



Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:

Schroeter, L;Schroeter, F;Jones, K

Title:

Do Emotions Represent Values?

Date:

2015-09-01

Citation:

Schroeter, L., Schroeter, F. & Jones, K. (2015). Do Emotions Represent Values?. *Dialectica*, 69 (3), pp.357-380. <https://doi.org/10.1111/1746-8361.12108>.

Persistent Link:

<https://hdl.handle.net/11343/56688>

DO EMOTIONS REPRESENT VALUES?

Laura Schroeter, François Schroeter, Karen Jones

Abstract: This paper proposes a framework for adjudicating debates about representationalism in the case of emotions, i.e. the claim that emotions attribute evaluative properties to target objects or events. We argue that representationalism faces a significant explanatory challenge: proponents must establish that a representation relation linking emotions and value is explanatorily necessary. We use the case of perception to bring out the difficulties in meeting this explanatory challenge.

Many theorists take the analogy between emotions and perceptions literally in one crucial respect: just as visual perception represents objects as having properties such as shape, trajectory and color, emotions represent objects and situations as having evaluative properties. Fear, for instance, represents situations as fearsome or dangerous, and indignation represents situations as unjust. In both perception and emotion, the representation of the relevant properties is supposed to be non-conceptual: it does not depend on the subject's having a concept dedicated to representing the relevant feature.¹

In the case of perception, theorists who invoke non-conceptual content take perceptual states to play a central role both in explaining competence with certain concepts (e.g. concepts of particular shapes, colors, or orientations) and in justifying perceptual beliefs that deploy those concepts (Peacocke 1992). If emotions are non-conceptual representational states, perhaps they could play similar semantic and epistemic roles. The suggestion would be that just as the non-conceptual content of visual experiences helps secure a determinate reference for the concepts expressed by 'red' or 'oval', the non-conceptual content of emotions might help secure a determinate reference for evaluative concepts expressed by evaluative terms like 'unjust' or 'contemptible'. And in both cases, conceptually articulated judgments might be directly epistemically justified by experiences or emotions with matching non-conceptual contents.²

The claim that emotions represent objects and situations as having evaluative properties has recently been challenged (e.g. Salmela 2011; Whiting 2012; Deonna and Teroni forthcoming). In this paper, we propose a framework for adjudicating debates between representationalists, who hold that emotions represent evaluative properties, and skeptics. Our primary interest is in whether emotions have representational properties that would allow them to play the sort of epistemic and

¹ For discussion and defense of the claim that emotions represent their targets as having evaluative properties, see (de Sousa 1987; Tappolet 2000; Johnston 2001; Döring 2003, 2007; Prinz 2004; Deonna 2006; Elgin 2008; Brady 2013; Montague 2014). For a critical survey of the literature on the connections between emotion and value, see (Todd 2014).

² For theorists who defend the epistemic roles allegedly played by the non-conceptual evaluative content of emotions, see (Tappolet 2000; Döring 2003, 2007; Elgin 2008).

semantic roles commonly attributed to non-conceptual contents of perceptual experiences. How plausible is it that emotions attribute evaluative properties in ways that ground the epistemology and semantics of moral and other evaluative domains?

We start by articulating the notion of representation as standardly understood in philosophy of mind and cognitive science (§1). Next we argue that a representationalist about emotions must establish that representation is explanatorily necessary. Drawing on the extensive literature in the philosophy of perception, we show that this is a significant explanatory burden (§2). In the final section (§3), we highlight the difficulties of discharging this burden in the case of emotions.

1. Representation: vehicles and underived accuracy conditions

What it takes for a mental state to count as a representation is a large and controversial question. But some specification of what exactly is at issue in claiming that emotional states represent evaluative properties is needed to ensure that the debate is not a purely verbal dispute. We'll take as our starting point the notion of representation as understood by the representational theory of mind (RTM). RTM is committed to internal representations that serve as the direct bearers of representational content.³ Although it is not uncontroversial, RTM is the dominant paradigm in philosophy and cognitive science for thinking about the representational properties of mental states and processes. Moreover, if emotions represent evaluative properties in the sense specified by RTM, then their evaluative contents would be well-suited to grounding the epistemology and semantics of evaluative judgment in the ways we've sketched. Of course, some theorists might wish to endorse a weaker notion of mental representation than the RTM model that's standardly invoked in philosophy of mind and cognitive science. But the robust notion of representation proposed by RTM nonetheless provides a useful reference point, which can help to demarcate the weaker senses of representation that these theorists may wish to endorse.

RTM can be contrasted with behaviorist and interpretivist views of mental states, which attribute representational content on the basis of global input/output dispositions without any commitment as to the nature of the internal structures that ground those dispositions (e.g. Davidson 1980, 1984; Dennett 1981; Stalnaker 1984). According to interpretivism, a robot connected to a giant look-up table correlating micro-level sensory inputs with micro-behavioral outputs could count as having mental states that represent dungeons, dragons, or danger, provided it was able to pass a Turing test. Not so according to RTM, since such a robot would lack the relevant internal structures that function as representations of these things.

³ Jerry Fodor is an influential contemporary advocate of representationalism (Fodor 1975, 1987, 1990), who combines representationalism with a computational account of mental states and processes. But representationalism is compatible with very different accounts of the nature of mental states. Early modern theorists, for instance, generally assumed that mental states were structured by phenomenally or intellectually introspectable ideas (Descartes 1641; Locke 1690; Hume 1740). And contemporary connectionists and dynamic systems theorists, who disagree with classical computationalists about the nature of mental states, often argue in favor of internal representations (Smolensky 1988; Rumelhart 1989; van Gelder 1995).

The interpretivist approach yields a highly deflationary account of mental representation: indeed the idea of mental representation without representations is perhaps better understood as ‘as-if’ representation. According to interpretivism, ascribing particular representational contents is a matter of overall resemblance in global behavioral dispositions to some paradigm.⁴ This means that there would be very little at stake in arguing over the truth of the claim that emotions represent values on an interpretivist account of representation. We can all agree that an animal’s fear of a predator resembles a judgment of dangerousness *to some degree* in the general pattern of input/output dispositions.

However, this observation about overall behavioral similarity carries little theoretical weight. In particular, it cannot ground the semantic and epistemic priority of representation by perceptual or emotional states. It makes no sense to say that evaluative concepts acquire their content from the non-conceptual content of emotions, if all these intentional contents are attributed as a package on the basis of global behavioral dispositions. Given this interdependence, it’s not clear how emotions could play any special role in explaining the acquisition of evaluative concepts. Similarly, one might worry that interpretivism does not accord individual mental states enough independence to vindicate a genuine epistemic basing relation between emotional states and evaluative beliefs. The notion of basing suggests a realist approach to mental states, in which individual states persist over time and have independent causal powers. On the interpretivist picture, however, mental states are theoretical abstractions from the totality of input/output dispositions, so they do not reflect the internal structure of the system. So it is hard to see how interpretivism can support the claim that some token beliefs are based on and justified by specific token perceptions or emotions.

So what are the marks of genuine mental representation, according to the representational theory of mind? We’ll emphasize two key features: (i) representational vehicles and (ii) underived accuracy conditions.

Mental states like beliefs, conjectures, hallucinations, or emotions count as genuinely representational states insofar as they include as constituents representational vehicles – i.e. internal structures that are the direct bearers of representational content. These vehicles are typically thought of as part of a complex representational system in which different elements bear systematic type relations to specific contents and can be recombined in systematic ways to represent distinct states of affairs. The standard paradigm in cognitive science is to identify representational vehicles with computational states. But RTM can accommodate more traditional views of the nature

⁴ Since resemblance is a matter of degree and since there will inevitably be trade-offs in how to prioritize among partially conflicting dispositions, the correctness of a particular content ascription will depend on the interpreter’s current interests and explanatory priorities (Davidson 1984; Dennett 1987; Stalnaker 2008). This suggests that interpretivist approaches take a purely instrumental approach to representational content. On such an instrumentalist approach, any non-semantic characterization of a system will radically underdetermine the semantic interpretation of that system. However, David Lewis has expressed confidence that there are interpretive constraints that suffice to eliminate this sort of radical indeterminacy and so vindicate underived non-instrumental accuracy conditions (Lewis 1974).

of representational vehicles – e.g. as phenomenal states or as irreducible conceptual structures (Pitt 2013).⁵

What about underived accuracy conditions? Your belief that Bruno is a dog, for instance, represents the world as being a certain way. It follows that this belief has *accuracy conditions*: some possible ways the world could be are in accordance with the content of the belief and other ways are not. We can also talk about the content of desires: some ways the world might be are in accordance with the content of the desire, and others are not. The content of constituent representational elements can likewise be specified in terms of accuracy conditions. For instance, the ‘dog’ concept that figures in your belief about Bruno accurately represents all and only dogs, and your circular visual percept when you gaze at the ball in Times Square on New Year’s Eve accurately represents all and only spheres of a certain size, shape, distance, etc.⁶

In this context we’re interested in the underived accuracy conditions of mental states. Virtually any state can be used as a representation: you can decide to use a red flag to represent danger and you can take your aching muscles as a sign that your workout was effective. But these states seem to derive their representational contents from our contingent dispositions to interpret them in certain ways. Since you could equally well take red flags or aching muscles to represent something different, neither of these states has underived representational content. In contrast, your propositional attitudes and perceptions don’t seem to derive their representational contents from other states in this way: they have their content originally, not in virtue of being interpreted via some other intentional state. On a representationalist account, for instance, the content of perceptual experience does not derive from any perceptual beliefs or dispositions to believe – the representational content is intrinsic to the perceptual state itself.

⁵ There are many different ways of fleshing out the core idea of a representational vehicle. According to advocates of the computationalist theory of mind, representational vehicles should be understood on the model of formally defined symbols in a computer program (Fodor 1975). So a desire for chocolate, for instance, is a dispositional state involving a mental symbol that stands for chocolate – a symbol akin to a mental word that’s individuated solely by its syntactic properties. In the teleosemantic paradigm, it is common to think of representations as signals that mediate between a producing mechanism and a consuming mechanism (Millikan 1984, ch. 6). Expressions in natural language play such a mediating role: linguistic signals are produced by a speaker and consumed by a hearer. Similarly, perceptual states are produced by sensory systems and consumed by action-guiding systems within an organism. A different way of understanding representational vehicles is as phenomenal states that are salient from the first person perspective. On Christopher Peacocke’s account of non-conceptual content, for instance, subjectively discernible properties of the visual field, such as phenomenal redness (red’), are assigned to points in a three-dimensional egocentric representation of the environment (a scenario), which is itself manifest in phenomenal consciousness (Peacocke 1992, 7-8, and ch. 3).

⁶ Accuracy conditions contribute to the evaluation of whole attitudes: when the world accords with the accuracy conditions of a belief or perception, the attitude as a whole is true or veridical; and when the world accords with the accuracy conditions of a desire, hope or intention, the attitude as a whole is satisfied.

If a state has underived accuracy conditions, it can misrepresent: i.e. it can represent the world as being a way that it's not. It follows that underived accuracy conditions cannot be explained in terms of any direct relation between a representational vehicle and the world – a vehicle with accuracy conditions can exist even if what it represents does not, whereas two things stand in a direct relation only if they both exist.

What's crucial to RTM about the emotions, then, is that emotions have as components representational vehicles with underived accuracy conditions of a specific kind: emotions attribute specific evaluative properties to the state of affairs or objects targeted. For instance, the emotion of fear that Teddy feels towards the (alleged) monster under his bed and that Sally feels towards the idea of tumbling over the cliff she's standing on represents a specific evaluative property (e.g. dangerous) and attributes that property to the emotion's target. The fear itself (or the phenomenology of fear) constitutes a representation of an evaluative property that combines with a representation of the target to form a complex representational state that attributes the property to the target. So emotions have underived accuracy conditions such as:

- Teddy's fear: [The monster is dangerous.]
- Sally's fear: [My proximity to the cliff is dangerous.]

Teddy's and Sally's fears are veridical or falsidical depending on whether these accuracy conditions are satisfied.

On the version of RTM we're exploring here, the content of emotion is non-conceptual in the sense that the subject need not possess the concept of DANGER in order to represent danger via fear. So the representationalist can appeal to the non-conceptual evaluative content of fear to explain competence with the corresponding evaluative concept, by linking the possession of the concept of DANGER with the emotional state of fear in such a way that the concept inherits its representational content from the content of the emotion.⁷ Moreover, on this approach, the emotion and the evaluative judgment are psychologically real, causally independent states. So

⁷ If evaluative judgments inherit their content in this way, there may be untapped resources for solving the disagreement problem in metaethics. A central problem for realists in metaethics is to explain how it's possible for different individuals to pick out the same property with evaluative predicates like 'right', 'unjust', 'fearsome', or 'contemptible' despite widespread, apparently irresolvable, disagreement about which things fall into the extension of such terms. On the face of it, we seem to ascribe properties when we use such terms in evaluative judgments. Indeed, we seem to be ascribing the *same* property when we disagree about such questions as whether abortion is wrong, military intervention unjust, or François Hollande contemptible – otherwise we wouldn't be making contradictory claims about the world, we'd be just talking past each other. Many metaethicists think the problem of disagreement is insuperable, so that the initial assumption that speakers co-refer with evaluative terms must be rejected. Indeed, the problem of disagreement has played a large role in motivating a range of anti-realist accounts of the semantics of evaluative terms, including expressivism, contextualism, new-wave relativism, error theory, and fictionalism (Schroeter and Schroeter 2013). If different subjects' fears have the same evaluative property as content, realists could appeal to this content to explain how different speakers manage to co-refer in their judgments about danger, despite apparently irresolvable disagreements about what it takes to be dangerous.

fear can stand in a genuine epistemic basing relation to the evaluative judgment that shares its representational content.

How do we adjudicate the debate between representationalists, who hold that emotions represent evaluative properties, and skeptics, who may grant that emotions represent objects or situations, but deny that they represent them as having an evaluative property? In the next section, we turn to a standard strategy for adjudicating this sort of question: appealing to explanatory necessity.

2. The explanatory challenge

When does an internal state count as a genuine representation with specific accuracy conditions? Consider a much-discussed example: magnetotactic bacteria. Certain anaerobic bacteria possess internal structures – magnetosomes – that are sensitive to magnetic fields. In the bacterium’s home environment, these structures normally align with the earth’s magnetic field, causing the bacterium to move downwards towards oxygen-poor water, which is essential to its survival. So there is a reliable co-variation – backed by causal laws and selected for by evolutionary processes – between the alignment of the magnetosomes with the direction of a magnetic field, the North Pole, the ocean depths, oxygen-free water, survival-friendly conditions, and so on. Should we say that the alignment of the magnetosomes in these bacteria is a representation with the content, say, *magnetic north is in that direction*, or *oxygen-poor water is over there*, or *it’s safe down there*?

A major objection to the attribution of such representational properties is that they don’t seem to be doing any real explanatory work. As many theorists have pointed out, we can give a full explanation of the movements of the bacterium in one direction rather than another by providing a mechanical account of (i) why magnetosomes are sensitive to magnetic fields, (ii) how the presence of a magnetic field controls their alignment within the organism, (iii) how alignment of the magnetosomes affects the orientation of the organism as a whole, and (iv) how the organism’s flagella propel it in the direction in which it’s oriented. There is no need to mention representational states with underived accuracy conditions in this explanation – the mechanical account is a complete proximal explanation the organism’s reaction to the stimulus.⁸

This sort of example suggests a simple and powerful criterion for determining whether a state counts as a representation – i.e. a state with underived accuracy conditions: Does ascribing representational properties to the state play a genuine

⁸ Of course, speaking loosely, we call the magnetosome’s alignment a ‘perception’ of oxygen-free water or the direction of safety – and this may make the explanation more colorful, just as we may say our car doesn’t want to start in cold weather. But such colorful intentional descriptions don’t add anything to the causal explanation of the underlying processes that produce the relevant behavior. Moreover, if taken literally such intentional descriptions suggest more complex psychological capacities than the target system really has. One sign that the content ascription is playing no essential explanatory role is that there is no serious scientific dispute about precisely which content should be ascribed: it makes no difference to the underlying explanation if we label the magnetosomes as representing *oxygen-free water that way* or *no danger down there* or *magnetic field stronger in this direction*. On this point, see especially (Schulte 2015).

explanatory role? If all the real explanatory work can be done without mentioning the alleged representational content, by simply citing independently specifiable mechanisms or dispositions, then we should conclude that there is no representation involved.⁹

The task for representationalists about emotions, then, is to show that construing emotions as representing specific evaluative properties is part of the best, most economical explanation of the psychological processes underlying perception, behavior, and cognition. In particular, RTM is committed to the following three claims as explanatorily indispensable:

First, emotional states are constituted in part by *representational vehicles*. Emotions have internal structural elements that function as representations of evaluative properties that combine with representations of objects or states of affairs. In other words, the emotional system involves a combinatorial system of internal representations, which can be interpreted as attributing properties to individuals. According to proponents of RTM, construing emotions as mere attitudes towards targets – i.e. mere dispositions towards representations of objects or situations targeted by the emotion – would leave crucial psychological phenomena unexplained.

Second, the representational component of an emotion has *underived accuracy conditions*. The properties attributed by emotions do not depend on the contingent explanatory interests of an external interpreter, nor are they a projection of the contents of subject's own conceptual system. Instead, it's crucial to understanding the nature of the emotional system considered in itself that emotions attribute specific properties to their targets. According to a proponent of RTM for the emotions, any account of the emotion of fear that failed to take this accuracy condition into account would be explanatorily inferior.

Third, emotions represent *evaluative properties*. Assuming that fear has underived accuracy conditions, why think it attributes danger? Why not suppose it attributes a different evaluative property like fearsomeness? Or sui generis evaluative properties that do not match any properties picked out by our commonsense evaluative concepts? Or a purely descriptive property? Or perhaps the accuracy conditions are indeterminate between several different properties. Once again, a proponent of RTM for the emotions must show that the precise evaluative contents they attribute play an essential explanatory role. This challenge is especially acute for theorists who claim that emotions play the same epistemic roles as perceptual experiences (Tappolet 2000; Döring 2003, 2007; Elgin 2008). If there isn't an appropriate match between the contents of emotions and the evaluative judgments based upon them, then evaluative judgments cannot acquire their justification (or inherit their content) directly from the emotional states that prompt them. And so the desired analogy with perception would fail.¹⁰

⁹ For similar appeals to explanatory necessity as a criterion for positing representational content see (Dennett 1981; Sterelny 1995; Neander 2006; Burge 2010; Schulte 2012).

¹⁰ The alleged role of emotions in the epistemic justification of evaluative judgments has been challenged in the literature (see (Brady 2013) for a critique and (Pelser 2014) for a defense). One prominent objection has focused on the unreliability of

To clarify the nature of the three explanatory tasks facing representationalists about emotions, it's helpful to consider the analogous challenge in the comparatively less controversial case of visual perception of features like shape, color, or distance. This will allow us to get a feel for how the challenge might be met, and for the difficulties one may encounter.

First, does visual perception involve representational vehicles? The claim is disputed. Certain enactivists, for instance, hold that perception is a dynamic interaction between the visual system and the environment: visual perception does not rely on any static, stable internal states that encode information about distal features of the environment such as shape, color or distance (Port and van Gelder 1995; Chemero 2009). In response, representationalists about perception argue that representations are needed to explain certain empirical phenomena. Consider a case of multi-modal perception: you may perceive a fire truck as having both visual and auditory properties – you perceive a single thing, the truck, both *as bright red* and *as emitting a blaring noise*. Arguably we cannot explain how specific visual properties get bound to specific auditory properties without positing cognitive processes in which visual and auditory representations that are matched and stably bound together (Nanay 2014). Thus, internal representational vehicles seem necessary to explain the structure of perceptual experience.

Second, does visual perception have underived accuracy conditions? This claim has been challenged from two very different directions. Relationalists hold that perceptual states are constituted in part by direct, non-intentional relations to the perceived objects (Brewer 2006; Campbell 2002; Martin 2002; Travis 2004). So a perception of a fire truck and a subjectively indiscernible hallucination of a fire truck are not members of any fundamental psychological kind defined by shared accuracy conditions. On this view, veridical perceptions needn't have underived accuracy conditions: they can be explained purely in terms of a direct engagement with certain features of the world.¹¹ Syntacticists reject underived accuracy conditions on very different grounds. In their view, psychological states, including perceptions, are individuated by the internal functional roles they play in a computationally defined cognitive system (Stich 1983; Egan 1995). So there is no need to bring in semantic properties like representation, or accuracy conditions in the definition of the states of computational systems.

In response, many representationalists about perception argue that ascribing underived accuracy conditions to perceptual states plays an ineliminable explanatory role. Tyler Burge (2010), for instance, argues that the best explanation of human

emotions in tracking evaluative facts. We're targeting a different type of worry about the alleged justificatory role of emotions – the mismatch in content.

¹¹ Charles Travis, for instance writes, '[...] perception, as such, simply places our surroundings in view; affords us awareness of them. There is no commitment to their being one way or another. It confronts us with what is there, so that, by attending, noting, recognizing, and otherwise exercising what capacities we have, we may [. . .] make out what is there for what it is—or, again, fail to [. . .] in perception things are not presented, or represented, to us as being thus and so. They are just presented to us, full stop.' (Travis 2004, 65).

vision, from mainstream vision science, posits genuine representations with underived accuracy conditions. Burge draws a sharp distinction between mere sensation and sensory perception: whereas sensation simply registers external stimulation, sense perception involves genuine representations of stable external features. Sensory systems like smell or taste, Burge argues, merely serve to register proximal stimuli of the sensory organs: your nose registers aspects of the chemical profiles of air it comes in contact with. Your sense of smell simply registers the passing show. But there's no explanatory gain in supposing your sense of smell constructs representations of stable, objective properties that are distinct from these immediate sensory registrations, and constitute the accuracy conditions for those representations. It's implausible, for instance, to posit olfactory representation to explain how an animal can locate food via smell: the animal can simply follow its nose towards the greater intensity of a smell pattern (415). According to Burge, even the recognition of specific smells – like the smell of banana – doesn't require a dedicated representation within the olfactory system itself. There's no reason to suppose olfactory sensory system generates representations of perspective-independent features of objects as opposed to simply registering immediate appearances. Our awareness of objective smell properties, Burge suggests, is more plausibly explained by invoking a concept of the property that's associated with olfactory pattern recognition (415).

In contrast, Burge argues that the visual system does involve representations of perspective-independent features. The input into the visual system is a pattern of retinal stimulation that directly registers a two-dimensional projection of incoming light. According to mainstream vision science, the visual system uses this retinal input to construct a three-dimensional representation of the environment attributing perspective-independent properties such as shape, color, size, texture, distance, object continuity, etc. (342-366). The visual system registers the passing show of light; but the visual system also abstracts from this immediate sensory input to create a visual model of these perspective-independent features of the world. Thus unlike olfaction, the visual system itself constructs representations of properties that diverge from the immediate sensory input registered by the system. According to Burge, this is what makes it plausible to attribute underived accuracy conditions to these visual sensory states. Thus, the key to defending underived accuracy conditions for perceptions is that accuracy conditions are necessary in order to explain the nature of the sensory system considered in itself, independently of more complex conceptual systems.

Consider now the third challenge: Assuming that perceptual states do have underived accuracy conditions, what exactly do they represent? A frog snaps at a passing fly. Assuming that its visual system is representing some feature of the world, what exactly is the content of that representation? Does the frog's visual system represent something as specific as the content that *there is a house fly (musca domestica) at location x*? Or does it represent that *there is a packet of frog food at location x*? Or does it just represent that *there is a small, dark, moving object at location x*?

There is a good deal of disagreement about this question, and alleged pre-theoretic intuitions about the case vary. However, the key to resolving the issue surely lies in showing that one's favored interpretation is superior from an explanatory point of view. In particular, the correct interpretation must underwrite the type of explanation that requires attributing representational content in the first place. Several theorists have argued that this explanatory strategy favors attributing the relatively superficial

perceptible property, *small, dark moving object* (e.g. Pietroski 1992; Sterelny 1995; Neander 2006; Schulte 2012).

It's important to note that the frog is doing more than simply responding to immediate visual stimulation: its visual system extracts information about object continuity, trajectory, rough shape and color from the immediate input. Moreover, this limited information exhausts the output forwarded to the frog's behavioral systems. When suitably motivated and undistracted, a frog will snap at any small, dark, moving object within a certain spatial range; and the frog will continue to snap at foreign insects or indigestible bee-bees *ad infinitum*, provided that the objects exhibit the relevant pattern of perceptible properties. Thus, the frog's visual system, considered in itself, seems sensitive to the perspective-independent perceptible property, *being a small, dark moving object*, but insensitive to more abstract functional or structural properties of that object. It seems wrong, therefore, to say that states of the frog's visual system represents more abstract properties such as *having the power to nourish Northern leopard frogs (rana pipiens)* or *being a housefly (musca domestica)*. The worry is that such interpretations are not supported by the explanatory projects of cognitive science, cognitive ethology, or commonsense psychology: the functioning of the perceptual system itself and the way it impacts on the organism's behavior and cognitive dispositions can be fully explained at the psychological level without attributing any abstract functional or natural kind contents. These abstract content ascriptions would be explanatorily idle.

In the philosophical literature on content determination, 'consumer-based' teleosemantic theories assign contents to perceptual representations on the basis of how they've been used by other 'consuming' subsystems within the organism (or the organism's ancestors), and ultimately how this use contributed to the survival and reproduction of the organism as a whole (Millikan 1984, 1991). Since the frog's visual representation leads to a snapping response, and ultimately to the digestion of food, consumer teleosemantic theorists argue that the visual representation has the content *frog food in location l*. An important worry about such approaches is that they are not suited to psychological explanations of the actual cognition and behavior of an individual organism (Pietroski 1992; Sterelny 1995; Neander 2006; Schulte 2012).

To see the intuitive force of this point, consider possible psychological explanations of a five-year-old's choice of ice cream over broccoli:

1. The child chooses ice cream over broccoli because in the past she perceived ice cream *as sweet* & she likes *sweet things*.
2. The child chooses ice cream over broccoli because in the past she perceived ice cream *as having concentrated nutrients* and she likes things with *concentrated nutrients*.
3. The child chooses ice cream over broccoli because in the past she perceived ice cream *as containing large proportion of mono or disaccharides* and she likes things *containing large proportion of mono or disaccharides*.

Whereas the first explanation seems apt, the second two explanations seem distinctly odd from a psychological point of view. Of course, it's plausible that natural selection favored certain gustatory sensitivities in the child's ancestors because they were correlated often enough with high nutrient foods and mono- and disaccharides fulfilled this nutritional role. But we should not confuse an historical explanation of why certain capacities originally evolved with a psychological explanation of an

individual's current behavior. From a psychological point of view, attributing a gustatory perception as of *concentrated nutrients* or of *containing mono- and disaccharides* suggests a different pattern of dispositions than the child actually exhibits: we'd expect a desire for concentrated nutrients to lead to a rejection of artificially sweetened foods or some (perhaps cognitively mediated) ability to distinguish between nutritious and non-nutritious foods (cf. Schulte 2015).

The discussion of perception illustrates the explanatory burden that a proponent of RTM for a domain must discharge. To support RTM about the emotions, one must show that the best explanation of the emotional system must posit representational vehicles with underived accuracy conditions. And if emotions are to provide epistemic warrant for evaluative beliefs, one must show that there is an appropriate match between the precise contents mandated by our best theories and the contents of the evaluative judgments based on them.

3. Meeting the challenge

The nature and force of the explanatory challenge facing representationalism about the emotions has not been carefully articulated in the literature. As a consequence, there has been no systematic attempt to meet it. To highlight the difficulty of the challenge, we explore two broad strategies representationalists might employ, drawing on either cognitive science or on phenomenology.

But first we'd like to consider an intuitive argument for representationalism, which may seem to provide some direct support for RTM without addressing the explanatory challenge we've articulated. It's uncontroversial that we assess emotions like fear for *appropriateness* or *fittingness*. It is tempting to conclude on the basis of this observation that emotions represent the properties specified by their appropriateness conditions: a fear is accurate just in case it satisfies its appropriateness conditions.

Now, emotions can be assessed along multiple dimensions: for example, we can ask of feeling an emotion in a particular context whether it is likely to further or hinder the agent's ends (Frank 1988), whether it is morally appropriate (D'Arms and Jacobson 2000), or whether it is well enough grounded in evidence available to the agent (Greenspan 1988, ch. 4). Distinct from each of these dimensions is the question of whether it is a response that is appropriate to, or fits the situation. It is typically assumed that this dimension is one of representational accuracy (D'Arms and Jacobson 2000, 66; Tappolet 2000, 170-1). However, the mere fact that emotions are liable for assessment in terms of fittingness does not establish that fittingness is to be interpreted along representationalist lines. Only if we have antecedently established that emotions function as representations can we conclude that fittingness assessment is assessment for representational accuracy. Consider, for example, the relationship between salivation and the ingestion of food. The salivary response evolved to enable the digestion of food. It is a fitting response just in case the organism is presented with food. But it's highly implausible to conclude that salivation literally represents properties like *being human food* or *being a reason to activate digestion*. Any biological state with an evolutionary function can be assessed for fittingness in this way. From the fact that fleeing behavior evolved to enable predator escape we can conclude that such behavior is fitting in the presence of a predator, but not that it literally represents the predator's presence. Some theorists may want to *argue* that a

certain type of fittingness captures the accuracy conditions of representational states. But if offered as an argument in favor of representationalism in the first place, an appeal to assessability for fittingness is circular.

Let's now turn to the cognitive science strategy for addressing the explanatory challenge for RTM about the emotions. One important strand in the psychological literature on emotions takes 'appraisals', understood as representations, to play a central constitutive role in emotions. Roughly, appraisals are evaluations of how features of the environment affect one's interests. Early theorists took appraisals to be conceptually articulated evaluative judgments such as: *this situation is offensive* (Lazarus 1984). More recent proponents have adopted a more flexible account, according to which appraisals may be automatic, non-conceptually articulated states (Scherer 2009; Moors et al. 2013; Moors 2013). At the micro-level, nonconceptual appraisals are held to represent features like urgency, novelty, relevance, pleasantness, or goal-congruence, whereas macro-level appraisals would represent 'core relational themes' like being a threat to one's survival or being an affront to one's values. The key point for present purposes is that appraisal theorists claim that these nonconceptual states are genuine representations.

Can appraisal theory meet our explanatory challenge? The first question is whether appraisal theory attributes representational vehicles in the sense required for RTM. According to appraisal theory, appraisals are representations because they are internal states that play an organizing role in shaping flexible responses to the environment (Scherer 2009; Moors 2013). However, this characterization of representation is significantly weaker than is standard in the philosophy of mind: on this definition, it it's an analytic truth that flexible behavior is the product of representational states (Parkinson 2013). But as we saw in the case of perception, positing internal representational vehicles is controversial. Enactivists about perception deny that there are static internal states that encode information about the environment – instead they seek to explain perception through dynamic interaction with features of the environment, rather than stable internal representations of those features. A similar enactivist approach may be even more plausible in the case of emotions. As we noted earlier, proponents of RTM for sensory perception argue that there are cognitive processes that seem to require operations on representational vehicles, rather than direct engagement with the environment. But it's not clear that appraisal theory posits any processes that require stable internal representations rather than direct engagement with the targets of emotions. Indeed, critics of appraisal theory within psychology push a more enactivist approach to the emotions, which does not require stable representational vehicles that guide the unfolding of emotion (Parkinson 2013; Colombetti 2014).

Second, it's not obvious that appraisal theory is committed to attributing underived accuracy conditions. Of course, it's natural to *describe* emotions or more basic appraisal states in terms of evaluative properties that make those states appropriate: to describe fear, for instance, as recognition of danger or a subpersonal state as sensitivity to goal-conduciveness. However, such ascriptions may just reflect our interests as external interpreters, or they may reflect the fact that emotional states tend to prompt the application of evaluative concepts within conceptually sophisticated subjects. To vindicate RTM about the emotions, one must explain why underived accuracy conditions play an essential role in the best explanation of the emotional

system itself. Why not suppose that fear is just a complex way of processing representations of target objects or events? Or that subpersonal appraisals are just nodes in an information-processing system?

Recall how Burge responds to this sort of challenge in the case of perceptual experience. He argues that the visual system doesn't merely register incoming information, it constructs representations of non-perspectival features of the environment such as color constancy or object continuity, which are independent of variations in sensory registrations of light and shadow, eye movement, etc. On the face of it, appraisals seem to directly register information from internal and external monitoring systems, but they don't seem to construct anything like a stable internal model of non-perspectival evaluative properties of triggering objects and events. Is there anything within the system posited by appraisal theory that requires us to posit an appearance/reality distinction within that system? Do appraisals abstract non-perspectival evaluative constancies from the passing show, the way the visual system abstracts shape constancies (like circularity) that remain constant across the perspectival distortions registered by the visual system? If not, then there is no explanatory gain in attributing underived accuracy conditions to appraisals. In such a system, the possibility of misrepresentation plays no essential explanatory role.

Third, assuming that appraisals have underived accuracy conditions, why think they pick out evaluative properties? Why not think that they represent proximal states, like physiological changes? Jesse Prinz suggests that adding a consumer-based teleosemantic constraint to his Dretskean theory of content determination can vindicate the assignment of evaluative properties ('core relational themes') as the representational content of emotions as opposed to non-evaluative properties ('bodily changes') (Prinz 2004, 55-60). In a similar spirit, Caroline Price uses Ruth Millikan's consumer-based teleosemantic theory to determine the content of an emotional state of appraisal (Price 2006). Teleosemantic theories claim that the representational content is determined, not by actual dispositions of a signaling system, but by the biological function it was selected to perform. According to consumer-based teleosemantics, the relevant function is determined by the role played by emotional signals within the system that 'consumes' the signal – the action-guiding system. Appealing to action-guiding mechanisms allows consumer teleosemantic theorists to ascribe representational contents that are relevant to the organism's goals: the frog's visual percept, for instance, represents frog-food rather than a superficial perceptual kind, and the frog's fear appraisal represents an imminent threat to the frog's survival (Price 2006).

However, such appeals to consumer-based teleosemantics will face the explanatory worry we flagged in section 2. These theories attribute highly abstract contents that seem to go beyond what's normally invoked in psychological explanations of the functioning of the perceptual or emotional systems considered in themselves. Thus they risk glossing over a genuine empirical difference between complex conceptually infused systems that are capable of representing abstract properties like *having the power to nourish Northern leopard frogs* and simple perceptual and emotional systems that cannot.

Let's now turn to the second strategy for meeting the explanatory challenge. Whereas the psychological literature on appraisal has emphasized that appraisals can be

unconscious, many philosophers have argued that the conscious aspect of emotions supports a representationalist account (Goldie 2000; Johnston 2001; Döring 2003, 2007; Prinz 2004; Montague 2014). The idea is that the phenomenal character of emotional experience presents the target object or event as having an evaluative property. A conscious episode of anger seems to present the object targeted as offensive in much the same way that a conscious visual experience seems to present a targeted object as red. This apparent attribution of a property does not seem to be traceable to a propositional attitude like a judgment or a desire that correlates with or helps constitute the emotion – for one can feel anger without any such attitude. So the attribution of a property must be accomplished by the emotional experience itself (Goldie 2000; Döring 2007).

How does the explanatory challenge apply to the phenomenological approach? The first question is whether emotional phenomenology constitutes a genuine representational vehicle. It's standard to distinguish between pure and impure representationalism about phenomenal experience: pure representationalists hold that every aspect of phenomenal experience has a distinct representational content, whereas impure representationalists hold that some aspects of phenomenology can be traced to the *manner* in which a content is entertained – e.g. in a particular sensory modality like vision or hearing or in a particular attitude like belief or hope or imagination (Chalmers 2004). There are strong reasons to favor the more moderate impure representationalism, even for those who accept the view that all phenomenal experiences involve representation (Chalmers 2004; Crane 2009). Given the distinction between manner and content, moreover, it's natural to take the distinctive phenomenology of emotions like anger to be a reflection of the *manner* in which a representation of the target object or event is entertained (Deonna and Teroni 2012, ch. 7; forthcoming). What's distinctive of the phenomenology of anger, after all, is a specific orientation or attitude towards the target: as many authors have emphasized, emotional phenomenology includes an experience of bodily arousal and action-preparedness, together with a distinctive pattern of attention and salience associated with the target. Your phenomenal fear of the bear in front of you, for instance, makes the bear more salient, it creates a bodily sensations characteristic of anxiety and urgency, it directs your attention towards a range of options for responding to the bear, and it mobilizes epistemic resources towards properties of the bear that are relevant to your response (de Sousa 1987; Brady 2013; Deonna and Teroni forthcoming). While all of these orientations involve representations of target objects, events or actions, close attention to the phenomenal experience of fear does not obviously involve the attribution of a specific property to the bear. So it's not obvious from the phenomenology alone why we should construe emotions as representations rather than an attitude directed towards target objects or events. To defend a representationalist position, we need to know what explanatory advantage there is in supposing that the phenomenology of fear functions as a vehicle that attributes properties to these targets.¹²

¹² In the phenomenological tradition, the term 'content' is commonly used in a very broad way: *phenomenal* content includes everything the subject is aware of in experience. In our example, we assume that your orientation towards the bear has phenomenal content. However, it does not follow that your orientation towards the bear involves a representational vehicle over an above representations of the target, with its own distinctive *representational* content (i.e. accuracy conditions). On the

The second question focuses on accuracy conditions. Even if we assume that the phenomenology of emotion functions as a representational vehicle, it is not clear how the phenomenology helps establish the claim that these vehicles have underived accuracy conditions. It's worth keeping in mind that a lot of what we're aware of in phenomenal experience is mere registration of immediate stimuli. In both olfactory and emotional experience, for instance, we are aware of subtle variations in intensity, but these phenomenal variations by themselves don't differentiate perspective-independent features of the world from perspectival factors. To vindicate underived accuracy conditions, the representationalist needs to establish that the emotions abstract from the passing show to construct a model of perspective-independent properties. It's not obvious that emotional phenomenology does this. In particular, the mere fact that emotional phenomenology directs our attention outwards does not establish that it has underived accuracy conditions, attributing a specific property to its target. (A startle response, for instance, is outward-directed, but it doesn't involve any model of a perspective-independent property of being startling, or dangerous, or important, or...)

It's worth emphasizing here that there's a general methodological worry about the phenomenological strategy in establishing *underived* accuracy conditions. Reflective verdicts about the phenomenal character of the emotions are always filtered through conceptually articulated judgments: your first-person reflection on how things seem when you experience an emotion inevitably brings into play a whole range of sophisticated concepts and reflective capacities that are not essential to the emotional system. But representationalists make a claim about what emotions, considered in themselves, represent. So the worry is that appealing to phenomenology isn't well suited to isolating what's intrinsic to the emotional system by itself, when it's not enriched by conceptual representation and reflective judgment.

The third question concerns which properties are represented. Assuming that emotional experiences have underived accuracy conditions, can emotional phenomenology help explain why emotions represent specific evaluative properties? An important worry is that the phenomenal aspects of experience are supposed to be occurrent psychological states (like a visual experience as of a red apple) rather than dispositional states (like a standing belief that apples grow in Tasmania). Wittgenstein famously pointed out that an occurrent state by itself cannot settle how it is to be associated with specific accuracy conditions: simply holding up a sign cannot determine how that sign should be interpreted or applied. Similarly, simply having an occurrent phenomenal experience of reddishness does not suffice to determine which worldly things the experience is accurate about: we need some way of projecting from the occurrent experience to a general property of things in the world that constitutes its accuracy conditions. So it is not clear how phenomenology alone can ground a representation relation to a specific property. Securing representation of a property

distinction between phenomenal and representational content see e.g. (Montague 2014). Whether fear plays the role of a vehicle of representation or of a manner in which representations of targets are entertained, its phenomenological signature will be fully integrated with other phenomenological aspects of the experience. The unity and integration of the phenomenological experience is thus no guide to the precise role fear plays in representation (Deonna and Teroni forthcoming).

seems to require an extra element for introducing generality – something like a *disposition* to identify things as falling into the relevant class or a *general principle* of interpretation. But appealing to dispositional or interpretive states takes us outside the realm of phenomenal appearances: the way fear feels to you right now doesn't reveal what else you are disposed to fear, much less when fear is accurate.

A natural response is to appeal to *causal dispositions*: certain properties are reliably causally correlated with phenomenological states like a fear experience or a reddish experience. While causal dispositions provide a way of generalizing beyond the occurrent state of the subject, however, they are much too promiscuous for a genuine representation relation: for instance, they can't distinguish between distal and proximal causes of a response and they don't allow any possibility of misrepresentation.¹³ To avoid these difficulties, naturalistic theories of intentionality have proposed different ways to enrich the dispositional account – for instance by appealing to proper function and natural selection. We won't need to go into these details here. The key point is that such dispositional theories seem to single out non-evaluative properties – e.g. what's disposed to cause fear under certain circumstances, rather than what warrants fear – as the ones represented by emotional experience. So dispositional theories seem ill-suited to vindicating the representationalist claim that emotions represent evaluative properties.

A different strategy for generalizing from occurrent phenomenology to the properties represented is to appeal to *resemblance*. Michelle Montague (2014) suggests that the representational content of phenomenal experience is determined an intrinsic resemblance relation between reddish phenomenology and redness in the world: i.e. we take our phenomenological reddish experience to reveal the nature of a property spread across the surface of a red ball. Similarly, she suggests, the phenomenology of sadness intrinsically resembles the objective character of sadness in the world: we take the badness of our sad experience to reveal the nature of the badness of the sad event of a friend dying (47). The first worry about this proposal is that it generates an error theory, which would undercut the appeal of the representationalist position: there is no phenomenal property spread over the surface of objects the way that red experiences suggest, and events in the world do not really have feelings of sadness. The second worry is that the resemblance account seems much less plausible in the case of emotions than in the case of sensations. Perhaps we naively assume that the intrinsic nature of redness resembles our reddish experiences. But in the case of emotions, the property we're tempted to attribute seems to differ both in kind and degree from our emotional phenomenology. Does it really seem like anger intrinsically resembles the offensiveness of an insult? Although your reaction is anger, the emotion you take an insult to express seems to be contempt – not anger. And does it really seem, as Montague suggests, that the *badness* of your sad feelings resembles the *badness* of the death of your friend? There is a radical difference in degree between the two: your feelings may be bad, but their badness pales in comparison to the badness of the death. What seems to be true in both of these examples is that there is a *reliable correlation* between the emotional experiences and certain evaluative features of the events that provoke them. But following up this

¹³ For overviews of the problems faced by co-variation-based accounts of representational content, see (Loewer 1997; Neander 2006; Adams and Aizawa 2010).

insight would lead us back towards a dispositional account of the alleged representational relation.¹⁴

We'll close this section by returning to the epistemic and semantic roles that representationalists about the emotions hope to vindicate. Many theorists claim that evaluative judgments can acquire their justification directly from the emotional states that prompt them, just as perceptual judgments acquire their justification from perceptual experience (Tappolet 2000; Döring 2003, 2007; Elgin 2008). If emotions are to play this role, it's not enough to that emotions represent some evaluative property or other – there must be an appropriate match between the properties represented by emotions and those represented by the evaluative judgments based upon them. Consider the case of vision: your visual experience represents an apple as being a determinate shade of red, and your perceptual judgment that the apple is red is directly justified in virtue of being based on that experience. Note that this direct justificatory relation depends on a specific type of match in the representational contents of the two states: the belief content attributes a *determinable* property of redness, and the perceptual content attributes a *determinate* shade of that determinable. If emotions are to play a similar epistemic role, we should expect a similar determinate/determinable relationship between the contents attributed by emotions and evaluative beliefs. But neither of the strategies we've considered in this section seem well suited to meeting this challenge. Even if we grant that emotions represent evaluative properties, both the appraisal theory and the phenomenal strategy seem likely to generate relatively superficial observable properties as what's represented by emotional systems considered on their own. As we have seen, if appraisal theories construe emotional systems as attributing abstract contents they seem to go beyond what's normally invoked in psychological explanations of the functioning of these systems. Similarly, it is hard to see how emotional phenomenology considered independently of any reflective cognitive system could fix reference to anything but superficial perceptually salient properties of the target object or event. But evaluative judgments about what's dangerous, offensive, or lovable seem to represent highly abstract, non-observational properties, which seem very different in kind from the superficial properties that might be represented by emotional systems on their own. There does not seem to be the sort of determinate/determinable relationship, for instance, between what anger might represent and what really is offensive from a reflective evaluative point of view.

4. Conclusion

To assess the proposal that emotions attribute evaluative properties to the targeted objects or events, we have made a lengthy foray into the literature on perception and representation. We think this foray has paid off, not because it shows that representationalism about the emotion is wrong, but rather because it shows just how hard representationalism is to defend, especially in the case of emotion. We have considered two strategies for defending representationalism, one based on cognitive science and the other based on an appeal to phenomenology. Neither strategy, we've argued, has yet produced an adequate response to the challenges facing representationalism. Perhaps such a response will be forthcoming, but if not, it does not follow that emotions have no epistemic significance or do not make an important contribution to the practical rationality of cognitively limited beings like ourselves.

¹⁴ On this disanalogy between color and value, see (Wiggins 1987; McDowell 1998).

On any theory of the nature of emotions, emotions reliably co-vary with evaluative facts under a variety of normal conditions, and so emotions can be recruited to play an evidential role in justifying evaluative beliefs. However, if representationalism about emotions is not true, emotions cannot play the sort of epistemic and semantic roles that perceptual experiences are thought to play: emotions cannot provide immediate justification for evaluative beliefs or ground the acquisition of evaluative concepts in the way that perceptual states are supposed to do in the case of perceptual beliefs and perceptual concepts.¹⁵

References

- ADAMS, F., and AIZAWA, K. 2010, "Causal Theories of Mental Content", in: E. Zalta, ed, *Stanford Encyclopedia of Philosophy*, pp. 1-47.
- BRADY, M. 2013, *Emotional Insight: the Epistemic Role of Emotional Experience*, Oxford: Oxford University Press.
- BREWER, B. 2006, "Perception and Content", *The European Journal of Philosophy*, **14**, pp. 165-181.
- BURGE, T. 2010, *Origins of Objectivity*, Oxford: Oxford University Press.
- CAMPBELL, J. 2002, *Reference and Consciousness*, Oxford: Oxford University Press.
- CHALMERS, D. 2004, "The Representational Character of Experience", in: B. Leiter, ed, *The Future for Philosophy*, Oxford: Oxford University Press, pp. 153–181.
- CHEMERO, A. 2009, *Radical Embodied Cognitive Science*, Cambridge, MA: MIT Press.
- COLOMBETTI, G. 2014, *The Feeling Body: Affective Science Meets the Enactive Mind*, Cambridge, MA: MIT Press.
- CRANE, T. 2009, "Intentionalism", in: B. P. McLaughlin, A. Beckermann and S. Walter, ed, *Oxford Handbook to the Philosophy of Mind*, Oxford: Oxford University Press, pp. 474-493.
- D'ARMS, J., and JACOBSON, D. 2000, "The Moralistic Fallacy: On the 'Appropriateness' of Emotions", *Philosophy and Phenomenological Research*, **61**, pp. 65-90.
- DAVIDSON, D. 1980, *Essays on Actions and Events*, Oxford: Clarendon Press.
- DAVIDSON, D. 1984, *Inquiries into Truth and Interpretation*, Oxford: Clarendon Press.
- DE SOUSA, R. 1987, *The Rationality of Emotion*, Cambridge, MA: MIT Press.
- DENNETT, D. 1981, "True Believers: The intentional strategy and why it works", in: A. Heath, ed, *Scientific Explanation: Papers Based on Herbert Spencer Lectures*, Oxford: Clarendon, pp. 151-167.
- DENNETT, D. 1987, *The Intentional Stance*, Cambridge, MA: MIT Press.
- DEONNA, J. 2006, "Emotion, Perception, and Perspective", *Dialectica*, **60**, pp. 29-46.
- DEONNA, J., and TERONI, F. 2012, *The emotions: A philosophical introduction*, London: Routledge.
- DEONNA, J., and TERONI, F. forthcoming, "Emotions as Attitudes", *Dialectica*, pp.
- DESCARTES, R. 1641, *Meditations on First Philosophy*, Edited by J. Cottingham, Cambridge: Cambridge University Press, 1996.

¹⁵ Thanks to Tristram Oliver-Skuse, Gianfranco Soldadi, Martine Nida-Rümelin, Davor Bodrozic, and three anonymous referees for *dialectica* for helpful comments on previous drafts of this paper.

- DÖRING, S. 2003, "Explaining Action by Emotion", *The Philosophical Quarterly*, **53**, pp. 214-230.
- DÖRING, S. 2007, "Affective Perception and Moral Motivation", *Dialectica*, **61**, pp. 363-394.
- EGAN, F. 1995, "Computation and Content", *The Philosophical Review*, **104**, pp. 181–203.
- ELGIN, C. 2008, "Emotion and Understanding", in: G. Brun, U. Doguoglu and D. Kuenzle, ed, *Epistemology and Emotions*, Farnham: Ashgate, pp. 33-50.
- FODOR, J. 1975, *The Language of Thought*, New York: Thomas Y. Crowell.
- FODOR, J. 1987, *Psychosemantics*, Cambridge, MA: MIT Press.
- FODOR, J. 1990, "A Theory of Content II", in: J. Fodor, ed, *A Theory of Content and Other Essays*, Cambridge, MA: MIT Press, pp. 89-136.
- FRANK, R. 1988, *Passions Within Reason: The Strategic Role of Emotions*, New York: W.W. Norton.
- GOLDIE, P. 2000, *The Emotions: A Philosophical Exploration*, Oxford: Clarendon Press.
- GREENSPAN, P. 1988, *Emotions and Reasons*, New York: Routledge.
- HUME, D. 1740, *Treatise of Human Nature*, Edited by L. A. Selby-Bigge, Oxford: Oxford University Press, 1978.
- JOHNSTON, M. 2001, "The Authority of Affect", *Philosophy and Phenomenological Research*, **53**, pp. 181-214.
- LAZARUS, R. 1984, "On the primacy of cognition", *American Psychologist*, **39**, pp. 124-129.
- LEWIS, D. 1974, "Radical Interpretation", *Synthese*, **23**, pp. 331–344.
- LOCKE, J. 1690, *An Essay Concerning Human Understanding*, New York: Dutton, 1961.
- LOEWER, B. 1997, "A Guide to Naturalizing Semantics", in: B. Hale and C. Wright, ed, *A Companion to the Philosophy of Language*, Oxford: Blackwell, pp. 108-126.
- MARTIN, M. F. G. 2002, "The Transparency of Experience", *Mind and Language*, **17**, pp. 376–425.
- MCDOWELL, J. 1998, *Mind, Value, and Reality*, Cambridge, MA: Harvard University Press.
- MILLIKAN, R. G. 1984, *Language, thought, and other biological categories*, Cambridge, MA: MIT Press.
- MILLIKAN, R. G. 1991, "Speaking Up for Darwin", in: G. Rey and B. Loewer, ed, *Meaning and Mind: Fodor and his Critics*, Oxford: Blackwell, pp. 151–164.
- MONTAGUE, M. 2014, "Evaluative Phenomenology", in: S. Roeser and C. Todd, ed, *Evaluation and Value*, Oxford: Oxford University Press, pp. 32-51.
- MOORS, A. 2013, "On the causal role of appraisal in emotion", *Emotion Review*, **5**, pp. 132-140.
- MOORS, A., ELLSWORTH, P. C., SCHERER, K. R., and FRIJDA, N. H. 2013, "Appraisal theories of emotion: State of the art and future development", *Emotion Review*, **5**, pp. 119-124.
- NANAY, B. 2014, "Empirical Problems with Anti-Representationalism", in: B. Broogart, ed, *Does Perception Have Content?*, New York: Oxford University Press, pp. 39–50.
- NEANDER, K. 2006, "Content for Cognitive Science", in: G. Macdonald and D. Papineau, ed, *Teleosemantics. New Philosophical Essays*, Oxford: Clarendon Press, pp. 167-194.

- NEANDER, K. 2006, "Naturalistic Theories of Reference", in: M. Devitt and R. Hanley, ed, *The Blackwell Guide to the Philosophy of Language*, Oxford: Blackwell, pp. 374-391.
- PARKINSON, B. 2013, "Comment: Journey to the Center of Emotion", *Emotion Review*, **5**, pp. 180-184.
- PEACOCKE, C. 1992, *A Study of Concepts*, Cambridge, MA: MIT Press.
- PELSER, A. C. 2014, "Emotion, Evaluative Perception, and Epistemic Justification", in: S. Roeser and C. Todd, ed, *Emotion and Value*, Oxford: Oxford University Press, pp. 106-122.
- PIETROSKI, P. 1992, "Intentionality and Teleological Error", *Pacific Philosophical Quarterly*, **73**, pp. 267-282.
- PITT, D. 2013, "Mental Representation", in: E. N. e. Zalta, ed, *Stanford Encyclopedia of Philosophy* pp. 1-41.
- PORT, R., and VAN GELDER, T. 1995, *Mind as Motion: Explorations in the Dynamics of Cognition*, Cambridge, MA: MIT Press.
- PRICE, C. 2006, "Fearing Fluffy: the content of an emotional appraisal", in: G. MacDonald and D. Papineau, ed, *Teleosemantics*, Oxford: Oxford University Press, pp. 208-228.
- PRINZ, J. 2004, *Gut Reactions. A Perceptual Theory of Emotion*, New York: Oxford University Press.
- RUMELHART, D. E. 1989, "The Architecture of the Mind: A Connectionist Approach", in: M. I. Posner, ed, *Foundations of Cognitive Science*, Cambridge, MA: MIT Press, pp. 133-159.
- SALMELA, M. 2011, "Can Emotion be Modelled on Perception?", *Dialectica*, **65**, pp. 1-29.
- SCHERER, K. R. 2009, "The dynamic architecture of emotion: Evidence for the component process model", *Cognition and Emotion*, **23**, pp. 1307-1351.
- SCHROETER, L., and SCHROETER, F. 2013, "Normative realism: co-reference without convergence?", *Philosophers' Imprint*, **13** (13), pp. 1-24.
- SCHULTE, P. 2012, "How Frogs See the World: Putting Millikan's Teleosemantics to the Test", *Philosophia*, **40**, pp. 483-496.
- SCHULTE, P. 2015, "Perceptual Representations: A Teleosemantic Answer to the Breadth-of-Application Problem", *Biology & Philosophy*, **30**, pp. 119-136.
- SMOLENSKY, P. 1988, "On the Proper Treatment of Connectionism", *Behavioral and Brain Sciences*, **11**, pp. 1-74.
- STALNAKER, R. 1984, *Inquiry*, Cambridge, MA: MIT Press, Bradford.
- STALNAKER, R. 2008, *Our Knowledge of the Internal World*, Oxford: Clarendon Press.
- STERELNY, K. 1995, "Basic Minds", *Philosophical Perspectives*, **9**, pp. 251-270.
- STICH, S. 1983, *From Folk Psychology to Cognitive Science*, Cambridge, MA: MIT Press.
- TAPPOLET, C. 2000, *Emotions et Valeurs*, Paris: Presses Universitaires de France.
- TODD, C. 2014, "Emotion and Value", *Philosophy Compass*, **9/10**, pp. 702-712.
- TRAVIS, C. 2004, "The Silence of the Senses", *Mind*, **113**, pp. 57-94.
- VAN GELDER, T. 1995, "What Might Cognition Be, if not Computation?", *Journal of Philosophy*, **91**, pp. 345-381.
- WHITING, D. 2012, "Are Emotions Perceptual Experiences of Value?", *Ratio*, **25**, pp. 93-107.
- WIGGINS, D. 1987, *Needs, Values, Truth*, Oxford: Oxford University Press.

