification to all representations).

21. There is also the apparatus developed by Crane (2001), who analyzes talk of intentional object as talk of things that are not really objects.
22. I am assuming here, with the eternalist, that all mothers exist, not only the live ones.
23. Having admitted both a non-relational representation property and a representation relation, it is again a matter of mere bookkeeping whether we would like to reserve the verb "to represent" for the latter. If one hears "represents" as a success verb, one certainly ought to do so, and construe the non-relational property in terms of attempting-to-represent. If one does not hear "represents" as a success verb, then there is no necessity to thus shuffle around the terminological scheme.
24. Since this is an objection to adverbialism about sense perception, and sense perception is a form of intentionality, the objection could be marshaled against generalized adverbialism about all forms of intentionality.
25. Trying to make use of the fact that the white rectangle and the purple circle are located at different parts of the subject's visual field will not help. For it will still be impossible to distinguish "you perceive a white rectangle to the left of a purple circle" from "you perceive a purple circle to the left of a white rectangle," since both would be analyzed into "you perceive whitely and rectangularly and purply and circularly and to-the-left-ly."
26. In the case of some kinds of representation, their compositionality may even be necessary for their acquirability (see Davidson 1984). Thus, on many views the compositionality of natural language is necessary for acquiring the ability to represent via natural language—i.e., the ability to use natural language. According to Fodor (1975), this is true also of the "language of thought."
27. To be sure, to the extent that it is cashed out in terms of standing in a representation relation to Bigfoot it could be made sense of. But that would defeat the purpose of the adverbial move.
28. Moreover, once we produce a solution to the problem for those other forms of intentionality, it might turn out that the relevant solution applies also to the forms of intentionality that are amenable to adverbial treatment. In that case, the other solution would obviate the need for the adverbial solution. Even if the other solution does not apply to the forms of intentionality amenable to adverbial treatment, there would be something odd about the view that the solution to the problem of intentional inexistence splits into two parts, one an adverbialist solution for some forms of intentionality and one a different solution for other forms. After all, the underlying problem is likely to be the same—or we would have to see why it is not—and so the solution would presumably have to be homogenous for different kinds of intentionality.
29. The second version seems unworkable. Perhaps Dennett's (1987) interpretational semantics is a form of this view: every representation represents in virtue of standing in the right relations to interpreters. But as we will see in §7, an obvious problem for Dennett's interpretationism is that the interpreters' acts of intentional ascription are themselves intentional, and so interpretationism appears to lead

us in a circle. Dennett is happy with a circle big enough, since the whole web of representational states and systems is in his view a mere useful fiction. If it be accepted that intentionality is a useful fiction, the second version of rejecting (c) might be viable. Here I proceed on the assumption that intentionality is not a useful fiction, and therefore that the second version of rejecting (c) is not really believable.

30. Moreover, Jackson's objection and the externalist objection are, I will argue, uncompelling as objections to adverbialism about phenomenal intentionality.
31. There is a question, of course, as to what the relevant relation is between non-conscious and conscious representations; but the answer to that question has to be provided by a general account of how derivative intentionality actually derives from non-derivative intentionality, and the question does not arise only in the context of the problem of intentional inexistence.
32. This is a general solution to the problem of intentional inexistence in that it does not leave any form of intentionality outside its scope. There is no need for a different solution which might eventually obviate this one.
33. A related view is that phenomenal intentionality is the only kind of intentionality. Strictly speaking, this is a version of the claim that all non-derivative intentionality is phenomenal, though a radical version. It is defended by Georgalis (2006) and Strawson (2004).
34. Of course, one not unimportant reason to embrace the two theses may be precisely that they afford a solution to one of the perennial problems of philosophy. In the next two sections, however, we will also get a glimpse of independent reasons for adopting the two theses.
35. The two theses appealed to here offer concrete ways of understanding the notion that phenomenal intentionality is special and basic; namely: that it is *special* in that it is non-relational or "adverbial," and that it is *basic* in that all other intentionality derives from it.
36. Ideas to this effect have recently been pursued by Georgalis (2006), Horgan and Tienson (2002), Horgan et al. (2004), Loar (2002), McGinn (1988, 2008), Searle (1991, 1992), Siewert (1998), Strawson (1994, 2004), and myself (Kriegel 2003a), among others.
37. There are several ways to characterize "phenomenal consciousness." Perhaps the most widely used one is in terms of "what it is like" (Nagel 1974): a mental state is phenomenally conscious iff there is something it is like for its subject to be in it. My own preference is to characterize phenomenal consciousness as the property for which there is (at least the appearance of) an explanatory gap: a mental state is phenomenally conscious iff we cannot entirely wrap our mind around the notion that all the state's properties are physical.
38. For present purposes, we can say that the "phenomenal character" of a mental state M is that property of M's in virtue of which M is the phenomenally conscious state it is and a phenomenally conscious state at all.
39. This formulation assumes that if a property F constitutes a property G, and F is non-relational, then it follows that G is non-relational. If this assumption is denied, we would have to strengthen condition (iii), so that the relation between M_i's being P and M_i's being I be such that the non-relationality of P would entail the non-relationality of I. Obviously, one appropriate relation is identity: if M_i's

being P is identical to M_i's being I, and P is a non-relational property of M_i, then I is a non-relational property of M_i.

40. Here too, the "in virtue of" relation should be understood very strongly, as implying at least something like a constitution or realization relation.
41. The phrase "phenomenal intentionality" was coined by Loar around the mid-nineties (Loar 2002), but was being independently used in the second half of the nineties by Horgan and Tienson (Horgan and Tienson 2002).
42. There may be good reason to believe that all *conscious* states are phenomenally intentional, but what we need here is only to deny that *no* conscious states are phenomenally intentional, that is, to claim that *some* conscious states are phenomenally intentional.
43. Adapted from Byrne (2001).
44. In this paper, I use "qualia" in the innocuous sense in which it means the same as "phenomenal properties." Sometimes the term is used to designate properties that by definition cannot be accounted for in terms of functional role or representational content. I am not using the term in this more restrictive sense.
45. The transparency claim was articulated clearly enough already by Armstrong (1968), and was probably aired for the first time by Moore (1903).
46. There are kinds of conscious state for which the transparency claim is less compelling. Thus, certain moods and emotional experiences often require a more theory-laden "spin" to align with transparency. (For such theoretically loaded treatment, see Armstrong (1968) and Tye (2000). Both argue that emotional experiences' phenomenal character consists in their representing bodily changes in their subject. Pre-theoretically, however, when one turns one's attention away from an enraging event or state of affairs and onto one's episode of rage regarding it, one does become aware of a feature that does not seem to belong to the enraging event or state of affairs.) But this should not matter for our present purposes, since the premise being defended is not that *all* phenomenal character is intentional, only that *some* is. We can grant for the sake of argument that certain moods and emotional experiences have phenomenal characters that do not make them intentional, as long as we insist that other experiences—such as visual perceptions—have phenomenal characters that do make them intentional. And that is hard to deny.
47. The claim is most naturally construed as universal: *all* phenomenal characters are non-relational (i.e., all properties of having a certain phenomenal character are non-relational properties of their bearers). But it is worth noting that we may also allow that only *some* phenomenal characters are non-relational, as long as there is overlap between the class of non-relational phenomenal characters and the class of intentional phenomenal characters. That is, what is essential to our present purposes is the following claim: some phenomenal characters that make the states that have them intentional are non-relational properties of those states; or: there are properties $F_1, \ldots, F_n$ of having a certain phenomenal character, such that $F_1, \ldots, F_n$ are both intentional and non-relational. Nonetheless, we will defend here the stronger, universal claim.
48. At least it is a non-relational property of the brain (or the brain's property of having experiences with these phenomenal characters is a non-relational property

of its). This is not yet to say that phenomenal character is a non-relational property of the *experiences*, since it leaves open the possibility that a mental state has phenomenal character in virtue of standing in a relation to another cerebral item in the subject. However, it is hard to see what the relevant cerebral item could be. In particular, there are familiar difficulties with the thought that a mental state could have phenomenal character in virtue of standing in certain relations to other mental states, as per functionalist accounts of phenomenal consciousness. These difficulties involve so-called inverted and absent qualia, as well as the simple fact that a mental state's functional role is a dispositional property, whereas phenomenal character is a manifest, categorical property.

49. If the claim is that the brain's state is indistinguishable from its duplicate's from a subjective first-person point of view (i.e., is indistinguishable *to the brain*), but that we should nonetheless count it as a different phenomenal state for broadly top-down theoretical reasons, then the response is that this would be a merely verbal move: there would still be a kind of "subjective intentionality" fully determined by what phenomenal state a conscious creature seems to itself to be in from a subjective first-person point of view, and that that subjective intentionality would be non-relational. If the claim is that the brain's phenomenal state is different from its duplicate's in a substantive, non-verbal way, such that the two states are subjectively distinguishable, some reason would have to be given for believing this, and for my part I cannot imagine what that reason might be.
50. In addition, it is arguable that phenomenal character must be in principle introspectible by the subject. However, relational properties of the subject are not introspectible by the subject, since the external relatum is not accessible to introspection. The subject may access that relatum, and therefore the relation, but not via introspection. Since the external relatum is external, any awareness of relations involving it would *ipso facto* be extrospective awareness. If so, phenomenal character cannot be relational, for it must be in principle introspectible, and the relational is not in principle introspectible. One might try to undermine this consideration by adapting a Burgean (1988) approach to self-knowledge that makes it constitutively related to external objects and apply that approach to introspection.
51. Loar (2002) explicates the notion of phenomenal directedness in terms of the distinction between *what* a phenomenal state represents and *how* it represents. If you look at a lemon while stretching your eyes, Loar says, it is not so accurate to say that your visual experience represents a blurry lemon as to say that it has a blurry way of representing a lemon. In this case the blurriness is clearly seen as a way the experience represents (the how of representation), but the same holds in all cases of phenomenal directedness: they all involve a way an experience represents rather than something to do with what the experience represents. This explication is in line with our adverbial account of phenomenal intentionality, inasmuch as talk of ways-of-representing is very much in line with talk of representing *x*-wise. On the other hand, the adverbial account does not work with a dualism of what- and how-representation. In a way, the whole point of the adverbial move was to reduce what-representation to how-representation. It was not to distinguish two sorts of representation, side by side, and somehow give precedence to one over the other. Loar's what/how distinction parallels the traditional Fregean

distinction between reference and sense (or mode of presentation), and his claim reads most naturally as stating the primacy of sense over reference. The adverbial account presented here, however, makes a different claim: that (phenomenal) reference is just as non-relational as (phenomenal) sense. Horgan and Tienson (2002) have another gloss on the notion of phenomenal directedness. For them, a conscious state's phenomenal directedness at a state of affairs is a matter of the state presenting apparent objects apparently instantiating apparent properties and apparently bearing apparent relations to each other. Presumably, the peppering of such a description with appearance terms is meant to convey that conscious experiences present a system very much like the system of objects, properties, and relations that is the external world, but without the ontological commitment to entities being presented. Thus it is worth noting that although "there is an x, such that x appears to instantiate the property of being F" entails "there is an x, such that x does instantiate the property of appearing to be F" (i.e., that the appearance of instantiating a property implies the instantiating of an appearance property), "there appears to be an x, such that x is F" does not entail "there is an x-appearance, such that the x-appearance is F" (i.e., the appearance of a particular existing does not entail the existence of an appearance particular). Talk of what is represented via phenomenal directedness in appearance terms may thus afford us a way of talking of what is represented without committing to any existents.

52. It may be worth stressing that phenomenal directedness is a good candidate for representing x-wise independently of the possibility that conscious representation is pictorial. For one thing, it is debatable whether phenomenal character is indeed pictorial. Thus some philosophers have held that no mental activity is ever pictorial, and the view that conscious states feature intrinsic phenomenal directedness should be open to them. For another, phenomenal character may not be restricted to mental states that might even initially plausibly be construed as pictorial. Thus some philosophers have claimed that cognitive and doxastic states can have a proprietary phenomenal character that outstrips any accompanying pictorial imagery (Horgan and Tienson 2002; Pitt 2004), and again these philosophers should be able to hold that such conscious states exhibit intrinsic phenomenal directedness.
53. For plausible accounts of narrow content, see Chalmers (2002, 2003) and Segal (2000).
54. It may be worth noting that there are hypothetical scenarios for which the other response—denying phenomenal indistinguishability—is appropriate, and yet other for which a combination of the two responses might be apt. Thus, consider an objection of the following form. Suppose that your department chair has been replaced by an alien duplicate; the phenomenal character of your dealings with your chair is the same, but their intentional content is different, at least if enough time has passed. In addressing *this* objection, it would be important to distinguish two cases: where you are aware that the replacement has taken place, and where you are not. In the first case, the correct response is to deny that the phenomenal character has remained the same. Perhaps the purely sensory component of your phenomenal character is unchanged, but the overall phenomenology has certainly changed: what it is like for you to deal with the

chair is not the same. This case is akin to the case of switching from visually experiencing a duck-rabbit figure as a duck and experiencing it as a rabbit: the sensory dimension of phenomenology is unchanged, but overall what it is like is different. (Indeed, there is something it is like to undergo the very event of aspect switch.) In the second case, the converse response is plausible: the phenomenal character is the same, but so is the intentional content. Even if your chair, Chris, has been replaced by an alien, Jimamba, as far as you are consciously concerned you are still dealing with Chris. It does not matter how long ago the replacement occurred. Twenty years after the replacement, you are still under the impression that you are dealing with Chris, never under the impression that you are dealing with Jimamba; and so your relevant conscious experiences represent Chris, not Jimamba. To be sure, they all *misrepresent*. But this is precisely because their representational content is unchanged, while the world is changed.

55. In one sense, this is to avoid genuinely singular (phenomenal) representations altogether. But we have to be clear on the sense in which this is so. What is rejected here is the possibility of *de re* phenomenal intentionality. But we still have representation of concrete particulars (as opposed to bundles of properties).
56. Moreover, while lack of compositionality as such may be an embarrassment for an adverbial account of representation in natural language, and perhaps something like a language of thought, it does not seem so embarrassing for an account of phenomenally conscious representation, since conscious representations do not constitute a syntactically well-organized, well-behaved representational system in the way a language does. It is hard to see what theoretical fallout there would be if it turns out that conscious representations do not exhibit compositional structure if we can after all perform the kind of inferences Jackson says adverbialism prevents. In particular, conscious representations do not have to be acquired as a system in the way linguistic representations do.
57. Note, however, that the above response to Jackson's objection does not depend on the peculiarities of phenomenal intentionality. The response would have worked just as fine for Chisholm and Ducasse. There are some kinds of intentionality for which compositional inferences seem more appropriate, but non-compositional inferences are also available in the case of all kinds of intentionality.
58. There may still be a skeptical worry about the outlook on representation that I am about to defend: if all representation derives from a kind of representation that is not grounded in a relation to external reality, then all representation may be ultimately disconnected from reality in a way that opens up a skeptical gap between reality and our representation thereof. However, it is doubly inaccurate that on the present outlook, conscious representation is not grounded in a relation to external reality. First of all, an adverbial account may allow many relations to reality, as long as these are *causal* rather than *constitutive*. Furthermore, there are also constitutive relations involved in *veridical representation*. When a representation is veridical, there are necessarily relations it bears to reality. All an adverbial account denies is that the very *existence* of a representation involves bearing relations to reality.
59. This distinction is drawn most clearly in Searle (1992) and Dretske (1988). But it is inspired mainly by Grice's early work on natural and non-natural meaning (Grice 1957).

60. Suppose a so-called natural sign S of an entity E cannot possibly be recognized by anyone to stand for E. S bears a certain systematic causal relation to E, but for whatever reason, the relation cannot possibly be uncovered by any sentient creature. It is natural to think of S as carrying information about E, but it would be strained to say that S *represents* E. For S to genuinely represent E, the fact that S carries information about E should be recognizable and usable by someone.
61. Suppose that, due to brain lesion, John is unable to recognize ducks. To him too, the figure represents a rabbit but not a duck. As it turns out, however, John's brain lesion incapacitated the ventral stream in his visual cortex (which produces conscious representations), but spared the dorsal stream (which produces non-conscious representations). (See Milner and Goodale 1995 for the hypothesis regarding the two informational pathways in the visual system.) As a result, John is fully capable of hosting dorsal-stream representations of ducks, in spite of his inability to have ventral-stream ones. These represent ducks, but not *to John*.
62. More accurately, conscious representations always represent to the representing agent or subject, whereas non-conscious ones rarely do, at least non-incidentally. Thus, tree rings may represent the tree's age to a botanist, but clearly, they do not represent the tree's age *to the tree*.
63. Several philosophers have held that all and only phenomenally conscious states are self-representing, or at least self-representing in the appropriate way (Caston 2002, Kriegel 2003b; see Kriegel and Williford 2006 for discussion). For these philosophers, it is plausible that only conscious representations are self-representing. When the view is that conscious representations represent themselves to have a certain representational content, it follows that conscious representations are self-interpreting.
64. For comments on a previous draft, I would like to thank Ben Blumson, David Chalmers, Terry Horgan, George Graham, Farid Masrour, and Kristie Miller, and for helpful conversations Ben Blumson, Mike Bruno, Jordi Fernández, and Kristie Miller. I have also benefited tremendously from presenting this paper at the University of Arizona and the University of Fribourg; I would like to thank the audiences there, especially Stephen Biggs, Brian Fiala, Marc Johansen, Martine Nida-Rümelin, Laura Schroeter, Houston Smit, Gianfranco Soldati, Juan Suarez, and Joseph Tolliver.

References

Armstrong, D. 1968. *A Materialist Theory of the Mind*. London: Routledge.
Burge, T. 1988. "Individualism and Self-Knowledge." *Journal of Philosophy* 85: 649–663.
Byrne, A. 2001. "Intentionalism Defended." *Philosophical Review* 110: 199–240.
Caston, V. 2002. "Aristotle on Consciousness." *Mind* 111: 751–815.
Chalmers, D.J. 2002. "The Components of Content." In D.J. Chalmers (ed.), *Philosophy of Mind: Classical and Contemporary Readings*. Oxford and New York: Oxford UP.
Chalmers, D.J. 2003. "The Nature of Narrow Content." *Philosophical Issues* 13: 46–66.
Chalmers, D.J. 2006. "Two-Dimensional Semantics." In Smith and E. LePore (eds.), *Oxford Handbook of the Philosophy of Language*. Oxford and New York: Oxford UP.
Chisholm, R. 1957. *Perceiving: A Philosophical Study*. Ithaca: Cornell UP.
Crane, T. 2001. "Intentional Objects." *Ratio* 14: 336–349.

Davidson, D. 1984. "Enquiries into Truth and Interpretation." Oxford: Clarendon Press.
Dennett, D.C. 1987. *The Intentional Stance*. Cambridge MA: MIT Press.
Dennett, D.C. 1990. "The Myth of Original Intentionality." In K.A. Mohyeldin Said, W.H. Newton-Smith, R. Viale, and K.V. Wilkes (eds.), *Modeling the Mind*. Oxford: Oxford UP.
Dretske, F.I. 1988. *Explaining Behavior*. Cambridge MA: MIT Press.
Dretske, F.I. 1996. "Phenomenal Externalism." *Philosophical Issues* 7: 143–159.
Ducasse, C.J. 1942. "Moore's Refutation of Idealism." In P.A. Schlipp (ed.), *The Philosophy of G.E. Moore*. La Salle IL: Open Court.
Fodor, J.A. 1975. *The Language of Thought*. Cambridge MA: Harvard UP.
Georgalis, N. 2006. *The Primacy of the Subjective*. Cambridge MA: MIT Press.
Goodman, N. 1976. *Languages of Art*. Indianapolis and Cambridge: Hackett Publishing Company.
Grice, H.P. 1957. "Meaning." *Philosophical Review* 66: 377–88.
Harman, G. 1990. "The Intrinsic Quality of Experience." *Philosophical Perspectives* 4: 31–52.
Horgan, T. and J. Tienson 2002. "The Intentionality of Phenomenology and the Phenomenology of Intentionality." In D.J. Chalmers (ed.), *Philosophy of Mind: Classical and Contemporary Readings*. Oxford and New York: Oxford UP.
Horgan, T., J. Tienson, and G. Graham 2004. "Phenomenal Intentionality and the Brain in a Vat." In R. Schantz (ed.), *The Externalist Challenge: New Studies on Cognition and Intentionality*. Amsterdam: de Gruyter.
Jackson, F. 1977. *Perception: A Representative Theory*. Cambridge: Cambridge UP.
Jackson, F. 1998. "Reference and Description Revisited." *Philosophical Perspectives* 12: 201–18.
Kriegel, U. 2003a. "Is Intentionality Dependent upon Consciousness?" *Philosophical Studies* 116: 271–307.
Kriegel, U. 2003b. "Consciousness as Intransitive Self-Consciousness: Two Views and an Argument." *Canadian Journal of Philosophy* 33: 103–132.
Kriegel, U. and K.W. Williford (eds.) 2006. *Self-Representational Approaches to Consciousness*. Cambridge MA: MIT Press.
Loar, B. 2002. "Phenomenal Intentionality as the Basis for Mental Content." In D.J. Chalmers (ed.), *Philosophy of Mind: Classical and Contemporary Readings*. Oxford and New York: Oxford UP.
Lycan, W.G. 2001. "The Case for Phenomenal Externalism." *Philosophical Perspectives* 15: 17–35.
McGinn, C. 1988. "Consciousness and Content." In *Proceedings of the British Academy* 76: 219–239. Reprinted in N.J. Block, O. Flanagan, and G.Güzeldere (eds.), *The Nature of Consciousness: Philosophical Debates*. Cambridge MA: MIT Press, 1997.
McGinn, C. 2008. "Consciousness as Knowingness." Forthcoming in *The Monist* 91.
Milner, A.D., and M.A. Goodale 1995. *The Visual Brain in Action*. Oxford: Oxford University Press.
Moore, G.E. 1903. "The Refutation of Idealism." *Mind* 12: 433–453.
Nagel, T. 1974. "What Is It Like to Be a Bat?" *Philosophical Review* 83: 435–450.
Pitt, D. 2004. "The Phenomenology of Cognition; or *What Is It Like to Think that P*?", *Philosophy and Phenomenological Research* 69: 1–36.
Putnam, H. 1975. "The Meaning of 'Meaning'." In his *Mind, Language, and Reality*. Cambridge: Cambridge UP.
Putnam, H. 1981. *Reason, Truth, and History*. Cambridge: Cambridge UP.
Quine, W.V.O. 1960. *Word and Object*. Cambridge MA: MIT Press
Searle, J.R. 1983. *Intentionality: An Essay in the Philosophy of Mind*. Cambridge: Cambridge UP.
Searle, J.R. 1991. "Consciousness, Unconsciousness, and Intentionality." *Philosophical Issues* 1: 45–66.

Searle, J.R. 1992. *The Rediscovery of Mind*. Cambridge MA: MIT Press.
Segal, G. 2000. *A Slim Book About Narrow Content*. Cambridge MA: MIT Press.
Siewert, C.P. 1998. *The Significance of Consciousness*. Princeton NJ: Princeton UP.
Strawson, G. 1994. *Mental Reality*. Cambridge MA: MIT Press.
Strawson, G. 2004. "Real Intentionality." *Phenomenology and the Cognitive Sciences* 3: 287–313.
Tye, M. 2000. *Consciousness, Color, and Content*. Cambridge MA: MIT Press.