

INTENTIONALITY AND NORMATIVITY

Uriah Kriegel
University of Arizona

Introduction/Abstract

One of the most enduring elements of Davidson's legacy is the idea that intentionality is inherently normative. The normativity of intentionality means different things to different people and in different contexts, however. A subsidiary goal of this paper is to get clear on the sense in which Davidson means the thesis that intentionality is inherently normative. The *central* goal of the paper is to consider whether the thesis is true, in light of recent work on intentionality that insists on an intimate connection between intentionality and phenomenal consciousness. According to several recent authors, there is a kind of intentionality—"phenomenal intentionality"—that is fully constituted by the phenomenal character of conscious experiences. I will argue that although Davidson's thesis, when correctly understood, is compelling for *most* intentionality, it is false of phenomenal intentionality. I start, in §1, with an explication of the notion of phenomenal intentionality; in §2, I elucidate Davidson's thesis and his case for it; in §3, I argue that the case does not extend to phenomenal intentionality; I close, in §4, with some objections and replies.

1. Psychological Intentionality and Phenomenal Intentionality

As treated in many mainstream discussions in the philosophy of mind, the phenomenon of intentionality is conceived as a theoretical posit of sorts, a property ascribed from the third-person perspective in the context of trying to explain and predict the behavior of persons and other intelligibly-behaving

systems. The paradigms of intentionality, so conceived, are the propositional attitudes, especially belief and desire.

This is not at all how intentionality was conceived by Brentano (1874) and his followers when they reintroduced the term into modern philosophy. Rather, it was conceived as a first-person phenomenon, a property we ascribe to conscious states because they present themselves to us in our personal experience as *directed* in the relevant sense. The paradigms of intentionality so conceived were typically not belief and desire, but manifestly intentional conscious states, such as perceptual experiences and conscious occurrent thoughts.

That the phenomenon is conceived now otherwise than then is hardly a condemnation of the current conception. Conceptions of phenomena are routinely updated through the development of inquiry. It may well be that the present conception is simply the superior framework for the study of intentionality. But it may also be that in the process of developing a new framework, something essential to the original phenomenon has been left out.¹ It may therefore be useful to introduce a *conceptual* distinction between two kinds of intentionality, corresponding to each conception—without supposing at the outset that this conceptual distinction corresponds to a “real difference” between two numerically distinct properties.

To distinguish the two, I propose that we use the labels *psychological intentionality* and *phenomenal intentionality*. This is supposed to follow the spirit of Chalmers’ (1996 Ch.1) distinction between two conceptions of mind, the psychological and the phenomenal. The psychological conception of mind characterizes mental phenomena third-personally in terms of their causal relations to the environment and to each other; the phenomenal conception characterizes them first-personally in terms of their subjective feel. Thus the former focuses on the *mechanical* dimension of mental life, the latter on its *experiential* dimension. It is generally true of mental terms, says Chalmers, that they lead a “double life” as psychological terms and as phenomenal terms. The term “pain” is a case in point: mental events can be classified as pain either because (roughly) they are caused by harmful stimulation and cause aversive reaction, or because they feel that particular unpleasant way—they hurt. To keep track of the difference between these two notions, it is useful to distinguish between psychological pain and phenomenal pain. My suggestion is that we do the same with intentionality. Psychological intentionality will be characterized third-personally in terms of a state’s causal or mechanical profile, while phenomenal intentionality will be characterized first-personally in terms of the state’s subjective or experiential feel.²

A mental state’s psychological intentionality (henceforth: *ps-intentionality*) is an abstraction from the state’s total long-armed causal role. A mental state’s total causal role is the set of all its causes and effects; its *long-armed* role includes *distal* causes and effects (see Harman 1987). Different accounts of *ps-intentionality* focus on different subsets of the state’s total

long-armed causal role. Broadly causal or informational accounts tend to focus on distal causes; teleological accounts tend to focus on both distal causes and distal effects; inferential role accounts focus on proximal causes and effects; verificationist accounts on proximal causes only; and so on.³ What all these accounts have in common is that they appeal to a portion of the state's (long-armed) causal profile as that in virtue of which the state has its ps-intentional properties.

A mental state's phenomenal intentionality (henceforth: ph-intentionality) is an abstraction from its total phenomenal character, the overall way it is like for its subject to be in it. It is a matter of considerable debate whether or not some of a mental state's phenomenal properties are non-intentional, but if some are, they would have to be abstracted from in homing in on the state's ph-intentional properties.⁴

The claim I am making here is that ph-intentionality is conceptually independent from ps-intentionality, which means that there is no conceptual entailment from facts about the long-armed causal role of a conscious state to facts about that state's phenomenal character. It may turn out, of course, that as a matter of fact, phenomenal properties *just are* psychological properties, or that what it is like for a subject to be in a conscious state is metaphysically necessitated by the state's total long-armed causal role. But this would not undermine the *conceptual* distinction between ph- and ps-intentionality: the concepts can be separate even if the property they pick out is the same.⁵

The claim of conceptual independence is thus compatible with lack of metaphysical independence. I focus on the former because it is the one that will matter for my purposes later in the paper. As we will see in §2, the alleged normativity of intentionality pertains (in the first instance) not to intentional *properties* but to intentional *ascriptions*. Since in the relevant sense ascriptions have concepts rather than properties as constituents, a conceptual distinction is all that will be needed to consider whether both ps- and ph-intentionality are inherently normative. Thus the only substantive claim we have to make here is that there *exist* both intentional properties mental states have in virtue of their *psychological* properties (long-armed causal role) and intentional properties mental states have in virtue of their *phenomenal* properties (what it is like). Whether the two are metaphysically distinct will not concern us.

That there exist intentional properties mental states have in virtue of their psychological properties is widely accepted. It is a more delicate matter whether there exist intentional properties mental states have in virtue of their phenomenal properties. Here I will assume that there does, for broadly the reasons adduced by Siewert (1998 Ch.7). It is a remarkable fact about a visual experience of a table, notes Siewert, that purely in virtue of its phenomenal character, and without need of interpretation, the experience is *assessable for accuracy*. Thus the experience has accuracy conditions purely in virtue of what it is like to undergo it. And just as truth conditions are

naturally thought of as constituting an intentional content for truth-apt mental states such as belief, so these accuracy conditions should be thought of as constituting an intentional content for accuracy-apt mental states such as perceptual experiences.

One question that arises immediately concerns the notion of phenomenal character at play here. It is no secret that “phenomenal character” is an essentially contested term in the philosophy of mind. But it seems to me that there are two main notions the term is used to express. The first is the notion of a sensuous quality. The second is the notion of a property for which there is an explanatory gap, or at least an appearance of one. Elsewhere, I argue that the latter notion is the more theoretically profitable, because more neutral (Kriegel 2009). Importantly, it is quite plausible that, in this second sense, there is a proprietary phenomenal character associated with conscious occurrent thoughts: a phenomenal character of thoughts that goes beyond the character of accompanying imagery and silent speech (see Pitt 2004; also Goldman 1993, Strawson 1994, Siewert 1998, and Horgan and Tienson 2002).

The notion of a kind of intentionality best appreciated from a subjective first-person perspective has attracted the interest of several recent authors. This kind of intentionality is variously referred to as “subjective intentionality” (Loar 1987), “conscious intentionality” (McGinn 1988, Searle 1992, Kriegel 2003, Georgalis 2006), “experiential intentionality” (Strawson 1994, 2004), and indeed “phenomenal intentionality” (Loar 2003, Horgan and Tienson 2002, Kriegel 2007). In light of the recent interest in this notion, it is intriguing to consider whether some of the more traditional claims sometimes made about intentionality apply also to phenomenal intentionality. For example, some of the above authors have claimed that although Putnam’s (1975) and Burge’s (1979) thought-experiments support an externalist view of (what I called) ps-intentionality, they do not apply to ph-intentionality, which is internalistic (Loar 2003, Horgan and Tienson 2002, Horgan, Tienson, and Graham 2004, Georgalis 2006, Kriegel 2007). In a similar vein, I want to consider here whether Davidsonian claims about the normativity of intentionality apply to ph-intentionality as much as to ps-intentionality. I will argue that they do not.

2. The Normativity of Intentionality

According to Davidson, ascription of propositional attitudes to persons is supposed to maximize the intelligibility of their overt behavior, and is therefore governed by a cluster of normative principles sometimes loosely referred to as “the principle of charity.”⁶ These include the principles that (by the interpreter’s own lights) persons’ beliefs are mostly true and coherent, that their desires are mostly good and mutually satisfiable (and/or suitably

prioritized), and that their beliefs and desires mostly constitute good reasons for action. As Davidson (1970: 97) puts it, in constructing a theory of someone's behavior, "we will try for a theory that finds him consistent, a believer of truths, and a lover of the good (all by our own lights, it goes without saying)."

Note that, so understood, the normativity of intentionality pertains in the first instance not to intentional *properties*, but to *ascriptions* of intentional properties. The Davidsonian case for the normativity of intentional ascription could be represented as the following piece of reasoning: 1) intentional ascription is governed by principles of charity, in the sense that a competent interpreter must use principles of charity in ascribing intentional states; 2) these principles are normative; therefore, 3) intentional ascription is governed by normative principles.⁷ In the bulk of this section, I want to remind the reader of the Davidsonian reason for embracing the first premise. Before doing so, let me say a few words about the second premise.

The term "normativity" is used in a sometimes bewildering variety of senses in the literature on meaning and content. It is thus important to get clear on the sense in which the principles of charity are "normative." Compare the following two statements:

- N1) Old friendships are rare and valuable.
- N2) Call your mom on Mother's Day.

N1 is normative in the sense that it employs a normative term ("valuable"). N2 does not employ any normative term, but it is recognizably normative in another sense, namely, that it is an imperative rather than an indicative—it sets a norm for us to follow. We may say that it involves a normative "force." I will call the normativity involved in N1 *evaluative* normativity and that involved in N2 *action-guiding* normativity. Superficially at least, these are two different kinds of normativity,⁸ though it is a controversial matter what the ultimate relationship between them is.⁹

The principles of charity are normative in the action-guiding sense. They are principles roughly of the form "Try to take subjects to have mostly true beliefs," "Try to take subjects to have mostly coherent beliefs," etc. These principles do not employ normative terms, but they are imperatives which impose norms on (and thus guide) the ascriptive activity of competent interpreters.¹⁰

Having clarified the sense in which the principles of charity are normative, let us consider Davidson's reason for holding that intentional ascription is governed by them. The reason is brought out nicely in this passage (Davidson 1974: 18):

If you see a ketch sailing by and your companion says, "Look at that handsome yawl," you may be faced with a problem of interpretation. One natural possibility

is that your friend has mistaken a ketch for a yawl, and has formed a false belief. But if his vision is good and his line of sight favorable it is even more plausible that he does not use the word “yawl” quite as you do, and has made no mistake at all about the position of the jigger on the passing yacht.

From which Davidson concludes (*Ibid.*; italics mine):

[I]f we merely know that someone holds a certain sentence to be true, we know neither what he means by the sentence nor what belief his holding it true represents. His holding the sentence true is thus *the vector of two forces*.

Confronted with a situation in which someone exclaims “look at this tiger!” while pointing at (what we know is) a pigeon, two coherent interpretations are open to us: (i) that this person believes a tiger is present, and takes the word “tiger” to express the concept of a tiger; (ii) that the person believes a pigeon is present, and takes the word “tiger” to express the concept of a pigeon.¹¹ Two points should be appreciated in this example. First, any competent interpreter would choose the second interpretation. Indeed, we may reasonably consider it a constraint on the competency of an interpreter that s/he choose the second interpretation in a situation such as this. Secondly, there is nothing in the data at the interpreter’s disposal (i.e., the exclamation) to favor either interpretation. What favors (ii) is only the principle of charity. In other words, the reason to hold that intentional ascription is governed by principles of charity—that is, that principles of charity must be operative in competent interpretation—is that without them interpretation is wildly underdetermined.

In the case just examined, we use charity to *revise* our initial hypothesis about how the person uses words. (The initial, default hypothesis is that by “tiger” she means tiger; the charitably revised hypothesis is that by “tiger” she means pigeon.) More often in everyday life, we use charity to *form* an initial hypothesis about how a person uses words. If an impressive car drives by and our interlocutor exclaims “that’s a fly ride right there!” we will form the hypothesis that by “fly ride” she means (is endeavoring to express the concept of) an impressive car. This too is based on ascribing to her the (true) belief that the car is impressive and the (good) desire to speak the truth.¹² Without assuming that she believes truly and desires well, the competent interpreter may virtually have nothing to go on in interpreting her utterance.

This observation applies not only to verbal behavior, but to all behavior. As Dennett (1971) notes, if we see a person opening her umbrella when it starts raining, we instantly ascribe to her the desire to stay dry and the belief that opening the umbrella will further that cause. But our data, or evidence, are fully consistent with ascribing to her a desire to get wet and a belief that opening the umbrella will help. What makes us ascribe to her the first belief-desire pair rather than the second has nothing to do with the

behavioral data, which are also the vector of two forces (as Davidson puts it in the passage above). It has to do rather with the fact that (by our lights) the first belief is true and the first desire good (in the relevant sense), while the second belief is false and the second desire bad.¹³ That is to say, what makes us ascribe to her the first belief-desire pair is charity.¹⁴

In fact, it is natural to treat the case of verbal behavior as just a special case. In the general case, the data for competent intentional ascription (i.e., that on the basis of which we ascribe intentional states) are the vectors of two intentional forces, a cognitive force (in the form of a belief or some other doxastic attitude) and a conative force (in the form of a desire or some other pro attitude). The verbal case is one in which the desire is to perform a certain linguistic act (e.g., express or communicate a belief that *p*) and the belief is that uttering certain words is likely to achieve that (e.g., manage to express the belief that *p*).¹⁵ Thus, under the interpretation we would naturally adopt, insofar as we are competent interpreters, the person who says “look at that tiger!” while pointing at a pigeon wants to give voice to her belief that a pigeon is present and believes that uttering “look at this tiger!” will achieve that.

The reason to take intentional ascription to be inherently normative, then, is what we may call the *vector-of-two-forces observation*. The observation is that the data upon which intentional ascriptions are based cannot typically decide between a number of competing ascriptions. In an ordinary situation calling for the ascription of an intentional state, there is typically one correct (competent) ascription to make, but a great many alternative ascriptions fully and equally consistent with the data at the interpreter’s disposal. That is, many very different intentional states can be ascribed consistently with the data, but typically only one can be ascribed *competently*.¹⁶ Davidson infers—very reasonably, it seems to me—that principles of charity must be operative in such a situation, and serve to rule out inappropriate ascriptions that are nonetheless consistent with the data.¹⁷

Note, however, that our discussion in this section was based entirely on the psychological conception of intentionality, with its focus on beliefs and desires and their explanatory relation to behavior, including verbal behavior. This is the only type of (mental) intentionality Davidson considered, and I followed him in the exposition of his reasoning concerning the normativity of its ascription. This reasoning constitutes a case for saying that *ps*-intentional ascription is inherently normative, that is, that the ascription of *ps*-intentional states must be governed by normative principles. In the next section, I argue that that case does not carry over to *ph*-intentional ascription. More precisely, I argue for the following thesis: For any intentional state *S*, if *S* is *ph*-intentional, then there is a competent interpreter *C*, such that possibly, *C* ascribes *S* without employing principles of charity (or any normative principles).

3. Intentional Ascription and Phenomenal Intentionality

Let us distinguish four kinds of intentional ascription (see Figure 1):

- (a) ascription of ps-intentional states to another creature (third-person ps-intentional ascription)
- (b) ascription of ps-intentional states to oneself (first-person ps-intentional ascription)
- (c) ascription of ph-intentional states to another (third-person ph-intentional ascription)
- (d) ascription of ph-intentional states to oneself (first-person ph-intentional ascription)

My main thesis in this section and the next is this: while (a)-(c) are based on data that are the vectors of two forces, (d) is based on data that are the vectors of single forces.

The importance of this thesis is clear: it would mean that the Davidsonian reason to take intentional ascription to be governed by principles of charity, and hence to be inherently normative, does not apply to first-person ph-intentional ascription. There might, of course, be an *independent* reason to think that first-person ph-intentional ascription is governed by principles of charity, or an independent reason to think that it is inherently normative. But pending such reasons, we should construe first-person ph-intentional ascription as unbound by principles of charity, and ultimately as lacking any normative dimension. This would mean that while *all* ps-intentional ascription is normative, *some* ph-intentional ascription is not. In other words, while *ps-intentionality* cannot be competently ascribed without employing normative principles (such as the principles of charity), *ph-intentionality can* be competently ascribed without employing such principles.

	ph-int	ps-int
1 st person	1ph intentional ascription	1ps intentional ascription
3 rd person	3ph intentional ascription	3ps intentional ascription

Figure 1. Four types of intentional ascription

But why should we believe that, unlike all other intentional ascription, first-person ph-intentional ascription is *not* based on data that are the vectors of two forces? The remainder of this section explains why. I want to examine more closely each of the four types of intentional ascription. In each case, I will consider simple examples of intentional ascription that bring out particularly crisply the central model for how the relevant type of ascription works.¹⁸ We will then see that the mechanics of first-person ph-intentional ascription are crucially different from those of other intentional ascription.¹⁹

Let us start with third-person ps-intentional ascription. Suppose a person says “it’s a nice day” and on that basis you ascribe to her the desire to express the belief that it is a nice day and the belief that just that concatenation of sounds will achieve that. In what sense are your data for this ascription the vectors of two forces? The only relevant datum here is the person’s utterance, her verbal behavior. On a natural model, others’ beliefs and desires are posited as behind-the-scene causes of behavior, including verbal. Ascriptions of beliefs and desires to others is based on causal inference: we make inferences about others’ beliefs and desires on the basis of observing behavior in the same way we generally make inferences about hidden causes on the basis of observed effects. Just as on the basis of observing smoke you infer that there is unobserved fire, so on the basis of observing the person’s utterance you infer that she has certain unobserved beliefs and desires.

What is important here is this: the utterance of “it’s a nice day” is the vector of two forces in that neither the belief by itself nor the desire by itself could causally explain it. Combined with a different desire, the same belief would cause a different utterance, and combined with a different belief, the same desire would cause a different utterance. For example, combined with the desire to mislead, the same belief would result in uttering something other than “it’s a nice day”; combined with the belief that ‘day’ means night and ‘night’ means day, the same desire would result in uttering “it’s a nice night.” It is only the conspiracy of a belief and a desire that can cause it, never a belief in isolation from any desire or a desire in isolation from any belief. Accordingly, the datum at your disposal offers no support for ascription of the belief in isolation from the desire or the desire in isolation from the belief.²⁰ More generally, it is impossible to use the datum at your disposal to ascribe to the person any *single* state; only pairs of states can be ascribed by way of explaining her verbal behavior. It is in this sense that your datum is the vector of two forces.

Contrast now the case of first-person ph-intentional ascription. Suppose you ascribe to yourself a visual experience as of a table. This kind of ascription works very differently. Here the basis on which you make this ascription, at least in central cases, does not consist in behavioral data of any sort. It is not as though you notice a certain behavior on your part and infer that you must be undergoing a visual experience as of a table. Perhaps there

are cases of first-person ph-intentional ascription that do work like this, but those are surely the pathological or highly unusual cases. In the central case, your ascription of a visual experience to yourself is not based on inference from your behavior. More generally, you do not observe putative causal effects of your experience and infer on their basis the existence of hidden causes. Rather, you seem to make the ascription on the basis of observing, in some admittedly elusive sense, the experience itself—observing, that is, the very state which you ascribe.²¹

Here the natural model is along the lines of what is sometimes referred to in the perceptual case as *endorsement*. Just as you judge that there is a chair in front of you because (a) it *perceptually* seems to you that there is and (b) you endorse this perceptual seeming, so you judge that you are having an experience as of a table because (a) it *introspectively* seems to you that you are and (b) you endorse this introspective seeming. In both cases, something *seems* to be the case and you endorse the seeming. In this kind of intentional ascription, your data are certain deliverances of introspection. Once the deliverances are endorsed, you ascribe a ph-intentional state to yourself—or perhaps the endorsing of such deliverances *constitutes* the relevant ascription. In any case, these deliverances are *not* the vectors of two forces: your experience of the table *can* yield the relevant deliverance all by itself and in isolation from any other mental state.²² Therefore, you *can* on the basis of the datum at your disposal ascribe to yourself an experience as of a table *without at the same time ascribing to yourself any other state*. More generally, you can on the basis of your datum ascribe to yourself a *single* mental state; no conspiracy with other mental states is needed to explain the datum.²³ The datum is thus the vector of a single force.²⁴

To repeat, this is the *central* model for how first-person ph-intentional ascription works. There may well be cases of such ascription that do not fit this model. But to establish the asymmetry between ps- and ph-intentional ascription in which I am interested, it is sufficient that *some* ph-intentional ascriptions are not based on data that are the vectors of two forces (given that *all* ps-intentional ascription is so based).²⁵

Since what we have been concerned to establish is that *some* ph-intentional ascription is based on data that are not the vectors of two forces, and not that *all* such ascription is, we need not worry overmuch about third-person ph-intentional ascription. For the asymmetry can be established regardless of how third-person ph-intentional ascription turns out to work.²⁶ Nonetheless, let me consider briefly this kind of ascription.

Suppose you ascribe to *another person* a visual experience as of a table. One way you might come to ascribe such an experience is by positing it as cause of table beliefs which, in conspiracy with conversation desires, cause this person's table conversation. In such a case, the ascription behaves essentially as third-person ps-intentional ascription does, and the data are clearly the vectors of two forces.²⁷ Another way you might come to ascribe

to another person a table experience is by observing the table itself and the position of the person's visual organs, and concluding, perhaps through simulation, that the person must be seeing the table. This kind of ascription may or may not behave completely differently, but in any case it seems to me to be even more immediately imbued with charity: you are assuming that the person's perception is *veridical*—and assuming so without evidence. (You could certainly produce evidence, but such evidence would consist in just more behavioral data, which would of course be the vectors of two forces.) It appears, then, that third-person ph-intentional ascription is either based on data that are the vectors of two forces or even more immediately governed by principles of charity. Either way, such ascription is inherently normative.

As stressed above, however, this is of no great consequence to our present concerns. In particular, it does not undermine the asymmetry between ps- and ph-intentional ascription. What *would* undermine the asymmetry is if first-person ps-intentional ascription turned out to work in the same way as first-person ph-intentional ascription.

Let us consider, then, the case of first-person ps-intentional ascription. Are its data the vectors of two forces or one? This kind of intentional ascription involves ascribing to oneself ps-intentional states, paradigmatically unconscious beliefs or desires. Importantly, phenomenally conscious beliefs and desires do not qualify as ps-intentional states, as I use the term, but as ph-intentional states.²⁸ So it is not with ascription of such states that we are concerned. We are only concerned with the ascription of phenomenally *unconscious* states. Such states seem to fall into three categories: (i) dispositional or tacit beliefs and desires; (ii) Freudian suppressed or repressed beliefs, desires, and emotions; (iii) occurrent sub-personal states typically posited in the context of cognitive-scientific explanations of behavior. Let us consider the central models for ascription of states in each category.

As regards dispositional beliefs and desires, however, my own view is that there are no such states. It is not true that a person *dispositionally believes* that Michael Jordan is not a three-headed kangaroo; what is true is that she is *disposed to believe* that. That is, there are no dispositional beliefs and desires, only dispositions to believe or desire. More generally, there are no dispositional intentional states, only dispositions to enter (non-dispositional) intentional states. The argument for this is somewhat complicated, but its essential point is this: any explanatory burden dispositional intentional states are called forth to meet can be equally met, and more parsimoniously, by dispositions to enter intentional states (see Manfredi 1993, Audi 1994).²⁹ (There may be an *intuitive* cost in denying the existence of dispositional states, but there is no *explanatory* cost.³⁰) Recall now that on the psychological conception of intentionality, intentionality is conceived of primarily as an explanatory posit. So unless an explanatory gain can be found in positing dispositional states, there is no reason to suppose that some ps-intentional states are dispositional. Accordingly, when we say that someone believes that

Michael Jordan is not a three-headed kangaroo, we are not ascribing to them an intentional state; rather, we are ascribing to them a disposition to enter an intentional state. So on this view, these are not intentional ascriptions after all and need not concern us here.^{31,32}

Consider now the case of Freudian states. When ascribing such states to ourselves, we typically rely on an array of indirect behavioral and introspective data. For example, you might infer that you must be unconsciously expecting bad news in a looming phone conversation from the fact that you have been postponing picking up the phone for a few hours already (a behavioral datum) and have been finding it difficult to concentrate during that time (an introspective datum). Inasmuch as you rely on behavioral data (not picking up the phone), the mechanics of ascription will be the same as in *third*-person ps-intentional ascription, so the data will be the vectors of two forces. Thus, expecting bad news would explain not picking up the phone only in conspiracy with a desire for *good* news. More interestingly, the same applies to the introspective datum: expecting bad news would explain the difficulty to concentrate only in conspiracy with a desire for good news. If you desired bad news, the expectation of bad news would not cause difficulty to concentrate. Thus even introspective data for first-person ps-intentional ascription are the vectors of two forces.

(How come the introspective data behave so differently in ph- versus ps-intentional ascription? The key difference is that in the former case, but not in the latter, the data pertain directly to that which is ascribed, in the sense that the relevant deliverances of introspection are that one seems to be in the very intentional state one ascribes to oneself. Because of this, first-person ph-intentional ascription involves nothing but endorsement of seemings, whereas first-person ps-intentional ascription involves a more substantial inference – essentially, a causal inference from broadly phenomenological effects to behind-the-scene ps-intentional causes.)

Finally, let us consider first-person ascription of non-phenomenal sub-personal intentional states. There are many examples of such states posited in cognitive science: Marr's (1982) 2.5D sketches, Milner and Goodale's (1995) dorsal-stream visual representations, blindsight states, and subliminal perceptions would be examples drawn just from vision science. It is worth noting that ascription of such states does not normally take place in the folk, since the folk are typically unaware that such states exist (in contradistinction with Freudian states, which have entered psychological lore long ago). In any case, since subjects have no introspective access to such states, they can ascribe such states to themselves only on the basis of behavioral data. If so, the model for first-person ascription of such states would be the same as the model for third-person ascription of ps-intentional states: it would be based on inference from observed behavioral effects to unobserved intentional causes. In both cases, then, the ascription is based on data that are the vectors of two forces.

I conclude that all first-person ascriptions of ps-intentional states are based on data that are the vectors of two forces. Given that the same holds of third-person ps-intentional ascription, I further conclude that all ps-intentional ascription is based on such data, and that therefore there is an important asymmetry between ps- and ph-intentionality. In fact, there are three different asymmetries here. In the first instance, there is the following asymmetry: while *all* ps-intentional ascription is based on data that are the vectors of two forces, *some* ph-intentional ascription is based on data that are the vectors of single forces. On its basis we can probably assert a second asymmetry: while *all* ps-intentional ascription is governed by principles of charity, *some* ph-intentional ascription is not (pending independent reasons to think that it is). And on this basis we can assert a third one: while *all* ps-intentional ascription is inherently normative (in the sense that its competent performance requires the employment of action-guiding principles), *some* ph-intentional ascription is not (again, pending independent reasons to think that it is). I offer this third asymmetry as the central thesis of this paper.

4. Comparisons and Objections

I will consider a number of objections momentarily, but first let me note that my response to claims about the normativity of intentional ascription is structurally similar to an early reaction due to McGinn (1977). McGinn's discussion is complex, but let me crush the subtleties and present what I take to be its upshot: ascription of *de dicto* intentionality may be based on data that are the vectors of two forces, but ascription of *de re* intentionality is not. Consider the person who utters the words "this balloon is yellow" while pointing at a yellow refrigerator. As interpreters, we may be concerned to ascribe to her either beliefs *de dicto* or beliefs *de re*. Suppose we wish to ascribe beliefs *de dicto*. Our data are consistent with two interpretations: (i) she believes that the refrigerator is yellow and by "balloon" she means refrigerator; (ii) she believes that a balloon is yellow and uses words as we do. But now suppose we wish to ascribe to this person beliefs *de re*. While the data are consistent with saying that the person believes, of a refrigerator, that it is yellow, they are *inconsistent* with saying that the person believes, of a balloon, that it is yellow. If there is no balloon, it is false that the person believes anything of a balloon.³³

This reaction is structurally similar to ours, in that it indicts the normativist about intentional ascription with failure to make certain distinctions, and subsequently failure to appreciate that the vector-of-two-forces observation applies only to some kinds of intentional ascription and not others.

However, there are important differences between this reaction and ours. One is that the *reasons de re* intentional ascription and ph-intentional ascription are not two-force affairs are very different. For my purposes,

however, the most important difference is that, at least as construed in Quine's (1956) seminal discussion, *de re* intentional ascription is extensional. This is in contrast with *ph*-intentional ascription, which is clearly intensional.³⁴ That is, the former does, but the latter does not, support existential generalization and truth-preserving substitution of co-referential terms. From "she believes, of the refrigerator, that it is yellow," we can infer both "there is something, such that she believes of it that it is yellow" and "she believes, of the only rectangular object in the house, that it is yellow." One *cannot* similarly infer, from "I am thinking that the refrigerator is yellow," either "there is something, such that I am thinking that it is yellow" or "I am thinking that the only rectangular object in the house is yellow."³⁵

This is significant, because it means that even if McGinn's reaction is well justified (as it strikes me), the normativist could insist that nonetheless all *intensional* intentional ascription is based on two-force data. It is widely thought that it is only with the appearance of intensionality that the most distinctive and most important features of intentionality come onto the scene.³⁶ So the fact that there is a type of extensional intentional ascription that evades the two-force predicament may be less surprising and less significant than the corresponding point about *ph*-intentional ascription, which is intensional. In any case, the following still holds: the only *intensional* intentional ascriptions based on data that are not the vectors of two forces are *ph*-intentional ascriptions.

I close with consideration of some objections. In the previous section, I have treated first-person *ph*-intentional ascription as based on the endorsement of introspective seemings, in which intentional states are in some sense "observed." It may be objected that this commits me to a perceptual model of introspection, and that such a model is hopeless. By "perceptual model," I mean a model of introspection that construes introspection as crucially analogous to perception. Typically, the crucial analogy would be that in both perception and introspection, there is a kind of direct epistemic contact with the perceived or introspected. Some philosophers have objected to such a perceptual model from the left, arguing that the relationship between introspection and the introspected is less intimate than that between perception and the perceived. Others have objected from the right, arguing that the relationship is *more* intimate. My response in both cases is that the perceptual model is *not* in fact crucial to the epistemological claims I made; those claims can be made while accommodating the objector's favored model of introspection.

The left-wing objector is the philosopher who holds that there is nothing at all like perceptual immediacy in introspection. Perhaps the main version of this is the view that introspection is always doxastic or intellectual: introspecting my table experience is simply *thinking* that I am having that experience, or that I *seem* to have it (see, e.g., Rosenthal 1993). However, we can readily accommodate this intellectualist view of introspection. Even if

introspecting an experience is just thinking about it, the relevant thought would be formed on the basis of (perhaps be inferred from) data that are the vectors of single forces: it is possible to think that I am having a table experience without at the same time thinking that I am in any other intentional state. This would still contrast with thoughts that ascribe *ps*-intentional states (whether to oneself or to others), since those thoughts are based on data that are the vectors of two forces.

Another version of the left-wing objection is the view that introspection is fundamentally directed at the world, rather than at experience, due to the so-called transparency of experience (Harman 1990). Thus, according to Dretske's (1995) "displaced perception" model of introspection, just as one hears that the postman has arrived by hearing the dog bark, so one introspects that one has a table experience by seeing the table. There is no direct epistemic contact with the postman, and none with the table experience. It is not immediately clear to me what this model implies for the mechanics of first-person *ph*-intentional ascription, but in any case I do not think the model can be successfully applied across the board. Perhaps it can be made to work as a model of our introspective access to perceptual experiences, but surely it is hopeless as a model of our introspective access to certain *imaginative* experiences. When a visual image of a smiling octopus pops up in my mind arbitrarily and involuntarily, I can introspect it, but no displaced perception can take place. It is not as though I become aware that I am having an experience as of a smiling octopus on the basis of being perceptually aware of an octopus—or even seeming to myself to be perceptually aware of an octopus. It might be claimed that I *pretend* to be perceptually aware of an octopus, but of course pretending is itself an imaginative exercise, so this account would be circular and vacuous.³⁷

Consider now the right-wing objector. This is the philosopher who holds that, unlike in perception, the relata of introspection are not "independent existences," i.e., not entities which may persist in the other's absence (see, e.g., Shoemaker 1994). The term "acquaintance" is sometimes used to denote such an epistemic relation. Again, however, as long as acquaintance with an experience is based on data that are the vectors of single forces, first-person *ph*-intentional ascription retains its distinctiveness. And indeed it seems to be: if there is such a thing as acquaintance (in this sense), then surely I can be acquainted with my table experience in isolation from any conative component; the conspiracy of a cognitive and conative element is not necessary to give rise to the state of acquaintance.

A related objection is that the model of intentional ascription as endorsement of seemings applies also to some forms of third-person ascription. It has sometimes been claimed that just by attending to facial expressions, for example, one can *see* that a friend is nervous or distressed about an upcoming meeting with a romantic prospect he has been talking about for a few days. It is natural to interpret such ascription as involving nothing

but endorsement of a *perceptual* seeming. This in turn might render the data of ascription “single-forced.” My response is that although this kind of intentional ascription is immediate in a way that may tempt us to assimilate it to first-person ph-intentional ascription, the sense of immediacy at play is actually different, and does not support such an assimilation. When we see that a friend is in distress, our ascription of distress to the friend is immediate in the sense of not being mediated by conscious inference, and instead being rather automatic. However, it is clearly not immediate in the sense that we observe the distress itself. For we do not observe the distress itself, only its facial manifestation. Accordingly, the facial manifestation is the vector of two forces: it indicates distress only against the background of attribution to the friend of a strong desire that the meeting with the romantic prospect be successful. If one knew for a fact that the friend was utterly uninterested in the person he is about to meet, one would not be in a position to (competently) ascribe distress or nervousness to the friend on the basis of the relevant facial expression.

There are probably other objections that can be raised to the argument of the previous section. But it strikes me that the argument ultimately rests on a relatively straightforward observation, namely, that there is an asymmetry between first-person ph-intentional ascription and other kinds of intentional ascription, inasmuch as the data for intentional ascription are crucially different, and differently culled. What I claim is that the data of first-person ph-intentional ascription alone are the vectors of single forces, and that therefore ph-intentional ascription is not inherently normative in the way ps-intentional ascription is.

Conclusion

The central thesis of this paper is that while ascription of psychological intentionality is inherently normative (in the sense that action-guiding principles are operative in its competent execution), because it is always based on data that are the vectors of two forces, ascription of phenomenal intentionality is *not* inherently normative (in that sense), because it is *not* always based on such data. If this is right, then the Davidsonian claim that intentionality is inherently normative is not correct across the board. It is true at most of (what I have called) psychological intentionality. There is nothing normative about phenomenal intentionality, whose ascription requires no appeal to charity.³⁸

Notes

1. This is especially likely when vaguely ideological pressures are exerted from the top down, as was plausibly the case in pursuit of naturalization of intentionality.

2. The cases of pain and intentionality are not analogous, of course, but the psychological/phenomenal distinction is useful for both (I claim).
3. There is also, of course, long-armed functional role semantics, which appeals to proximal and distal causes and effects (as in Harman 1987).
4. If they are not, then such abstraction would not be necessary. We can use the term “abstraction” as we use the term “part,” so that there is both proper and improper abstraction. Then if no phenomenal properties are non-intentional, reaching the ph-intentionality of a state would involve improper abstraction from the total phenomenal character of the state.
5. Whatever the metaphysical connections between psychological and phenomenal properties, the former are certainly not conceptually sufficient for the latter: there is no conceptual analysis of the phenomenal in terms of the psychological. Thus, it is relatively easy to conceive of worlds with ps-intentionality but no ph-intentionality and vice versa. We can imagine a world with zombies psychologically indistinguishable (though perhaps physically very different) from us, and they would presumably have ps-intentional states but no ph-intentional states; and we can imagine a world in which evil-demon-deceived disembodied souls have rich mental lives phenomenally indistinguishable from ours, and they would presumably have ph-intentional states but not ps-intentional ones.
6. I use the term “overt behavior” to refer to publicly observable, third-personally accessible behavior, i.e., bodily behavior. There is a sense of “behavior” in which many mental acts—such as calculating, deliberating, etc.—counts as behavior, namely, the sense in which calculating is something one does, not something that happens to one. It is these mental acts which are supposed to be excluded by the notion of *overt* behavior. We may treat them as *covert* behavior.
7. Note that this argument does not make any stronger claim about any essential normativity of mental phenomena. In particular, it does not conclude that the individuation of mental states appeals to normativity. It is thus not in tension with Schroeder’s (2003) claim that interpretations of Davidson’s theory of mind as normative rest on a confusion.
8. Accordingly, they contrast with two different ways a sentence could be *descriptive*: by employing no normative terms or by employing an indicative grammar.
9. There are different views about the ultimate relationship between these two kinds of normativity. According to some kinds of ethical internalism, evaluative normativity implicates action-guiding normativity: although N1 is an indicative, it entails (in some sense) certain imperatives, e.g., “Make an effort to keep old friendships.” We need not take a stand on the plausibility of internalism.
10. Being normative in the action-guiding sense, they are non-descriptive, in the sense that they are not in the business of “getting things right.” They do not “aim at the truth,” but rather guide the subject’s ascriptions. They have a world-to-mind direction of fit, not a mind-to-world one. Intentional ascription most certainly involves also descriptive principles (e.g., “beliefs that *p* tend to cause such-and-such actions”), but the fact that it involves non-descriptive principles as well means that the practice of intentional ascription is in part not responsive to the way the world is—precisely in the sense that it is, in part, not in the business of “getting things right.”

11. By saying that the person takes the word “tiger” to express the concepts of tiger and pigeon, I mean that she takes the word to express the concept we take the words “tiger” and “pigeon” to express.
12. The relevant belief-desire pair may actually be slightly different. I will later suggest that the relevant pair comprises the desire to express (or perhaps communicate) the belief that the car is impressive and the belief that uttering “that’s a fly ride right there” would achieve that.
13. A quick comment on the goodness of the desire: the desire is good not in the moral sense, but (we may say) in the prudential sense. In that sense, a desire is good when it is a desire for something that is good for the desirer. Compare the robber’s desire to get out of the bank as quickly as possible. This is a prudentially good desire but not a morally good desire.
14. There are also descriptive principles of ascription at play here, which favor both these options over, say, ascribing to the person the belief that the moons of Jupiter are nice and the desire to fly backwards. Descriptive principles ensure that this assignment is inferior, on this occasion, to the options considered in the text. But the options in the text are discriminated by normative principles, not descriptive ones.
15. Here too I follow Davidson (1963), now in his account of reasons for action as combinations of a desire that p obtain and a belief that ϕ -ing would make p more likely.
16. There may certainly be cases where there are several possible ascriptions that a competent interpreter might choose among while remaining competent. But those are bound to be atypical cases.
17. Davidson’s reasoning here may be represented as an argument by inference to the best explanation. What needs explaining is that competent interpreters manage, *qua* competent interpreters, to converge on a single interpretation (and can recognize that that is the correct interpretation) in most typical situations, even though the data at their disposal grossly underdetermine interpretation. The best explanation of this is that competent interpreters, *qua* competent interpreters, employ principles of charity that rule out possible interpretations otherwise consistent with the data. Alternative explanations would cite some other device(s) competent interpreters might employ, in their capacity as competent interpreters, to rule out those interpretations. But it is not clear what other device(s) there might be, so no plausible alternatives present themselves to the Davidsonian explanation that competent interpretation must employ, and is to that extent governed by, principles of charity. (Needless to say, all this applies only to *competent* interpretation. Faced with the pigeon-pointer who exclaims “this is a tiger,” we can, if we want, interpret her as believing that she is looking at a tiger. We can also interpret her, if we very much want, as believing that she is looking at a three-headed kangaroo. What we cannot do, however, is interpret her in this way *competently*.)
18. There certainly exist much more complex instances of intentional ascription. But my contention is that the complexity in those cases does not change the basic mechanics of ascription, though it may well serve to becloud those mechanics. This is a reason to focus on simple examples in which the mechanics come through clearly.

19. I start by contrasting first-person ph-intentional ascription with third-person ps-intentional ascription, and will then remark on the case of third-person ph-intentional ascription. I will end with the trickiest kind of ascription, case of first-person ps-intentional ascription, which I claim to be on a par with third-person ps-intentional ascription. As just noted, in every case I will construct a toy model of how the intentional ascription is formed. The models target artificially simplified and “cleaned up” situations, but they should apply to the messier reality just as well—it is just that in the messy reality their distinctive character are harder to discern.
20. There are probably different desires that could be ascribed together with the same belief, and different beliefs that could be ascribed together with the same desire, but it is impossible to ascribe the belief in isolation from *any* desire or the desire in isolation from *any* belief.
21. It is not a straightforward matter to elucidate the sense in which the epistemic relation you bear to your experience can be described as “observation.” What I have in mind is the thought that there is an epistemic relation we sometimes bear to our own mental states that is analogous in some respects to sensory perception of external entities. There is a difference between believing that it is raining on the basis of the weather report and believing this on the basis of *seeing* the rain. The latter case involves a kind of direct contact with the state of affairs believed to obtain that the former case does not. In an analogous way, there is a difference between the way you know that what you are visualizing right now is a smiling octopus and the way I know this. I know it on the basis of testimony, you know it on the basis of direct contact with the believed state of affairs. It is this kind of contact that I think would not be misleadingly described as “observation.”
22. I use the term “yields” in a way that is neutral between a causal and a constitutive reading, to accommodate different views of the nature of introspection. Later in the section I will address the objection that the perceptual model of introspection is misguided.
23. This formulation is useful for some purposes, but also has the disadvantage that it may invite the misguided objection that a perceptual experience is itself a combination of two states, say a sensation and belief, and so first-person ph-intentional ascription also ascribes only two states at once. Clearly, however, even if we accept this view of perceptual experiences as combinations of two states, there are differences between the reason we can only ascribe two states at once in this case and the reason in the case of third-person ps-intentional ascription. Also, the fact (if it were a fact) that we can only ascribe two states at once in first-person ph-intentional ascription would not entail that there are competing interpretations fully consistent with the data, as is the case with third-person ps-intentional ascription, and therefore it would not entail that principles of charity must be operative in choosing among these ascriptions.
24. It is important to note that, in saying this, we are not indulging in any mystery-mongering about first-person access. Introspection has been claimed sometimes to be infallible, such that if it delivers that *p*, then *p*; sometimes to be incorrigible, such that if it delivers that *p*, then the deliverance is incontrovertibly justified; sometimes to be self-intimating, such that if *p*, then it delivers that *p*; sometimes to be immune to error of a certain type (e.g., through misidentification), such that

if it delivers that p , and certain conditions are met, then p ; and so on. But none of this follows simply from the assertion that there exists introspective access to one's conscious experience of the sort that clearly does not exist to another's mental states. As it happens, I happen to believe that a strong kind of epistemic privilege does attach to the deliverances of introspection, and have argued so elsewhere (Horgan and Kriegel 2007), but that view is logically independent of the claims made here.

25. Since there are central cases that clearly fit a model in which data for ph-intentional ascription are indeed *not* vectors of two forces, this asymmetry arises. A non-mental analogy may be useful. Suppose you are presented with a color patch projected on a window pane. You are told that the color is produced by projection from two different light sources, one from the left and one from the right. You can see the patch, but not the light sources. Thus, when you see a purple patch on the pane, your data are compatible with at least three hypotheses: that the left light is blue and the right one red; that the left light is red and the right one blue; that both are purple. Since the color of the patch is the vector of two forces, and you do not have direct access to those forces themselves, ascribing colors to the light sources must involve an element of *decision*. (In the present case, that the decision must be arbitrary; in the case of ps-intentional ascription, it is based on charity.) Compare now a case in which you *can* see not only the patch, but also the light sources—at least *sometimes*. Here you do have direct access to the lights, and so ascribing colors to them is not the vector of two forces. This is what happens with some ph-intentional ascription.
26. If such ascription turns out to be based on two-force data, there is still an asymmetry in virtue of the fact that first-person ph-intentional ascription is centrally different. If third-person ph-intentional ascription turns out to be also not based on two-force data, then it becomes more plausible to envisage a case for saying that not only *some*, but *all*, ph-intentional ascription works in this way. That would be a sharper asymmetry, but not a significantly more important one. (So: if it is true that both kinds of ph-intentional ascription are based on data that are the vectors of single forces, then *all* ph-intentional is not inherently normative; if it is only true that one kind of ph-intentional ascription is not normative, then *some* ph-intentional ascription is not normative. The latter already amounts to a disanalogy between ps- and ph-intentionality.)
27. The causal structure of the hidden mechanism is a little different when the ascribed state is not a belief, but the kind of perceptual state that causes belief. In the former case, the envisaged mechanism involves two states, the belief and the desire. In the latter, it involves three states, the belief, the desire, and the perceptual state that causes the belief. In both cases, however, it is impossible to ascribe a single, non-conspiratory state.
28. I am assuming here that a belief can be phenomenally conscious, which is controversial. The issue cannot be pursued here seriously, but arguments for the existence of phenomenally conscious belief abound in the literature (see Strawson 1994, Peacocke 1998, Horgan and Tienson 2002, Pitt 2004). It has sometimes been claimed on conceptual-analysis grounds that beliefs cannot be conscious and occurrent (Crane 2001). I am somewhat sympathetic to this claim, understood as a claim about how the English word “belief” is normally used,

but it is clear that there is no substantive issue at stake here, inasmuch as a tacit belief could still transform into a ph-intentional state; it is just that that state would not qualify as a belief the way the word is normally used. A natural term to pick out the kind of state a tacit belief transforms into when no longer tacit is “a thought.”

29. There may also be another argument to the same effect, in which the central claim is that there is something incoherent about the notion of a dispositional state. Any respectable ontological assay of states would cast them as occurrent entities. To say that a *state* is dispositional is really a confusion. This argument may be harder to prosecute, but my hunch is that it is fundamentally sound.
30. It is a fair question whether there is indeed an intuitive cost. On the one hand, it does sound counter-intuitive to deny that most people want (right now) to stay alive. But on the other hand, it is not counter-intuitive to deny, and is in fact counter-intuitive to assert, that most Americans believe (right now) that Michael Jordan is not a three-headed kangaroo. Overall, I do think there is an intuitive cost here, and recognize this cost to be a liability on the view that there are no dispositional states. Nonetheless, I think the cost is worth paying for the sake of the economy in states posited.
31. I take this eliminativism about dispositional states to cover both tacit beliefs that were never occurrent and those that were once occurrent and then were stored in long-term memory. It may seem that with the latter type of states, there is more reason to be realist about them. But again there is no explanatory gain in preferring stored dispositional states over stored dispositions to enter states. (The only new pressure presented by stored states is that the intuitive cost of denying their existence seems greater.) In the case of beliefs stored in short-term memory, there seems to be genuine explanatory gain in admitting their existence, but this is mainly because such states are most certainly occurrent (short-term memory being what it is).
32. Furthermore, even if there existed dispositional states, their first-person attribution would certainly work in a very different way from the way the first-person attribution of ph-intentional states works. For dispositional states are certainly not observed. True, just as when I ask you whether you are having a table experience, you reply immediately that you do, so when I ask you whether you believe that there are more than four countries in Africa, you reply immediately that you do; and the immediacy of reply might suggest direct observation in both cases. But the idea that we observe our tacit beliefs is really incredible. Evans (1982) points out that when someone asks you whether you believe that there will be a third world war, you do not start monitoring your internal states in search of an item that fits the description “belief that there will be a third world war”; rather, you consider whether there will in fact be a third world war. You attend not to psychological facts, but to geopolitical ones. That is, you seek evidence not for your believing the proposition, but for the proposition itself. Clearly, then, the immediacy of reply is not explained by the fact that you observe your belief. My guess is that the right explanation of the immediacy is that when the proposition you are asked whether you believe is obviously true, in that the evidence you have for it is overwhelming and readily available, you will answer the question immediately because the question is so easy. The explanation for this immediacy is

therefore not that you have observed the believing, but rather that the proposition believed is obvious (again, in the sense that the evidence for it overwhelming and readily available). Thus there is no argument from immediacy-of-reply to direct observation.

33. It is worth noting that this point applies not only to the oddities to do with terms that appear in the subject ascription of utterances that express beliefs. If the person utters “this refrigerator is purple,” this is consistent both with saying that the person believes that the refrigerator is purple and with saying that the person believes that the refrigerator is yellow. But while it is also consistent with saying that the person believes, of the yellow instance, that it is a refrigerator, it is inconsistent with saying that the person believes, of the purple instance, that it is a refrigerator.
34. It is a matter of some debate whether *de re* intentionality should be characterized as intentionality ascribed extensionally, as in Quine’s seminal discussion. Burge (1977), for example, offers a more complex characterization of *de re* intentionality that makes some such intentionality intensional. The issues that arise in this connection are tremendously complicated and I will not go into them here, instead dogmatically adopting the seminal Quinean characterization. If that characterization is misguided in the ways claimed by Burge, then McGinn’s reaction to the Davidsonian position is not only structurally similar to ours but also of equal importance.
35. Likewise, one cannot infer, from “I have an experience as of a yellow laptop,” either “there is something, such that I have an experience as of it being yellow” or “I have an experience as of my most valuable possession.” I am using the “as of” locution in these formulations because it sounds more obviously intensional. Nothing substantial rides on this: the properties of the visual experience itself that make some reports of it intensional are there regardless of the language we use in the reports. We can therefore use whatever language will bring out those features out in reporting the experiences.
36. In fact, on some views—including in Chisholm’s (1957) original treatment of intentionality—intensionality is the mark of intentionality. In Chisholm’s discussion, a property is intentional just in case it is picked out by an intensional verb. On this view, if *de re* intentional constructions are indeed extensional, as Quine maintains, then the properties picked out by those verbs are not intentional properties.
37. How to do justice to the transparency observation (which I accept) given this fact is a good question, but the displaced perception model cannot be right. For my part, I am persuaded that experience is transparent in the sense that its vehicular properties are not available to direct introspection. But it seems to me not to follow that experience is itself unavailable to direct introspection. On the contrary, an experience is directly introspectible due to the fact that its content properties are directly introspectible.
38. For comments on a previous draft, I would like to thank David Chalmers and Farid Masrouf, and for useful conversations David Bourget, Jordi Fernández, and Declan Smithies. I have also benefited from presenting the paper to audiences at La Trobe University, University of Arizona, University of California at Riverside, and at a workshop on phenomenology and intentionality at the University of

Modena and Reggio Emilia. I am indebted to those audiences, in particular John Bigelow, Frank Jackson, Marc Johansen, Benji Kozuch, Michelle Montague, Jack Reynolds, Eric Schwitzgebel, Charles Siewert, Alberto Voltolini, and probably a number of other people I am forgetting.

References

- Audi, R. 1994. "Dispositional Beliefs and Dispositions to Believe." *Noûs* 28: 419–434.
- Brentano, F. 1874. *Psychology from Empirical Standpoint*. Edited by O. Kraus. English edition L. L. McAlister. Translated by A. C. Rancurello, D. B. Terrell, and L. L. McAlister. London: Routledge and Kegan Paul, 1973.
- Burge, T. 1977. "Belief *De Re*." *Journal of Philosophy* 74: 338–362.
- Chalmers, D.J. 1996. *The Conscious Mind*. Oxford and New York: Oxford UP.
- Chisholm, R. 1957. *Perceiving: A Philosophical Study*. Ithaca, NY: Cornell UP.
- Crane, T. 2001. *Elements of Mind*. Oxford: Oxford UP.
- Davidson, D. 1963. "Actions, Reasons, and Causes." *Journal of Philosophy* 60: 685–700.
- Davidson, D. 1970. "Mental Events." In L. Foster and J. Swanson (eds.), *Experience and Theory*. London: Duckworth
- Davidson, D. 1974. "On the Very Idea of a Conceptual Scheme." *Proceedings and Address of the American Philosophical Association* 47: 5–20.
- Dennett, D.C. 1971. "Intentional Systems." *Journal of Philosophy* 68: 87–106.
- Dretske, F.I. 1995. *Naturalizing the Mind*. Cambridge MA: MIT Press.
- Evans, G. 1982. *The Varieties of Reference* (edited by J. McDowell). Oxford: Oxford UP.
- Georgalis, N. 2006. *The Primacy of the Subjective*. Cambridge MA: MIT Press.
- Goldman, A. 1993. "The Psychology of Folk Psychology." *Behavioral and Brain Sciences* 16: 15–28.
- Harman, G. 1987. "(Non-Solipsistic) Conceptual Role Semantics." In E. Lepore (ed.), *New Directions in Semantics*. London: Academic Press.
- Harman, G. 1990. "The Intrinsic Quality of Experience." *Philosophical Perspectives* 4: 31–52.
- Horgan, T. and U. Kriegel 2007. "Phenomenal Epistemology: What is Consciousness that We Should Know It so Well?" *Philosophical Issues* 17: 123–144.
- 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.
- Kriegel, U. 2003. "Is Intentionality Dependent upon Consciousness?" *Philosophical Studies* 116: 271–307.
- Kriegel, U. 2007. "Intentional Inexistence and Phenomenal Intentionality." *Philosophical Perspectives* 21: 307–340.
- Kriegel, U. 2009. *Subjective Consciousness: A Self-Representational Theory*. Oxford and New York: Oxford University Press.
- Loar, B. 1987. "Subjective Intentionality." *Philosophical Topics* 15: 89–124.
- Loar, B. 1995. "Reference from the First-Person Perspective." *Philosophical Issues* 6: 53–72.
- Loar, B. 2003. "Phenomenal Intentionality as the Basis for Mental Content." In M. Hahn and B. Ramberg (eds.), *Reflections and Replies: Essays on the Philosophy of Tyler Burge*. Cambridge: MIT Press.
- McGinn, C. 1977. "Charity, Interpretation, and Belief." *Journal of Philosophy* 74: 521–535.

- 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.
- Manfredi, P.A. 1993. "Tacit Beliefs and Other Doxastic Attitudes." *Philosophia* 22: 95–117.
- Marr, D. 1982. *Vision*. San Francisco: WH Freeman Publishers.
- Milner, A.D., and M.A. Goodale 1995. *The Visual Brain in Action*. Oxford: Oxford University Press.
- Peacocke, C. 1998. "Conscious Attitudes, Attention, and Self-Knowledge." In C. Wright, B.C. Smith, and C. Macdonald (eds.), *Knowing Our Own Minds*. Oxford: Oxford University Press.
- Pitt, D. 2004. "The Phenomenology of Cognition; or *What Is It Like to Think that P?*", *Philosophy and Phenomenological Research* 69: 1–36.
- Quine, W.V.O. 1956. "Quantifiers and Propositional Attitudes." *Journal of Philosophy* 53: 177–187.
- Rosenthal, D.M. 1993. "Thinking that One Thinks." In M. Davies and G. W. Humphreys (eds.), *Consciousness: Psychological and Philosophical Essays*. Oxford: Blackwell.
- Schroeder, T. 2003. "Donald Davidson's Theory of Mind is Non-Normative." *Philosophers' Imprint* 3.
- Searle, J.R. 1992. *The Rediscovery of Mind*. Cambridge MA: MIT Press.
- Siewert, C.P. 1998. *The Significance of Consciousness*. Princeton NJ: Princeton UP.
- Shoemaker, S. 1994. "Self-knowledge and 'Inner Sense.' Lecture II: The Broad Perceptual Model." *Philosophy and Phenomenological Research* 54: 271–290.
- Strawson, G. 1994. *Mental Reality*. Cambridge MA: MIT Press.
- Strawson, G. 2004. "Real Intentionality." *Phenomenology and the Cognitive Sciences* 3: 287–313.