

Not in the Mood for Intentionalism

Davide Bordini

Centre for Philosophy of Time

Philosophy Department

University of Milan

davide.bordini@gmail.com

Published in *Midwest Studies in Philosophy*, Volume 41, “Phenomenology of Affective Life”

[Penultimate draft, July, 2017]*

Abstract. According to intentionalism, the phenomenal character of experience is one and the same as the intentional content of experience (e.g., Dretske, 1995, 1996; Tye, 1995, 2000). This view has a problem with moods (anxiety, depression, elation, irritation, gloominess, grumpiness, etc.). Mood experiences certainly have phenomenal character, but do not exhibit directedness, i.e., do not appear intentional. Standardly, intentionalists have re-described moods’ undirectedness in terms of directedness towards everything or the whole world (e.g., Crane, 1998; Seager, 1999). This move offers the intentionalist a way out, but is quite unsatisfying. More recently, Angela Mendelovici (2013a, b) has suggested something that looks more interesting and promising: instead of re-describing moods’ phenomenology, she accepts its undirectedness at face value and tries to explain it in intentionalist terms. In this paper, I focus on and criticize Mendelovici’s proposal. As I will show, despite its *prima facie* virtues, the view is poorly motivated. For, contrary to what Mendelovici argues, introspection does not support her proposal—arguably, it provides some evidence *against* it. So, the problem that intentionalism has with moods is not solved, but is still there.

Keywords: moods, emotions, intentionalism, transparency of experience, phenomenal character, intentionality.

1 Introduction

Experiences are phenomenally conscious states: there is something it is like to be in them. According to *intentionalism*, at a minimum, phenomenal character is entirely dependent (in one way or another) on what an experience *represents* or *is of* (or *about*), i.e. *intentional content*. Suppose I am now looking at my brown desk and am having a visual experience of it. On intentionalism, the fact that there is something it is like to undergo such a visual experience is *nothing over and above* the fact that such an

* Please, quote the published version available here <http://onlinelibrary.wiley.com/doi/10.1111/misp.12066/full>.

experience represents certain properties—e.g., being brown or having a certain shape—as belonging to a certain object, the desk in our case. The most widespread version of intentionalism is the identity-version:

(INT) For every experience E, E’s phenomenal character is one and the same as E’s intentional content.¹

In this paper, I will use “intentionalism” to refer to this view.²

So understood, intentionalism provides an answer to the question of what phenomenal character *is*: phenomenal character is essentially representation. Such an answer appealingly paves the way to a physicalistic reduction of phenomenal character: a naturalistic account of mental representation (e.g., Dretske, 1981, 1988; Fodor, 1990; Millikan, 1984, 1989) would suffice to get a naturalistic account of phenomenal character.

A major problem for intentionalism, however, comes from moods: experiences like anxiety, depression, elation, irritation, gloominess, grumpiness, etc. Usually, experiences that *are* intentional also *exhibit* their intentionality, their being contentful, in their phenomenal character. For example, when I look at my brown desk, my visual experience *is an experience of* the brown desk, but it also *feels like* an experience of a brown desk. I call *phenomenal directedness*, or simply *directedness*, this intentionality-feeling that seems inherently involved in what it’s like to undergo an intentional experience. So, my experience of the brown desk exhibits phenomenal directedness. In a slogan, if intentionality is the of-ness of experiences, directedness is its *as of-ness*.

Now, typically, moods seem to lack such as of-ness: they *feel* entirely *undirected*. The phenomenal qualities involved in depression, for example, do not appear to belong to any particular object in the world or to be localized in any specific point in space (or in one’s body). Indeed, when one is depressed, it is difficult for one to focus and tell exactly *what* one is depressed about. Quite often, one’s depression does not appear to be about anything at all. So, not only does the experience not feel like an experience of something as being a certain way, it also (more radically) does not feel like an experience of anything at all.

Such undirectedness leaves room for the denial that moods have intentional content (Searle 1983; Deonna & Teroni 2012). If so, given (INT), intentionalism would clearly turn out to be false. Moreover,

¹ Dretske (1995, 1996); Harman (1990); Tye (1995, 2000).

² I suspect, however, that my point applies to weaker versions of intentionalism too, *mutatis mutandis*.

granting the intentionality of moods makes it mysterious why moods feel phenomenally undirected. In short, then, moods' felt undirectedness seems to preclude the viability of intentionalism about moods and, thereby, undermines intentionalism *as such*. Accordingly, the challenge is to show how an intentionalist account of moods can be pursued.

The standard solution (e.g., Crane 1998; Seager 1999, 2002; Tye 2008) is to say that moods are experiences of *very general entities* such as the whole world or everything. For example, when one is depressed, one does not seem to experience this or that specific thing or event as pointless, but one does experience the *whole world* as pointless. This would explain the peculiar phenomenal character of moods: they do not exhibit directedness towards anything *in particular*, because they represent something *very general*. Accordingly, on this view, describing moods as phenomenally undirected is simply inaccurate: they are in fact better described as exhibiting a directedness of a more *generalized* (or unspecific) sort towards, e.g., *everything* or the *world*. However, as it will become clear soon, there are reasons to think that this is quite unsatisfying as it offers the intentionalist a way out at the high price of re-describing an intuitive datum, namely, moods' undirectedness.

Recently, however, Angela Mendelovici (2013a, b) has proposed something different that looks more promising. Unlike the proponents of the standard solution, she takes the phenomenology at face value and concedes that there are at least some genuinely undirected moods: she aims to provide an intentionalist treatment of those experiences on which moods are objectless but not contentless. I maintain that Mendelovici's strategy leads to the best option within the intentionalist field, precisely because it provides a tentative explanation of moods' undirectedness, as opposed to denying it. Nonetheless, as I will argue, her proposal is poorly motivated. For, contrary to what Mendelovici argues, introspection does not offer evidence in support of her view—arguably, it does offer reasons *against*. So, the problem of moods is still there.

The paper is structured as follows. I start with the problems of the standard account (§2). Then, I move on to Mendelovici's account and point out that it is the best option within the intentionalist field as it deals with moods' undirectedness, instead of denying it (§3). In §4, I illustrate the motivation problem affecting Mendelovici's proposal and her Introspective Argument to motivate it. In §5, I argue that such an argument fails. In §§6-7, I address two possible rejoinders to my objections. Finally, I conclude that Mendelovici's proposal is not properly motivated: at best, there is *no* reason to accept it; at worst, we have reasons *not* to accept it. Intentionalism, thus, has still a serious problem with moods.

2 Problems with standard intentionalism about moods

Adopting the standard view commits the intentionalist to denying moods' undirectedness and accepting directedness towards the world (or everything). In this section, I provide some reasons to find this move unconvincing.

The first reason is purely phenomenological. Consider the prominently subjective character that moods have. Usually, one feels in a certain mood *no matter what* one might encounter in experience. This suggests two things. First, a good description of mood experience does not (have to) appeal to the world or worldly features. Contrast this, for example, with the case of visual experience, where the only way of describing the experience seems to be by means of describing worldly objects and their worldly features.

Second, the world does not seem in any way “responsible” for one’s experience, quite the contrary: the mood seems to be there quite independently of it. Either way, mood experience does not really seem to inherently involve the world—or anything coming from the world-side, so to speak: being in a certain mood appears to be an entirely subjective matter, a matter of having certain feelings that have nothing to do with the way the world is.

An immediate reaction to this is to say that *at least* some moods do seem world-directed. After all, for example, when one is depressed, one tends to say things like “The world is awful.” The question is: Is this true in *every* case?

If one replies “Yes,” then a triviality issue arises. Indeed, if every mood is manifestly directed towards the world, then every mood is clearly contentful and represents that the world is F. So, it becomes somewhat mysterious why intentionalism has a problem with moods at all. This problem only arises if some moods do not exhibit any directedness (not even world-directedness) in their phenomenal character. So, in order to avoid trivialization, we have to acknowledge that at least some moods are *phenomenally* such that they do not manifestly involve the world.

This suggests that the proper answer to the question above is negative: there are at least some *genuinely undirected* moods and they are *phenomenally different* from those that manifestly involve the world. In these cases, a description in terms of world-directedness is simply *phenomenologically inadequate*—in this respect I agree with Mendelovici (2013a, b). (Importantly, these experiences are not theoretical posits that we postulate as a result of an argument, but the very intuitive datum that generates the problem for intentionalism.) It also looks *ad hoc*, since there appears to be no reason to make such a move independently of the desire to defend intentionalism.

Now, this generates a further serious difficulty for the intentionalist. (Again, I agree with Mendelovici on this point.) For if we accept that there are some genuinely undirected moods that are

phenomenally different from the world-directed ones, then we have a phenomenal difference between two different phenomenal types of moods that calls for an explanation. However, it is hard to see how the standard account could provide such an explanation since it can only appeal to one type of content—i.e., <the world is F>. This is clearly an issue because intentionalism implies that every phenomenal difference is a difference in content, and this would be openly false in this case.

To wrap up, the intentionalist faces the following dilemma: *Either* every mood is world-directed, but then the problem is trivial; *or* some moods are genuinely undirected, but then standard intentionalism is not able to explain the phenomenal difference between world-directed and undirected moods. The upshot is: the intentionalist has to acknowledge that there are at least some undirected moods, but a standard explanation of this phenomenal character in terms of representing the world as F is simply inadequate. And this gives us a second reason not to be satisfied with the standard account.

So, a different intentionalist solution to the problem of moods is needed, one capable of accommodating genuine undirectedness and explaining the difference between the two types of mood. This is what Mendelovici wants to do: instead of denying genuine undirectedness, she aims to show that genuine undirectedness is compatible with intentionalism.³ The next section presents her proposal.

3 Mendelovici’s non-standard intentionalism about moods

Roughly, Mendelovici’s idea is that undirected moods do not represent *any* particular or general object, but *unbound affective properties*. Call this *Non-standard Intentionalism About Moods* (NIAM). I come now to present it in more detail. Since Mendelovici’s account is an extension of her account of emotions, I will start with emotions (§3.1) and then turn to moods (§3.2).

3.1. Emotions

According to Mendelovici’s phenomenological description, there are three relevant points to be stressed about emotions. First, they always exhibit directedness towards a specific object, the object they represent. Second, they involve *affective qualities*—e.g., “the ‘fearfulness’ bit of an experience of fear” (Mendelovici, 2013a: 128). Third, affective qualities “behave” in *exactly the same way* as the visual qualities (e.g., colors or shapes) involved in visual experiences: they *fully qualify* the object of the experience. That is, they appear as properties of the object of the experience. For example, “a fear of a dog qualifies the dog itself as being a certain way, as being *scary*” (Mendelovici, 2013a: 128). In

³ This should be tempting even for intentionalists skeptical about genuine undirectedness: if the intentionalist is not forced to deny genuine undirectedness, her position in the dialectic with the anti-intentionalist is much stronger.

other words, emotions, like visual experiences, are *introspectively transparent* (Harman 1990): when one tries to introspectively focus on their phenomenal character, the only things one is aware of are the represented objects of the experience and their properties.⁴

For Mendelovici, the strict analogy between visual qualities and affective qualities strongly suggests that those affective qualities *are represented properties* of the represented objects of the experience. In both cases, all the phenomenal components that constitute phenomenal character entirely appear as properties of the represented object(s): introspectively, no “extra” phenomenal component is left. And this is the main reason to believe in an intentionalist account of emotions.⁵

3.2. Moods

What about undirected moods? Let me start, again, with a bit of phenomenology. First of all, there are *phenomenal similarities* between emotions and moods: they seem to involve the very same affective qualities, in that, for every emotion, there is a corresponding mood (irritation corresponds to anger, elation to joy, anxiety to fear, and so on). Nonetheless, there is also a relevant *difference*: while emotions always exhibit directedness towards something in particular, moods (at least sometimes) do not. Thus, here is Mendelovici’s proposal to account for undirected moods:

Undirected moods seem to be a lot like emotions, except that they lack objects at which they are directed. My suggestion is to take this appearance at face value. Moods are what we get when we have an emotion that is not directed at something: a representation of a mere affective property. ... We feel the fear, joy, or anxiety, but we don’t experience anything as having these affective qualities. (2013a: 131)

The suggestion is that the fact that undirected moods are objectless does not lead to conclude that they are also contentless. On the contrary, they represent instances of *affective properties* not bound to *any* (particular or general) object. This accounts for their peculiar phenomenal character—their undirectedness. For example, if one undergoes an experience of objectless depression, then one’s depression represents mere pointlessness. This is why one’s depression appears undirected.

Clearly, this account is grounded in the partial *phenomenal similarity* between emotions and

⁴ Mendelovici does not mention explicitly the word “transparency,” but her phenomenological description and the explicit analogy drawn with visual experience (the paradigm case of transparent experiences) strongly suggest the transparency terminology.

⁵ Mendelovici (2013b) defends an account of affective properties as “*Edenic*” *properties*, namely, *sui generis*, purely phenomenal, subjective properties, not instantiated in the outside world, but only represented in content (Chalmers 2006). However, the issue as to what is the best account of the nature of affective properties is left open in Mendelovici (2013a). This is not really going to play any substantive role in this paper, so I am mentioning it mainly for the sake of completeness.

moods—the fact that both involve *the very same* affective qualities. On the other hand, Mendelovici has an explanation for the *difference* between affective states too—i.e., directed *versus* undirected affective experiences. Finally, since the affective qualities involved in undirected moods are represented properties, the account preserves the intentionalist identification of phenomenal character and intentional content.

To sum up, NIAM defends the viability of intentionalism without committing the intentionalist to denying genuine undirectedness; on the contrary, it provides a coherent account of it. Moreover, NIAM is perfectly compatible with standard intentionalism’s story about the generalized directedness exhibited by some moods. (For example, when I undergo some world-directed happiness, my experience represents joyfulness as instantiated by the whole world.) So, it has all the benefits of the standard account without the drawbacks. NIAM, thus, *looks* like a proper solution to the problem of moods. As such, it raises the biggest challenge to the anti-intentionalist. Accordingly, NIAM is the best option for the intentionalist.

4 The motivation problem and Mendelovici’s reply

Many worries can be raised against NIAM. In particular, the non-propositional format of the (alleged) content of undirected moods, as well as its being composed only of unbound properties, is certainly problematic.^{6,7} The focus of my criticism here, though, will be on a set of issues having to do with whether or not NIAM is *properly motivated*. I present first the problem and Mendelovici’s reply to it (this section) and then (§§5-6) argue that such a reply is not enough to properly motivate NIAM.

4.1. The motivation problem

Recall that, according to Mendelovici, we have *introspective reasons* to accept intentionalism about emotions: like visual experiences, emotions are introspectively transparent. That is, the affective qualities involved in emotions appear to fully qualify the object of the experience. This suggests that

⁶ If undirected moods represent unbound properties, then they are non-propositional representations (a mere unbound property is not a proposition). Two issues are intertwined here. The first is whether we should allow for non-propositional representations at all. The second concerns the specific format of non-propositional representation that NIAM attributes to undirected moods. For non-propositional representations are standardly taken to be *objectual*—i.e., representations of objects. However, according to NIAM, undirected moods do not represent objects, but mere unbound properties. So, the question is whether we should allow for this new format of non-propositional representation, which is not objectual, but “*proprietal*,” as Kriegel (Ms) calls it. (Kriegel raises doubts concerning the very intelligibility of “proprietal” representation and its being part of our psychological repertoire.)

⁷ If one considers also the account of affective properties as *sui generis* Edenic properties, then further worries arise concerning questions of ontological parsimony or the very idea of Edenic content.

affective qualities are represented properties and, thereby, directly and strongly motivates intentionalism about emotions.

However, things are quite different with undirected moods: the affective qualities (a) do not appear to qualify any object, since no object (of any sort) is involved, and (b) do not even appear to be localized anywhere in outside space or one's own body. So, based on what is introspectively available, the affective qualities involved in undirected moods cannot be attributed to something other than the experience itself. If so, undirected moods are introspectively *opaque*: the affective qualities they involve appear to qualify the experience.

Accordingly, allowing for genuinely undirected moods means allowing for opaque experiences. But then a problem arises concerning the *motivation* of NIAM—call it the *motivation problem*:

(MOTIVATION PROBLEM) In the case of emotions, transparency is the direct motivation to believe that affective qualities are represented properties. Yet, undirected moods are opaque. So, *what is the reason* to believe that the affective qualities involved in undirected moods are represented properties as well, as opposed to intrinsic qualities of the experience?⁸

As presented here, this is a problem for NIAM, but it arguably affects intentionalism about moods as such.

Importantly, the motivation problem does not question the *viability* of the view—this has to do with telling a coherent story about what the content of undirected moods is. Rather, it concerns why we should believe such a story in the first place. For NIAM tells us that the affective qualities involved in undirected moods *are* unbound represented properties; but, still, they do not *appear* so, introspectively. This apparent lack of motivation threatens the proposal—no matter how internally coherent that proposal is—as it leaves in us the impression as of a mismatch between the explanation (what affective qualities are) and the *explanandum* (what affective qualities appear to be). And, given such a mismatch, the easiest and more natural explanation seems to be that affective qualities involved in undirected moods are not represented. Hence, to be a proper intentionalist reply to the problem of moods, NIAM not only has to be able to (i) provide a viable story about undirected moods' content, but it also (ii) has to give us a reason to believe that story.⁹

⁸ This is my own formulation of the problem. Kind (2013) raises a similar issue.

⁹ Mendelovici (2013a: 129) also refers to the fact that an intentionalist account has to deal with these two separate problems.

Although there might be a point of view *internal* to intentionalism from which to evaluate NIAM's motivation, the right perspective to take here is *external* to intentionalism. By "external perspective" I mean a theory-neutral point of view, i.e. the point of view of a philosopher who is not already an intentionalist or an anti-intentionalist. The reason for taking such a point of view has to do with the context in which the problem of moods arises. The problem of moods calls into question whether or not intentionalism is to be embraced at all, and clearly the reply to such an issue cannot speak to an intentionalist alone but has to address and persuade (at least) those who are neutral.

4.2. Mendelovici's Introspective Argument

Mendelovici is aware that a motivation problem affects NIAM and explicitly acknowledges that introspection on undirected moods does not provide any direct evidence for considering affective qualities to be represented properties. Yet, she maintains, the same introspective considerations that work as direct evidence for intentionalism about emotions can be exploited as *indirect* evidence to motivate intentionalism about undirected moods.

Here is the argument, as she herself formulates it (Mendelovici 2013a: 133):

INTROSPECTIVE ARGUMENT

- | | |
|--------------|--|
| (P1) | The affective qualities involved in undirected moods are involved in corresponding emotions. |
| (P2) | The affective qualities involved in emotions are represented affective properties.
Therefore, |
| (Conclusion) | The affective qualities involved in undirected moods are represented affective properties. |

The background assumption is that introspection is a reliable source of information concerning the nature of phenomenal qualities, so both the premises are justified on the basis of it. In particular, (P1) is justified by the partial phenomenal similarity between emotions and moods, (P2) by the transparency of emotions. Thus, introspective data on emotions indirectly motivate Mendelovici's proposal that the affective qualities involved in undirected moods are represented properties.

5 The reply does not work

I am going to argue now that this does not solve to the motivation problem. For if one endorses (P2) as Mendelovici does, then a dilemma arises: *Either* there is introspective justification for (P2), but the argument is not valid (§5.1); *or* the argument is valid but unsound because there is no introspective justification for (P2) (§5.2). So, introspection does not offer evidence in support of NIAM.

5.1. First horn: The argument is not valid

One way to read (P2) is as a claim about affective qualities *when involved in emotions*. That is:

(P2*) *When involved in emotions*, affective qualities are represented properties.

If so, then (P2) is justified by the introspective evidence coming from emotions, given that the latter are transparent. Assuming (P2*), then the argument should conclude:

(Conclusion*) *When involved in undirected moods*, affective qualities are represented properties.

However, (Conclusion*) does not follow from (P1) and (P2*): no conclusion concerning affective qualities and their represented-ness when involved in undirected moods can be drawn in this way. For the fact that affective qualities “behave” as represented properties in the case of emotions does not *guarantee* that they “behave” in the same way in the case of undirected moods, even if they are one and the same kind of properties. The same kind of things can “behave” differently in different contexts. For example, the property of *being tall* is the same property both when possessed by basketball players and when possessed by table tennis players. But, from the fact that being tall helps basketball players score more points it does not follow that it helps table tennis players score more points.

5.2. Second horn: (P2) is not justified

Of course, this is correct unless represented-ness is a property that affective qualities have *qua* affective qualities—and not simply when involved in emotions. This suggests a different reading of (P2):

(P2**) The affective qualities involved in emotions are, *qua affective qualities*, represented properties.

Now, assuming (P2**), the argument claims that something about the nature of affective qualities as such—i.e., that they *are* represented properties—is established based on what we gather from introspective observation of affective qualities involved in emotions. Introspection on emotion guarantees that affective qualities *are* represented properties, so they *are* represented properties in *any* case—even in the case of moods. Hence, the conclusion follows. So, this version of the argument *is* valid.

Now, the question is whether what we gather from introspection really justifies (P2) read as (P2**). I argue that it does not.

Let me start with the following remark. If observation on emotional experience makes it introspectively manifest that affective qualities are represented properties, then the fact that affective qualities are represented in moods as well should be evident from paying attention (introspectively) to our emotional experience (if not a trivial consequence of it). That intentionalism is true of moods is then merely an easy conclusion that one could draw without paying any attention at all to mood experience—some baby logic would suffice to establish this conclusion. But, then, why think that there is an interesting question about undirected moods? Why should one worry about them at all? The impression is that the argument as above simply transforms what seemed a difficult question (with no clear answer) into a question that is easily answered (and, perhaps, a question of little interest). These considerations suggest that something in the argument *is* wrong. But *what* exactly is it?

5.2.1. *Bad sampling*

I think the main problem has to do with the way (P2**) is justified. My first point is that, as it stands, the justification of (P2**) cannot be accepted because is based on some *sampling error* that prevents us to properly justify the claim. To see why, consider this passage from Fodor (1997: 152):

Suppose we've been considering whether oak trees shed their leaves in winter; and suppose it turns out, on re-examining the records, that all our positive instances are observations of oak trees on the north side of hills. Then we would no longer think of the generalization about oak trees losing their leaves in the winter as unambiguously well-confirmed; oak data confirm oak generalizations only if they are an *unbiased* sample of the oak population; which, on the current assumption, our data aren't. Maybe, in the present case, the generalization that the instances really confirm is that *oak trees on the north side of hills* lose their leaves in winter. But notice that discovering a sampling error of this sort would be no reason at all for doubting that 'tree' is a kind. Rather, the worry would be that maybe *oak tree on the north side of a hill is a kind too*. If it is, then our data are equivocal between two perfectly ok, projectible hypotheses: the one

that goes *blah, blah...oak trees...* and the one that goes *blah, blah.... oak trees on the north sides of hills...* . When we discover the sampling error, we regard neither generalization as unequivocally confirmed by data that are instance of both, and this is precisely because the data *are* instances of both. *The sampling error means that the data are equivocal, not that the hypotheses are unprojectible. ... Biased samples don't confirm anything.*

Now, I think that we are facing something analogous with our case. (P2**) is inductively justified roughly in the following way:

JUSTIFICATION OF (P2)**

- (J) Every introspectively observed instance of affective quality involved in emotion is a represented property. Hence, *by induction*,
- (Conclusion) (P2**) is true.

The sample that is supposed to confirm (P2**) is built up in such a way that considers as relevant *only* the way different affective qualities (e.g., scariness, joyfulness, etc.) “behave” in the case of emotions. The point is that this sample is not *unbiased*. For, as in Fodor’s oak tree case, we are dealing with *two* kinds, and not just one: *affective qualities* and *affective qualities involved in emotion*. And the instances we are considering are instances of *both* these kinds. So, for all we know, our data might merely confirm something about *affective qualities involved in emotion*—as opposed to confirming something about *affective qualities*. In other words, as it is, our sample is not able to tell us whether the affective qualities that we observe in emotions are represented *qua* affective qualities or *qua* affective qualities involved in emotions. Hence, we have no justification for (P2**).

*5.2.2. The right sample does not confirm (P2**). Quite the opposite...*

So, the right sample to justify (P2**) must be capable to confirm that being represented is a fact about the affective qualities *qua* affective qualities (and not *qua* involved in emotions). In other words, it is to be capable of confirming (AQ):

- (AQ) *Affective qualities* are represented properties.

This is to say that the justification of (P2**) depends, in fact, on the justification of (AQ). So, the real question becomes: Is (AQ) confirmed by introspection? My reply is: No.

To evaluate (AQ), we have to bring in moods. Recall, given (P1), moods are the other relevant case in which instances of affective qualities can be observed, and there are no good reasons to leave them out. Also, there are no reasons not to treat observations on moods as equally relevant as observations on emotions for establishing something about the nature of affective qualities.

However, as soon as we consider the sample, so construed, a new problem arises. Indeed, the opacity of undirected moods suggests that affective qualities are *not* represented properties, as much as the transparency of emotion suggests that they are. As far as introspection is concerned, then, one natural hypothesis is that represented-ness is just something accidental about affective qualities and only associated to them when they are involved in emotions. If so, affective qualities are *not* represented properties.¹⁰ Granted, this is not good news for the intentionalist, but it seems perfectly fine from a theory-neutral standpoint, given the deliverances of introspection that we have.

To sum up, the first problem is that (P2**) cannot be unequivocally confirmed, and thus justified, introspectively because of a sampling error. This led us to set the conditions for an unbiased sample to properly evaluate (P2**), which made clear that the justification of (P2**) depends on the justification of (AQ). At this point, a second problem arises. As soon as we consider the unbiased sample, which includes moods, we are left with no reasons to accept the claim that affective qualities are represented properties. Rather, the opacity of moods seems to offer us reasons *against* it, and thus against (P2**). So, there seems no way to find (*unbiased*) introspective evidence in support of (P2**), which comes down to say that introspection does not justify (P2). At best, what introspective data support is a claim about affective qualities involved in emotions—which takes us back to §5.1.

5.2.3. *What is wrong*

We are now in the position to see more clearly what is wrong with this version of the argument. To be properly justified, (P2**) presupposes that, in addition to observations of emotions, our observations of undirected moods confirm that affective qualities are represented properties. So, having an introspective justification for (P2**) presupposes that we already have introspective evidence for accepting NIAM. But the point is precisely that we have no such evidence because our introspective observations on undirected moods do not confirm that affective qualities are represented properties—quite the opposite. So, as soon as we notice the sampling error and set the right conditions for an

¹⁰ Plausibly, this suggests that the question about their represented-ness is not relevant for the identification of affective qualities. (In other words, affective qualities are not identical to a kind of represented properties.) They are simply identified based on different criteria/features—perhaps, certain purely phenomenal qualitative features that are shared by emotions and moods. To be sure, this is difficult to accept for an intentionalist, but I do not see any problem from a theory-neutral standpoint.

unbiased sample, we see clearly that (P2**) cannot be introspectively justified—and any attempt to do it would simply beg the question. For, given what we gather from introspection, that affective qualities are represented is not something that we can assume as introspectively manifest to justify one of the premises of our argument. Rather, it is what the argument itself should show.

5.3. Taking stock

In conclusion, there seems to be no way to (a) have (P2) justified and (b) preserve the validity of the argument, at the same time. If we interpret (P2) as (P2*), then we have no reason to *extend* our conclusions on affective qualities involved in emotions to affective qualities involved in undirected moods. Whereas if we interpret (P2) as (P2**), then we have no reason to *restrict* our sample to emotions alone to determine facts concerning the nature of affective qualities *qua* affective qualities—and the right sample leaves (P2**) unjustified. This generates the dilemma mentioned at the beginning of this section: *Either* there is introspective justification for (P2), but the argument is not valid; *or* the argument is valid but unsound because there is no introspective justification for (P2). Either way, introspection does not offer evidence to motivate NIAM, contrary to what Mendelovici argues. So, there is no reason to extend the motivation for intentionalism about emotions to undirected moods.

One final remark. The argument seems to implicitly assume that emotions have something deeper to tell about the nature of affective qualities than moods do. This is the rationale behind the move of extending the introspective motivation for intentionalism about emotions to the case of undirected moods. The lesson of the discussion above, thus, is this: if the point is to capture the nature of affective qualities as such, then we cannot really keep separated intentionalism about emotion from intentionalism about moods. And this means that we cannot even assume that the former is motivated unless the latter is motivated. Thus, far from allowing an extension of intentionalism about emotions to undirected moods, introspection actually suggests that the motivation problem affecting intentionalism about moods also extends to intentionalism about emotions. So, on the basis of what introspection provides, we do not get the motivation problem solved but, on the contrary, *extended*. On the other hand, if we restrict intentionalism about emotions to what happens in emotions, then it has little to say about the nature of affective qualities as such and, *a fortiori*, has little to say about the case of undirected moods.

6 Emotion reveals more?

One option left to the intentionalist is to give us a reason to accept that the case of emotions reveals more about the nature of affective qualities than moods do. This amounts to challenge the standard and *prima facie* plausible claim that moods are opaque. So, one rejoinder to my objection might go as follows:

In emotions, affective qualities qualify a represented object—this is introspectively manifest. So, we have evidence that they are represented properties. In undirected moods, instead, all we have is that affective qualities do not appear as qualifications of represented objects. However—and this is where this characterization would depart from a more standard one—this does not mean that *they appear to qualify the experience*. They simply do not appear to qualify *anything*. Based on this, *we cannot really tell* whether they are unbound represented properties or (non-represented) experiential features. *For all we know*, they can be represented as well as non-represented properties. And *this* is the reason why we need to appeal to emotion.

In short, given that introspection on emotions tells us *something*, whereas introspection on undirected moods tells us *nothing*, we do have reasons to consider the former as more revealing than the latter. If so, the Introspective Argument works.

I have two responses to give, they are quite independent from each other. The first one says something stronger: the rejoinder simply gets the phenomenology of undirected moods wrong. The right description is the one that supports the opacity of moods. Thus, once again, introspection offers direct evidence against the intrinsic represented-ness of affective qualities—and, so, against NIAM itself (§6.1). The second says something weaker: even if one accepts that introspection on undirected moods is silent, still this does not make emotions speak louder (§6.2).

6.1. First response

Let me first unpack the rejoinder. Both the proponent of the rejoinder and the anti-intentionalist agree on (PD):

(PD) The affective qualities involved in undirected moods do not appear to qualify any represented object.

In principle, (PD) admits two interpretations:

- (QE) The affective qualities involved in undirected moods appear to qualify the experience.
- (QN) The affective qualities involved in undirected moods do not appear to qualify *anything at all* (see Mendelovici 2013b: 148).

(QE) is the way of reading (PD) that favors the opacity of undirected moods. The first point that the rejoinder makes is, then, that the anti-intentionalist conflates (PD) with (QE); whereas (PD) is in fact best captured by (QN). So, the phenomenal character of undirected moods is such that they are *neither transparent nor opaque*.

Now, this phenomenal character—and this is the second point—is such that it does not allow us to determine which of the following two mutually exclusive *explanations* is true:

- (EX1) The affective qualities involved in moods are (unbound) represented properties (NIAM);
- (EX2) The affective qualities involved in moods are non-represented properties of the experience (anti-intentionalism).

Intuitively, the idea seems to be: since affective qualities do not appear to qualify anything, in experiencing them, it is not clear whether we are dealing just with the way our own experiencing is or with some unbound represented property. Hence, the rejoinder concludes: *because of their phenomenal character*, undirected moods are silent as to whether (EX1) or (EX2) is true. So, this phenomenology cannot be explained directly but requires us to appeal to what happens in emotions.

Let us now assume the whole package offered by the rejoinder: (i) the phenomenal character of moods is made of affective qualities that do not appear to qualify anything—that is, (QN); (ii) unbound properties *can* be intentional contents. I am going to argue that this delivers bad predictions that the proponent of the rejoinder herself does not want to have.

When we encounter in introspection something like phenomenal qualities that do not appear to qualify *anything at all* (if there is any such phenomenology), then it seems fair to say that we are dealing with phenomenal qualities that “float free” in-between us and an apparent mind-independent world. (Mendelovici herself (2013b: 148) suggests this: “in the case of affective properties, we *can* experience free-floating, or *unbound*, instances of the properties.”)

Now, if there is anything like a phenomenology of unbound-ness, this really seems what it should be like: a phenomenology of “free-floating” phenomenal qualities. But if we admit that unbound phenomenal qualities *can* be intentional contents, then we should have no problem in telling that those “free-floating” qualities are represented properties and, thus, the contents of our experience. So, as a matter of fact, undirected moods should count as transparent and there should not be lack of evidence: moods’ phenomenal character should tell us directly that affective qualities are represented.¹¹ (It might be that we do not *recognize* the transparency of moods. However, this can only depend on irrelevant facts—e.g., our ignorance about the relevant kind of content (unbound properties) or our reluctance in accepting that kind of content.¹²) But clearly, then, the question is: Why do we need an argument to indirectly motivate NIAM at all? If the rejoinder is right, the motivation problem is not solved: it is simply *not predicted*. That is, given the rejoinder, it should not even arise. This is the first bad prediction.¹³

The second bad prediction concerns undirectedness. There are two issues here. Firstly, given the phenomenological description that we are assuming, it is legitimate to expect that we should feel at least some directedness towards the “free-floating” affective qualities that we encounter. To cut a long story short, the point is the following. Despite its metaphorical flavor, the intuitive idea behind directedness should be relatively clear: in undergoing a directed experience, one has the feeling of being *confronted* with something that (at the very least) *appears* as other than oneself or one’s own experience.¹⁴ *Vice versa*, undirected experiences (if there are any) are those that lack this feeling. Now, if affective qualities involved in undirected moods do not appear to qualify anything at all, then this

¹¹ One might be worried that this is not really transparency. However, it is quite fair to say that what counts as transparent is not independent of what counts as content, after all. Recall, the whole point of transparency is that it offers introspective evidence for phenomenal character’s being intentional content. And it is evidence because, in transparent experiences, what introspection reveals about experience is that phenomenal character has a content-like structure. But the point is precisely what legitimately counts as content. Standardly, we consider as transparent only those experiences whose phenomenal character manifests a property-object structure. This does not have to depend on anything intrinsically special about such phenomenal structure, though; it rather depends on the specific notion of represented content we rely on—i.e., a standard notion on which content is propositional (or objectual). But, if we enlarge our notion of content, then this presumably affects also what phenomenal structures count as content-like structures. In this way, we are also enlarging our notion of what counts as a transparent experience. *If* there is a kind of content made by unbound properties, then undirected moods and emotions are transparent in the same sense, though not in the same way: both their phenomenal characters have a content-like structure.

¹² These count as irrelevant facts because the transparency of experience is not epistemic transparency. On the contrary, it has to do with the way the phenomenal structure of the experience *is* (see, e.g., Kind 2003 on this). So, it is a metaphysical matter. An ideal introspector with the relevant theoretical tools is able to recognize a certain phenomenal character as content.

¹³ Recall, the rejoinder presupposes the Introspective Argument (and so the motivation problem) since it is a reply that is supposed to make that argument work. More generally, part of the force of NIAM is that it acknowledges that the phenomenology of undirected moods generates a motivation problem.

¹⁴ I take this to be the most phenomenologically neutral, and thus less committal, way to describe phenomenal directedness. For more on this, see Frey (2013, 2015) and Kriegel (2011: Ch. 3).

suggests that they should not appear as qualifications of one's own experience either.¹⁵ But, if so, they should appear as other than one's own experience and our experience should feel directed towards them.

Secondly, and more generally, it is hard to imagine what it would be like for an experience to be transparent without involving any feeling of directedness: undirectedness and transparency just seem inconsistent. For directedness is likely to be the *ground* of the transparency of the experience.¹⁶

So, it seems that if the rejoinder is right, then undirectedness is not what we should expect from mood experiences: given the way their phenomenal character is described, we should expect them to be directed towards some (represented) affective quality. In other words, they should be experiences *as of* unbound affective qualities. However, *this* is not their phenomenology, and moods' undirectedness *is* accepted by NIAM as a basic introspective datum on independent grounds. More: it is what the account is engineered to explain in the first place. So, not only does the rejoinder not generate the motivation problem—first bad prediction; it also does not predict moods' undirectedness—second bad prediction.

Let me sum up. I argued that if we assume that the rejoinder is true, neither (i) the motivation problem nor (ii) moods' undirectedness are predicted. Since both the anti-intentionalist and the proponent of the rejoinder (and NIAM more generally) agree that the motivation problem and moods' undirectedness depend on the specific phenomenal character that undirected moods have, a good phenomenological description should be able to preserve them. Hence, the lesson that I draw is that the rejoinder gets the phenomenology of undirected moods wrong.¹⁷

The only way to preserve the motivation problem and moods' undirectedness is to admit that affective qualities *appear* to qualify the experience itself and, thereby, choose (QE) as the right way to look at (PD). This allows us to see clearly that the motivation problem as well as moods' undirectedness have a common source: the phenomenal character of undirected moods does not exhibit a content-like structure—whatever the relevant notion of content might be—as it involves affective qualities that appear to qualify the experience itself.¹⁸ If this is correct, there can be no introspective

¹⁵ After all, if they did, (QE) would be the right way of describing the phenomenal character of undirected moods, but the rejoinder claims that it is not.

¹⁶ It is natural to think that an experience is transparent *because*, in undergoing it, one feels to be confronting something other than oneself or one's own experience. Kriegel (2011: Ch.3 fn100) and Frey (2015: 8) make a similar point. According to Loar (2003), phenomenal directedness just *is* transparency (plus the idea that the intentional object does not have to be actual). Also, Kriegel (2007) takes phenomenal directedness and transparency to be strictly connected.

¹⁷ Of course, another option is to drop the idea that unbound properties *can* be represented. However, that would call into question the very intelligibility of the view and, thus, its viability. Kriegel (Ms) does that. However, here I am more interested in showing that the view is not supported by introspection. So, I do not follow this path.

¹⁸ The supporter of NIAM might drop undirectedness, but this seems to betray the very spirit of the view. After all, NIAM was engineered to take moods' phenomenal character, their undirectedness, at face value and explain it in intentionalist

way out of the motivation problem: introspection does not offer (indirect) evidence in favor of NIAM, but arguably direct evidence *against* it, given the opacity of undirected moods. In other words, introspection is the source of the problem, not its solution.

6.2. Second response

Let us now assume, for the sake of the argument, that the rejoinder is right: introspection on undirected moods is silent and, based on it, we are not able to tell whether affective qualities are represented properties or not. Still, I claim, the fact that undirected moods are a case of lack of evidence is no reason to think that we can use introspection on emotions alone to draw general conclusions concerning the nature of affective qualities.

To see this point more vividly, consider the following example. Suppose that I want to discover something about oak trees. I observe the oak trees on the north side of hills and discover that, every autumn, they lose their leaves. Now, I decide that I want to check whether losing leaves is a feature of oak trees *qua* oak trees. I know that oak trees are planted on the south side of hills too, so I observe them for one year and discover something interesting: they do not lose their leaves. I go on and observe them for a long time. Again, nothing happens: the leaves are still there well attached to the branches of the trees. This suggests to me that oak trees on the south side of hills do not lose their leaves. Hence, I conclude that losing leaves is not a feature of oak trees *qua* oak trees.

However, one day, a very good botanist comes to me with no reason or intention to lie and tells me: “It is very rare, but there is one species of trees, called X-trees, that lose their leaves once every one hundred years. Maybe oak trees on the south side of hills lose their leaves once every one hundred years too.” So, now the evidence seems to be compatible with two hypotheses: (H1) oak trees on the south side of hills do not lose their leaves; (H2) oak trees on the south side of hills lose their leaves once every one hundred years. Does what happens in the case of oak trees on the north side of hills help decide in favor of (H2) against (H1)?

I do not see why it should. The evidence from the north side of hills would only help if, for some reasons, I already knew that losing leaves is an intrinsic feature of oak trees *qua* oak trees. But this is precisely what I want to establish with my observations in the first place. So, I cannot assume it. To be sure, based on my observation and what I know, I can say that the evidence coming from the

terms. In doing that, it was supposed to grant a point to the anti-intentionalist: that moods are genuinely undirected. Now, if one drops genuine undirectedness or re-describes it in terms of (unbound) property-directedness, then one is not really granting anything to the anti-intentionalist. On the contrary, one just seems to be accommodating the phenomenology of moods in a way that is more suitable for an intentionalist treatment. But, if so, NIAM loses its force and its interest: it turns out to be as unsatisfactory as the standard intentionalist account.

south side of hills no longer disconfirms the hypothesis that oak trees *qua* oak trees lose their leaves. However, this is not enough to conclude that that hypothesis *is* confirmed. The hypothesis is still *not* confirmed by the evidence available.

The case of undirected moods is analogous in some relevant respects. Introspection on emotions suggests the hypothesis that affective qualities are represented properties. Undirected moods *neither* confirm *nor* disconfirm this hypothesis. But it would be wrong to conclude that our hypothesis is confirmed because we cannot prove that affective properties are non-represented in the case of moods. The option that affective qualities in the case of moods are not represented is still there: it is not removed by the evidence coming from emotions. The right conclusion to draw is that our hypothesis still stands as *not* confirmed by the introspective evidence available.¹⁹ For, as far as introspection is concerned, affective qualities can still be represented properties as well as non-represented properties. In a nutshell, the silence of undirected moods does not make emotions speak louder.

To conclude, appealing to the silence of undirected moods seemed a *prima facie* viable way to make the Introspective Argument work. This strategy seems to re-interpret the motivation problem as originating from some lack of evidence rooted in the phenomenal character of undirected moods. However, this move ultimately fails. In the worst case, the rejoinder gets the phenomenology of undirected moods wrong: undirected moods are opaque experiences. If this is correct, then the motivation problem arises not because we are not able to tell whether the affective qualities involved in undirected moods are represented properties, but because the phenomenal character of undirected moods gives us reason to think that they are *not*. However, even granting a lack of evidence scenario (best case), introspection neither confirms nor disconfirms the intrinsic represented-ness of affective qualities. Either way, introspection does not motivate NIAM.

7 Giving up on introspection?

Another rejoinder follows a different strategy: it gives up on introspection and tries to provide non-introspective reasons to motivate intentionalism about moods. I will argue that this strategy fails.

The rejoinder goes as follows:

To be sure, transparency is a motivation for intentionalism, but it is not the only one: there are at least two other. First, *theoretical unity*: intentionalism offers a unified account of phenomenal character. Moods would be the only exception to such an account, and this is weird. Second,

¹⁹ Importantly, “not confirmed” is not the same as *disconfirmed*.

physicalistic reduction: as pointed out in §1, given intentionalism, a naturalistic account of mental representation suffices to get a full physicalistic reduction of phenomenal character and close the explanatory gap. Other views, anti-intentionalism in particular, are not similarly helpful. So, even if undirected moods are introspectively opaque, we have two independent, non-introspective motivations to accept NIAM.

My response divides in two. Concerning theoretical unity, I have two things to say. First, moods are probably the case that raises the greatest challenge to intentionalism, but they are not the only case in which intentionalism looks non-obviously true. Others include orgasms, itches, and pain—not to mention emotions themselves or even certain visual experiences (e.g., blurry vision). Second, theoretical unity is an important virtue of a theory and is a reason to find intentionalism intriguing and worth pursuing, but it seems to me that beyond a certain point it cannot be used as a reply to counter-examples. Suppose we have many dubious cases of transparent experiences—as in fact we do—, so that we cannot appeal to transparency to directly motivate intentionalism in those cases. Appealing to theoretical unity in each single case sounds too much like deciding to ignore cases in which the theory does not work (instead of providing a real motivation to accept that theory).

As regards physicalistic reduction, this might be enough of a motivation to find intentionalism *promising*, but it is not enough of a motivation to buy intentionalism where transparency does not apply. As a matter of fact, intentionalism *promises* to solve the hard problem of consciousness, but has to prove that it is in fact *able* to do that, and that has to be independently evaluable. In other words, in order to prove that it is able to solve the hard problem, it has to show that it is true of *every* experience. So, the promise of solving the hard problem cannot be used as a motivation to accept the truth of intentionalism, but rather the other way around—otherwise, the motivation is circular.²⁰ While it is true that anti-intentionalism does not help with this issue, if intentionalism turns out to be false, it would be just as unhelpful.

8 Conclusion

Let me sum up. Moods' undirectedness potentially undermines intentionalism's project of reducing phenomenal character to intentional content. Mendelovici's proposal, NIAM, is the best candidate reply to this threat, since it is the only existing intentionalist account that really deals with undirected moods. The idea is to extend intentionalism about emotions to undirected moods as representing

²⁰ Kind (2013) also says something along (more or less) the same lines—she says that the physicalist reduction motivation is “loaded.” So, I basically agree with her on this point.

unbound affective properties. This solution really looks like a proper reply to the problem of moods.

However, allowing for genuine undirectedness means allowing for the introspective opacity of undirected moods. So, a motivation problem arises, since intentionalism is motivated precisely by introspective transparency. To solve this problem, Mendelovici puts forward the following Introspective Argument: (P1) the affective qualities involved in undirected moods are involved in corresponding emotions; (P2) the affective qualities involved in emotions are represented affective properties; therefore, the affective qualities involved in undirected moods are represented affective properties.

I have argued, against (P2), that the opacity of undirected moods undermines the justification for assuming intentionalism about emotions as a premise in the argument (§5). The argument only works if we can prove that emotions reveal more about the nature of affective qualities than moods do. One strategy to do this is to maintain that introspection on undirected moods is silent, whereas introspection on emotions tells us something. However, in §6, I argued that this strategy fails. Firstly, it relies on a wrong description of the phenomenal character of undirected moods. Secondly, and quite independently, even if we accept such a description, the silence of undirected moods does not make emotions speak louder. The upshot is that the attempt to introspectively motivate NIAM fails: (at best) there is no introspective reason to accept the view. Arguably, introspection gives us reasons *not* to accept it.

I have considered other non-introspective ways to motivate NIAM, the appeal to theoretical unity and physicalistic reduction, and found them unconvincing. Therefore, NIAM is not properly motivated and, thus, is not a proper reply to the problem of moods. For, even if NIAM shows that intentionalism about moods is *viable*, we still have no particular reason to think that it is *true*, given the phenomenal character of moods.

Therefore, since NIAM is the best candidate-reply to the problem of moods, the problem is still there and casts a shadow not only on intentionalism about moods, but on the whole intentionalist project of reducing phenomenal character to intentional content.

One final remark. A general lesson suggested by the discussion is that undirectedness raises a major difficulty for intentionalism. More precisely, it seems to force the intentionalist between Scylla and Charybdis. Scylla consists in conceding too little to the phenomenology of moods: in this way the intentionality is safe, but the phenomenology seems to get lost—this is what the standard account does. Charybdis consists in conceding too much: in this way the phenomenology is safe, but the intentionality seems to get lost—this is what NIAM does. This should not come as a surprise, after all.

Intentionalism is committed to treat undirectedness as a peculiar phenomenal character, ultimately rooted in some kind of intentionality. But this is *per se* quite odd. For intentionality is supposed to be the source of experience's directedness. Precisely because of that, at least *prima facie*, it does not really seem the right kind of thing to appeal to, if one is to explain *undirectedness*. To be sure, we can always say that there is some intentionality underlying a certain phenomenal character—as long as we provide a coherent story. But this does not seem to illuminate the phenomenon of undirectedness, it just leaves it as obscure, if not even more mysterious.²¹

References

- Chalmers, D. (2006). "Perception and the Fall from Eden" in T.S. Gendler & J. Hawthorne (eds.), *Perceptual Experience*. Oxford: Oxford University Press.
- Crane, T. (1998). "Intentionality as the Mark of the Mental" in T. Crane (ed.), *Contemporary Issues in the Philosophy of Mind*. Cambridge, UK: Cambridge University Press.
- Deonna, J. & F. Teroni (2012). *The Emotions: A Philosophical Introduction*. London: Routledge.
- Dretske, F. (1981). *Knowledge and the Flow of Information*. Cambridge, MA: MIT Press.
- _____. (1988). *Explaining Behavior: Reasons in a World of Causes*. Cambridge, MA: MIT Press.
- _____. (1995). *Naturalizing the Mind*. Cambridge, MA: MIT Press.
- _____. (1996). "Phenomenal Externalism, or If Meanings Ain't in the Head, Where Are Qualia?," *Philosophical Issues*, 7: 143-58.
- Fodor, J. (1990). *A Theory of Content and Other Essays*. Cambridge, MA: MIT Press.
- _____. (1997). "Special Sciences: Still Autonomous After All These Years," *Philosophical Perspectives*, 11: 149-163.
- Frey, C. (2013). "Phenomenal Presence" in U. Kriegel (ed.), *Phenomenal Intentionality*. Oxford: Oxford University Press.
- _____. (2015). "Sensuous Experience, Phenomenal Presence, and Perceptual Availability," *International Journal of Philosophical Studies*, 23(2): 237–254.
- Harman, G. (1990). "The Intrinsic Qualities of Experience," *Philosophical Perspectives*, 4: 31-52.

²¹ For constantly discussing these topics with me and helpful comments on previous drafts of this paper, I am grateful to Luca Barlassina and Uriah Kriegel. For very useful discussions and detailed comments, I would also like to thank Adriano Angelucci, Enrico Cinti, Nevia Dolcini, Samuele Iaquinto, Dave Ingram, Fred Kroon, Cristian Mariani, Mog Stapleton, Alfredo Tomasetta, Giuliano Torrenzo, Nick Young and Sandro Zucchi. Some of the materials that led to this final version of the paper were presented at conferences in Bucharest, L'Aquila, and Noto: thanks to the audiences for their comments and questions.

- Kind, A. (2003). "What's So Transparent about Transparency?" *Philosophical Studies* 115(3): 225-244.
- _____. (2013). "The Case Against Representationalism about Moods" in U. Kriegel (ed.), *Current Controversies in the Philosophy of Mind*. London: Routledge.
- Kriegel, U. (2007). "Intentional Inexistence and Phenomenal Intentionality," *Philosophical Perspectives*, 21(1): 307-40.
- _____. (2011). *The Sources of Intentionality*. Oxford: Oxford University Press.
- _____. (Ms). "The Intentional Structure of Moods." Unpublished manuscript.
- Loar, B. (2003). "Transparent Experience and the Availability of Qualia" in Q. Smith & A. Jolic (eds.), *Consciousness: New Philosophical Perspectives*. Oxford: Oxford University Press.
- Mendelovici, A. (2013a). "Intentionalism about Moods," *Thought*, 2: 126-136.
- _____. (2013b). "Pure Intentionalism about Emotions and Moods" in U. Kriegel (ed.), *Current Controversies in the Philosophy of Mind*. London: Routledge.
- Millikan, R.G. (1984). *Language, Thought, and Other Biological Categories*. Cambridge, MA: MIT Press.
- _____. (1989). "Biosemantics," *Journal of Philosophy*, 86: 281-97.
- Seager, W. (1999). *Theories of Consciousness: An Introduction and Assessment*. London: Routledge.
- _____. (2002). "Emotional Introspection," *Consciousness and Cognition*, 11: 666-687.
- Tye, M. (1995). *Ten Problems of Consciousness*. Cambridge, MA: MIT Press.
- _____. (2000). *Consciousness, Color, and Content*. Cambridge, MA: MIT Press.
- _____. (2002). "Representationalism and the Transparency of Experience," *Noûs*, 36(1): 137-151.
- _____. (2008). "The Experience of Emotion," *Revue Internationale de Philosophie*, 62: 25-50.