

Working Draft – Comments Welcome! Last Updated 2004.09.21
By Justin C. Fisher – University of Arizona, Department of Philosophy
Homepage: <http://www.u.arizona.edu/~jcfisher> E-mail: jcfisher@u.arizona.edu

Emotions as Modes of Cognition

Contents.

- I. Introduction.
- II. Ratiocination vs. Cognition.
- III. Emotions as Modes of Cognition.
- IV. Four Competing Proposals.
- V. The Impact of Emotion on Cognition.
- VI. The Kinematics of Ratiocination.
- VII. Competing Cognitive Theories.
- VIII. Why think Emotions are Beliefs?
- IX. The Intentionality of Emotions.
- X. The Kinematics of Emotions.
- XI. A Unified Account of the Emotions.
- XII. The Rationality of Emotions.

I. Introduction.

Common parlance often depicts emotions as demons which take up residence within a subject and which alter the ways in which she thinks and acts. For example, we commonly talk of subjects being *possessed by* rage, being *haunted by* grief, or being *in bad spirits*. We say that emotions *creep up on us* and *grab a hold of us* and that they then *refuse to let us go* or even *refuse to be dispelled*. Long lasting negative emotions are said to be *emotional baggage* or *demons* that we carry around with us. In the case of more positive emotions, we say that a feeling of happiness *descends upon* a subject, that love *wells up within* a subject, that a subject *becomes suffused with a spirit of awe* or of joy, or, more generally, that a subject is *in good spirits*.

There are several obvious drawbacks to the metaphor of emotions as demons. We think of demons as spooky things that enter an agent from outside, while our best empirical evidence suggests that emotions instead arise from natural mechanisms that play an integral role within a healthy animal's full cognitive system. We often think of demons as malevolent and destructive, while we accept that at least some emotions play a very positive, constructive role in our lives. But despite these drawbacks, I think that there is much to be gained by attending to this metaphor as we attempt to develop a good theory of the emotions.

I will propose that the central feature of emotions is the fact that they normally modulate the ways in which processes of ordinary epistemic and practical reasoning – processes which I will call 'ratiocination' – are performed. On this proposal, emotions really are like demons that possess us and alter the mode of our cognition. Like the hypothetical

demons, emotions draw our attention to certain things; they set priorities for us (whether we would ratiocinatively choose those priorities or not); and they introduce compelling urges in us to do various things. The goal of this paper is to spell out this ‘Modes of Cognition proposal’ and to give reasons for thinking that it is preferable to alternative understandings of the emotions.

I will compare my ‘Modes of Cognition’ proposal to three existing proposals about the nature of emotions. The central claims of the four competing proposals are as follows.

1. **Phenomenal Feels Proposal:** Emotions are certain sorts of phenomenal feelings.
2. **Somatic Proposal:** Emotions are (or essentially involve) clusters of physiological and behavioral changes that tend to be triggered together.
3. **Purely Ratiocinative Proposal:** Emotions are judgments which play a belief-like role in ordinary ratiocination.
4. **Modes of Cognition Proposal:** Emotions alter a subject’s mode of cognition and especially her mode of ratiocination.

I will argue that the Phenomenal Feels and Somatic proposals both fail to take into account the important effects that emotions have on our cognition. Emotions don’t just feel certain ways and/or make our bodies do certain things; they also impact upon what we think and how we think. To account for the centrally important effects that emotions have on cognition, one must either depict emotions as ordinary players in ratiocination (the Purely Ratiocinative proposal) or depict emotions as being otherwise capable of influencing ratiocination (the Modes of Cognition proposal).

I will then argue that there is quite strong reason to favor the Modes of Cognition proposal over the Purely Ratiocinative proposal. This discussion will raise important questions about the sense in which emotions are ‘intentional’ and ‘rational’, about whether and how emotions may be changed by socialization, and about whether one might hope for a unified account of the emotions which accommodates both immediate responses in the heat of the moment and long-term modulations in how one thinks about an object. I will argue that the Modes of Cognition proposal gives quite satisfactory answers to all these questions.

II. Ratiocination vs. Cognition

Now is a good point for several brief notes on terminology.

I will use the term ‘ratiocination’ to refer just to our introspectable deliberative processes of epistemic and practical reasoning. Hence, the paradigmatic ratiocinative elements are beliefs and instrumental desires. I will use the term ‘cognitive’ in the broad sense that it is used in ‘cognitive science’, as a way of capturing many (or perhaps all) of the interesting sorts of information processing that animals do in order to get around in the world. On this usage, everything that is ratiocinative is cognitive, but there are many

cognitive elements (e.g., elements of subconscious or non-deliberative information processing) which are not ratiocinative.¹

Using this terminology, my view is that emotions are *cognitive*, that they have essential (normal) effects on one's mode of *ratiocination*, but that they are *not*, strictly speaking, *ratiocinative*, because they do not have the sort of normal role in cognition that is possessed by the paradigmatic players (e.g., beliefs or instrumental desires) in ratiocinative processes.

III. Emotions as Modes of Cognition.

With this terminology in mind, let us now attempt to extract several important germs of truth from the demons-metaphor, even as we back away from its spookier aspects.²

First, there is the idea that, like demons, emotions are *external to* ratiocination. This idea may be supported by reflection upon one's own emotional experiences. It certainly *seems* to be the case that one does not come to have an emotion in the same way that one forms a belief or instrumental desire. Further, it certainly seems that *just* arriving at the ratiocinative conclusion that an emotion is out of place is often not enough to get that emotion to go away. For example, when I stand in high places, my *ratiocinative* processes can (seemingly) be univocally confident about my complete safety, and yet my dizzying fear of falling may persist unabated, and prevent me from thinking about other things in the comfortable way that I normally can.

These intuitive observations suggest that emotions do not have the role in cognition that is characteristic of the paradigmatic players in ordinary ratiocination. Instead, the intuitive conclusion is, emotions must be something external to the processes of ordinary ratiocination, and the *creation* and *dissolution* of emotions must be only contingently linked to ratiocinative events. At least from the perspective of ratiocination, emotions *might as well be* demons which come from outside to possess a cognitive system and alter the way in which it functions.

I think reflection on everyday experience strongly supports this intuitive conclusion. However, we will see below that there are people who do reject it, and we will consider their arguments in detail. For now, allow me to continue with an intuitive presentation of my positive view.

¹ My broad construal of 'cognitive' may be contrasted against the much narrower sense of 'cognitive' which often arises in meta-ethics, and which is intended to apply only to those representations that are truth-evaluable and belief-like. This meta-ethical use of the term 'cognitive' is better tracked by my term 'ratiocinative' than by my use of the term 'cognitive'. Similarly, Jesse Prinz's (2004, pp. 41-9) preferred use of the term 'cognitive' falls more closely in line with my use of the term 'ratiocinative' than it does with my term 'cognitive'. However, in briefly mentioning the potential defensibility of a view like mine, Prinz (2004, pg 10) appears to grant my broader usage of 'cognitive'.

² I will cite two germs of truth in the following paragraphs. A third will come to light in the discussion below, where I will note that the demons-hypothesis offers a conceptual guide towards understanding how one might give a unified account of both those emotions which last no longer than the heat of the moment, and those 'cooler' emotions that may last for long periods of time.

The demons-hypothesis also embodies a second important germ of truth, namely that having an emotion alters the manner in which a subject's cognition proceeds. Like the hypothetical demons, emotions draw our attention to certain things; they set priorities for us; and they introduce compelling urges in us to do various things. These examples give a taste for the many ways in which having an emotion might alter one's mode of cognition. As we consider the various modulating effects that emotion might have on other cognitive processes, it will be good to turn to existing work that is related to my proposal.

A fair number of existing proposals bear at least some relation to my own, including Oatley and Johnson-Laird (1987), de Sousa (1987), Frank (1988), Wright, Sloman, & Beaudoin (1996), Wollheim (1999), and Sizer (2000). However, for the present goal of cataloging the sorts of modulating effects that emotions might have, I will concentrate on the work of leading evolutionary psychologists Leda Cosmides and John Tooby. Cosmides, Tooby, and other evolutionary psychologists have developed a number of very good and very interesting ideas about the emotions – ideas that are well worth the trouble of extracting from the controversial baggage that so often comes bundled with the products of this research program.³ The following is a greatly abbreviated and charitably modified version of a quite comprehensive list of effects that Cosmides and Tooby (2000) think emotions probably have:

Goals: Emotions may influence what goals one sets for oneself, or the ways in which one goes about selecting goals from available alternatives.

Motivational priorities: Emotions may influence the relative priorities that are given to various goals, and they may influence the ways in which apparent progress or failures to progress toward a goal might affect one's inclination to continue working towards that goal.

Information-gathering motivations: Emotions may affect what one is curious about, what one finds interesting, or what one is obsessed with discovering.

Attention: Emotions may affect what sorts of things one attends to, both in perception and in ratiocination.

Perceptual mechanisms: Emotions may prompt perceptual systems to be especially sensitive to certain sorts of cues.

³ Particularly questionable is the general view that the human cognitive system consists of extremely many innate modules each of which is geared towards its own very specific purpose in a way that was quite ideally adapted to the social and environmental conditions that humans faced some number of thousands of years ago. Rather than committing ourselves to this constellation of controversial claims we should try to develop a view of the emotions that is open to alternatives. It may turn out that there are relatively few neatly identifiable modules in the human cognitive system. It may turn out that the human cognitive system is quite plastic and perhaps also significantly culturally influenced. Hence, it may turn out that there are important cognitive differences across people, across times, and across cultures, and it may turn out that human cognitive systems are reasonably well adapted to contemporary circumstances and contemporary cultures rather than being, as it were, stuck in the stone age.

Imposed conceptual frameworks: Emotions may prompt one to tend to construe the world in terms of particular concepts or categories.

Specialized inference and decision-making: Emotions may influence which sorts of reasoning schemes are likely to be applied in thinking about a problem. Some of these influences might be viewed as alterations of parameters or weightings in multi-purpose reasoning schemes. E.g., fear may license quick-and-dirty inferences which would be inappropriate ‘hasty generalizations’ at other times.

Memory: Emotions may affect which sorts of memories are potentiated for quick retrieval (e.g., those relevant to the thing that the emotion is about). Emotions may also affect the ways in which present events will be stored in memory.

Learning: Emotions may alter what someone learns from stimuli. This may be in part due to the allocation of attention, motivation, or situation-specific learning algorithms. Highly specialized learning mechanisms might also be activated, such as those that control food aversions (Garcia, 1990), predator learning (Mineka & Cooke, 1985), or fear conditioning (LeDoux, 1995).

Imagination: Emotions may affect what one imagines or the ways in which one imagines. For example, fear may make it relatively easy to imagine disastrous scenarios involving the object of fear, and relatively hard to imagine less pressing matters.

Energy level and effort allocation: Emotions may regulate one’s overall metabolic budget, as well as specific allocations to various processes and facilitation or inhibition of specific activities. The effort that it takes to perform given tasks will shift accordingly, with things being easier or more effortful depending on how appropriate they are to the situation reflected by the emotion (Tooby & Cosmides, 1990a).

Physiology/Behavior: An emotion may produce or influence the production of physiological or behavioral effects in a way that (at least ideally) is likely to be beneficial in times in which that emotion is triggered. In particular, many (but not all) emotions evoke expressions which may easily be recognized by conspecifics.

It may be worth noting that most of these potential modulating effects of emotions are purely cognitive (in the broad sense indicated above). Only the last two entries involve non-cognitive elements like (mere) behaviors or (mere) physiological changes.⁴ But, even in these cases, it is sometimes plausible to construe emotions as having these effects by modulating the flow of information from cognitive processes to these non-cognitive results. For example, fear may produce tentative, cautious behavior by altering the (cognitive) thresholds that impulses to behave must exceed before producing behavior. Hence, even a good portion of these latter two entries might be assimilating in a modes-

⁴ The ‘mere’s are there because, most (or all?) *cognitive* changes are also *physiological* changes, and, depending upon what one means by ‘behavioral’ they may also be *behavioral*.

of-cognition model. This will be in strong contrast to the view (discussed in the next section) that what is essential to emotions is their direct effects on (mere) physiology and behavior, rather than their effects on cognition.

We have just seen a long list of potential modulating effects that emotions may have on other cognitive processes. Metaphorically, this as a demon's handbook outlining the many potential ways different emotional demons might affect a cognitive system.

Quite plausibly, different types of emotions will have different clusters of characteristic causes and effects. For example, the various cases that we commonly label as 'fear' plausibly may be triggered by more-or-less reliable indicators of potential danger, and plausibly may involve modulations in all these categories in ways that are geared towards making an agent attend to and avoid a potential danger. Cases to which we commonly give other labels (like 'love', 'happiness' or 'jealousy') likely involve different characteristic clusters of triggers and of modulations on cognitive processes.

Hence, I suggest that emotions may be usefully categorized in terms of the clusters of things that will trigger them and the clusters of modulating effects that they have upon other cognitive processes. In suggesting this, I am not committing myself to the idea that, for something to count as an instance of a particular type of emotion (say, as an instance of 'fear') that it must display every single property in the cluster that is characteristic of that emotion. Instead, the categorization should demand that an instance of a category display only most (perhaps a *weighted* most) of the features in the cluster that is characteristic of that category. Like most categorizations in the life sciences these categorizations must be inherently fuzzy, and they must admit of a (hopefully not very large) number of borderline cases.

If pressed further, I would say that the categorization should not be made on the basis of *actual* triggers and effects, but instead on the basis of what *would* cause or be caused by that emotion, *if* the cognitive system *were* operating in accordance with the principles that describe normal successful cognition in that kind of system. At some risk of confusion, I will abbreviate this idea by saying that what is essential to fear is its *normal* role in successful cognition, or what characteristic triggers and modulating effects it *normally* has. This move allows that there may be pathological cases of an emotion which, in fact, are triggered in uncharacteristic ways (e.g., phobias), or which, in fact, bring about uncharacteristic effects. This general sort of move is advocated by teleofunctionalists (e.g., Millikan 1984) and by others (e.g., Lewis 1983).

In summary, my modes of cognition proposal holds that the central and most essential feature of emotions is that they normally have various modulating effects on cognitive processes, and that emotions may usefully be categorized in terms of the particular clusters of things that normally trigger them and of the modulating effects that they normally produce.⁵ This view retains the central insights of the demon metaphor – that

⁵ Below, I will suggest that there may also be useful ways of categorizing emotions in terms of their *intentional content*.

emotions are external to ratiocination but significantly impact upon ratiocination and other cognitive processes – while at the same time being empirically very plausible.

IV. Four Competing Proposals

At this point, it will be useful to contrast the view that I have just proposed against three other proposed views of the emotions.

First, there is the intuitive view that an emotion is a certain sort of phenomenal feeling that one might have. This view holds that what is essential to anger is just what it feels like to be angry. As a matter of contingent fact, it may be that anger often is caused in certain ways and often brings about certain cognitive effects, but, according to this view, these contingent facts are not in any (direct) way responsible for anger being an emotion, nor for its being the emotion that it is. This ‘phenomenal feels’ proposal is intuitively attractive. Jaak Pansepp (2000) reports that most subjects (interestingly with the exception of philosophy majors) rank feelings over thoughts and changes in the autonomic nervous system as being more important to emotions. This proposal was also influentially advocated by William James (1884) and Karl Lange (1885), who held that emotions were the experiences of various changes in one’s body.

James and Lange were also committed to a second proposal, namely that it is essential to emotions that they characteristically involve a certain cluster of physiological and/or behavioral changes. This ‘somatic’ proposal is often supported by noting research, tracing back to Darwin and more recently pursued especially by Paul Ekman, suggesting that there is a small number of such clusters of emotional displays – or ‘affect programs’ – each of which is quite universal across human cultures. The somatic proposal allows that, as a matter of contingent fact, anger may typically be caused in certain ways, may typically have a certain phenomenal feel, and may bring about various cognitive effects. But, according to (a pure version of) this proposal, these contingent facts are not in any way responsible for the fact that anger is an emotion, nor for its being the emotion that it is. Versions of the somatic proposal were defended by behaviorists (e.g., Watson 1919, Ryle 1949, Skinner 1962), and more recently by Griffiths (1997) and Rolls (1999). One may also interpret Jesse Prinz’s (2004) ‘embodied appraisal’ account of emotions as an (impure) form of a somatic proposal, as it holds that (perceptual registerings of) bodily states are essential to emotions.

Third, there is the view that what is essential to emotions is their having (or their being caused by things that have) some particular sort of belief-like or desire-like role as elements in ratiocination. This ‘purely ratiocinative’ proposal grants that emotions may often have certain phenomenal feels, certain physiological effects, and even certain effects on non-ratiocinative cognition, but it holds that these accidental features are not directly responsible for these emotions’ being the emotions that they are.

The purely ratiocinative proposal may be traced back to Aristotle and the stoics. Its recent defenders have included Robert Solomon (1999) and Martha Nussbaum (2001), who suggest that emotions are belief-like judgments regarding the value of an object within the scheme of things that one cares about.

This proposal also encompasses the predominate understanding of emotions among psychologists which depicts emotions as essentially being *products of* certain sorts of ratiocinative appraisals, in much the way that a sunburn is essentially the product of a certain sort of process involving the sun (Gordon 1987). Versions of this view are defended by Arnold (1960), Lazarus (1991), among many others. Like the view (more often defended by philosophers) that emotions *are* a certain sort of belief, the view (more often defended by psychologists) that emotions *are essentially caused by* a certain sort of belief is committed to saying that the presence, absence, and essential consequences of emotions are closely linked to the kinematics of ordinary ratiocination. In what follows, I will concentrate upon the philosophers' formulation of the purely ratiocinative view (as a thesis about the identity of emotions), but most of what I say about it may be adapted, *mutatis mutandis*, to the psychologist's version (as a thesis about the essential causal origins of emotions).

In summary, the central claims of the four proposals are as follows.

1. **Phenomenal Feels Proposal:** Emotions are certain sorts of phenomenal feelings.
2. **Somatic Proposal:** Emotions are (or essentially involve) clusters of physiological and behavioral changes that tend to be triggered together.
3. **Purely Ratiocinative Proposal:** Emotions are judgments which play a belief-like role in ordinary ratiocination.
4. **Modes of Cognition Proposal** (briefly sketched above): Emotions alter a subject's mode of cognition and especially her mode of ratiocination.

These four proposals amount to four sorts of ingredients that a theorist might take to play an essential role in an account of the emotions. There has been an understandable temptation among many theorists to choose just one of these ingredients and to attempt to spell out a general account of the emotions just in terms of it.

However, one needn't be committed to saying that *just one* of these four ingredients is the only essential one. For example, we have seen that James and Lange held that it was essential to emotions both that they involve various somatic changes *and* that they involve the phenomenal experience of those changes. Alternatively, one might follow Aristotle and hold that an emotion is essentially a belief that has both an appropriate content *and* an appropriate phenomenal feel. For a nice survey of various hybrid views see Prinz (2004, Chap 1).

An alternative way of combining different sorts of ingredients would be to say that there are several distinct classes of phenomena which have been lumped together under the heading 'emotion', and that different ingredients are required to account for each class. For example, Paul Griffiths (1997) suggests that one important class of emotional phenomena is well accounted for by the Somatic proposal, while a distinct class of phenomena which have also been labeled as 'emotions' might be reasonably well

accounted for by some form of the Modes of Cognition proposal.⁶ Insofar as both of these classes deserve the label ‘emotions’ (a point about which Griffiths wavers), Griffiths’ suggestion is that a full bipartite account of the emotions must make quite separate references to physiological and to ratiocinative elements.

Additionally, one might attempt to propose some fifth ingredient, and hold that any proposed account of the emotions must be incomplete without making essential use of this fifth ingredient. However, I don’t know of any plausible candidates.

V. The Impact of Emotion on Cognition

These four proposals may be naturally divided into two categories. The Purely Ratiocinative and the Modes of Cognition proposals are quite *cognitivist* – they both take it to be essential to emotions that they play a particular sort of role in the processing of information in the cognitive system, and especially that they play a particular sort of role *vis a vis* ratiocination. In comparison, the Phenomenal Feels and Somatic proposals are (relatively) non-cognitivist – they both (seem to⁷) allow that whether or not an agent has an emotion might vary quite independently of the sorts of information processing that are going on in her cognitive system.

This categorization reveals a potential ground for favoring one pair of proposals over the other. Insofar as we have good reason to think that emotions are essentially cognitive, we will have good reason to reject the Phenomenal Feels and Somatic proposals.

I am inclined to say that we may unproblematically reach this conclusion simply by reflecting on our pre-theoretic expectations of a theory of emotions. In designating the paradigmatic cases that such a theory must account for, we designate essentially cognitive phenomena. For example, in pointing towards paradigm cases of fear, we pre-theoretically can say that, if any state deserves to be called ‘fear’, then it must (essentially) be the sort of state that is normally responsible for producing characteristic cognitive effects: the urge to escape the situation, the tendency to imagine horrible things happening, and the difficulty of concentrating upon unrelated matters. If a proposed theory of fear relegates these cognitive phenomena to an insignificant side role, then, I am inclined to say, it doesn’t even deserve to be called ‘a theory of fear’.

Despite its intuitive force, the preceding argument will likely strike my opponents as a bald-faced attempt to beg the question. For this reason, it is worth noting that there are alternative strategies that the cognitivist might attempt to pursue here. One such strategy would be to seek empirical evidence that phenomenal and somatic factors can vary significantly across the various paradigm cases of a given emotion. Hence, Nussbaum

⁶ Although Griffiths does not consider my Modes of Cognition proposal, it seems to me that this proposal is reasonably consistent with the general approach that Griffiths takes towards these issues, and it resonates very nicely with some of Griffiths comments (eg., pg 121).

⁷ This qualifier is here because there are many competing accounts of ‘phenomenal feels’, and some of these accounts hold that one’s phenomenal feels are determined in part by the roles that the various parts of one’s cognitive system are playing. While such views may end up having implications regarding cognition (information-processing), this must be by a round-about and currently-quite-unclear route.

(2001, pp 56-64) suggests that fear may have quite different phenomenal feels on different occasions, or that one may be in fear without having any sort of distinctive phenomenology at all. Similarly, she suggests, instances of a particular emotion may have quite different physiological effects in different subjects or at different times in a given subject.

These empirical claims can be intuitively supported, at least to some extent, by reflection on one's own experience. I suspect that Nussbaum is right that we will discover that most interesting physiological, behavioral and phenomenal classifications cross-cut our pre-theoretic or folk-theoretic classifications of emotions; and hence that we will end up feeling little pull to attach our folk-labels to physiological, behavioral, or phenomenal kinds.

However, it's not *clear* to me that these empirical claims will actually be borne out. For all I know, it may turn out that fear really does regularly produce a certain cluster of (perhaps quite subtle or context dependent) physiological effects. And, for all I know, it may turn out that emotions really do have characteristic phenomenal feels, albeit ones that I'm not very good at recognizing. And even if it happens to turn out that emotions do typically have these features, I think we *still* should say that these aren't the sorts of features that are essential to emotions.

Compare: for all I know, it may turn out that whenever I get angry, some particular bulb lights up in some secret CIA office – but certainly that bulb's lighting up is not an essential feature of my anger. There is a strong inclination to say that (mere) somatic effects and phenomenal feels should have pretty much the same role in a theory of emotions as the CIA's light bulb. What's centrally important about emotions is not which muscles they might cause to contract, which phenomenal buzzes they might invoke, or which bulbs they might light up – instead what's centrally important about the emotions is that they change what we think and how we think. I think it is *obligatory* that any theory of the emotions must give center stage to these cognitive phenomena. The problem with the Phenomenal Feels and Somatic proposals is that they don't fulfill this obligation.

That is my 'hard-line' position. But I also recognize that there is a more conciliatory position available. All parties may agree that it would be nice to give a good account of the various clusters of physiological and behavioral changes that people undergo. All parties may agree that it would be nice to develop a theory of what phenomenal feelings people have in various circumstances. And all parties may agree that it would be nice (*really nice!*) to have a good theory of the different ways that cognition proceeds in various sorts of paradigmatically emotional phenomena. Given that we all agree that it would be nice to develop all these theories, we needn't begrudge any of the theorists who are working to develop them.

Still, I maintain that the most important phenomena to be accounted for here – the ones which will have the greatest impact upon questions about how we should live our lives, about how we should treat one another, and about how we should develop our social policies – are the phenomena that involve the effects that emotions have on cognition.

We want to know why unrequited love makes it so hard for us to get on with our lives, not why it makes our hearts go pitter-pat. We want to know whether homophobic rage impairs people's decision-making ability in a way that reduces their moral responsibility for their actions, not what this rage feels like. We want to know the numerous sources and cognitive consequences of misogynistic disgust, not the physiological and phenomenal changes that accompany it. For these reasons, I won't very easily back down from hard-line stance that a 'real' account of the emotions must centrally address these cognitive issues, and that any proposed account that fails to do this will thereby count as an account of the 'emotions' only in some quite attenuated sense.

Hence, I conclude that, at least for the purposes that I maintain are most important, we may lay aside the Phenomenal Feels and Somatic proposals. To account for the centrally important effects that emotions have on cognition and especially upon ratiocination, one must instead embrace some *cognitive* theory of the emotions.

VI. The Kinematics of Ratiocination

There is a wide range of logically possible cognitive theories of the emotions. I will approach these issues by contrasting two rather extreme positions – the Purely Ratiocinative proposal and an (overly) extreme version of my own modes of cognition proposal. These extreme positions stake competing claims regarding the sort of role that emotions play in ratiocination. To make clear the differences between these views, it will be useful to say a bit more regarding ratiocination and the patterns of interaction among beliefs and desires within ratiocination – patterns which I will call the *kinematics* of ratiocination.

I take two things to be essential to ratiocinative processes. First, these processes must involve representational tokens – 'beliefs' – whose normal job in cognition is to correspond to contingent states of affairs in one's environment, and other representational tokens – 'desires' – whose normal job in cognition is to tend to bring it about that certain sorts of states of affairs obtain.⁸

⁸ A notion of 'normal jobs' (or of 'natural purposes' or 'proper functions') is central to *teleofunctionalist* understandings of cognition (e.g., Millikan 1984). This notion allows that, in many cases, a thing may fail to successfully perform its 'normal job', even while it still is the case that making reference to that thing's normal job is essential to understanding that thing's role within a cognitive system.

I should perhaps note that, when push comes to shove, the notion of 'normality' being called upon here, and throughout this paper, isn't simple statistical normality, but rather (something like) accordance with general principles that explain how success was achieved *in past success cases*, where success is taken to be promoting the replication (and hence evolution) of similar entities (again see Millikan, 1984). Hence, a 'normal job' of a sperm is to fertilize an egg because that is how historically successful sperm have secured their evolutionary success, despite the fact that historically, only a vanishingly small fraction of all sperm have successfully performed this task.

A related (but in my opinion much less plausible) view is the popular *functionalist* view which holds that we should talk about the role that these tokens *actually are* disposed to play within a cognitive system, rather than about what role they *would play* in it *if things were going normally*. The committed functionalist may, with relatively little difficulty, translate my talk of 'normal jobs' into a form that is more acceptable to mainstream functionalism.

Second, it must be the case that these processes normally display a particular sort of *kinematics* – a particular sort of pattern in the ways that beliefs and desires normally are generated, retained or abandoned, and in the ways in which they interact with other cognitive elements. The kinematics of ratiocination quite closely reflects certain relations between the contents of various representations.⁹ For example, when one forms a belief on the basis of other beliefs, there is normally an appropriate relation of support between the content of the other beliefs and the content of the belief one forms. Similarly, whenever one actively considers beliefs whose contents are contradictory, one normally abandons one or the other of these beliefs (or both). Similarly there are patterns in the normal formation, retention, and abandonment of instrumental desires, and these patterns also track relations between the contents of the various representations in question.

It is the kinematics of ratiocination that we (at least attempt to) report upon when we adduce reasons for the various beliefs and instrumental desires that we have. It is likely that humans are the only Earthly species with the capacity to engage in long trains of ratiocination regarding abstract topics, or with the capacity to give overt reports about these processes. However, it is at least plausible that many non-human animals engage in processes that are more or less like short chains of human ratiocination. Hence, one might claim that emotions are essentially ratiocinative elements without committing oneself to denying that animals have states that deserve to be called emotions too.

VII. Competing Cognitive Theories.

Given the above notions of ratiocination and its kinematics, we may now state two extreme cognitive proposals regarding the nature of emotions.

At one extreme is the *purely ratiocinative proposal*, which holds that emotions are a special variety of one of the paradigmatic players in ratiocination – e.g., of beliefs, of desires, or perhaps of some complex combination of beliefs and desires. On this proposal, what distinguishes emotions from other ratiocinative players is the particular sort of content that emotions have; as one commonly taken line would have it, the characteristic content of emotions involves an evaluation of the role that some object might play within the scheme of things that one values. Finally, the purely ratiocinative proposal holds that the kinematics of emotions – the patterns in the ways that emotions interact with one another and with other cognitive elements – *just is* the kinematics of the paradigmatic players in ratiocination. I.e., emotional states are formed in the same sorts of ways that ordinary beliefs or instrumental desires are formed; emotional states impact upon cognition just by serving as premises, goals or defeaters in the ratiocinative processes by which further beliefs or instrumental desires are formed or abandoned; and emotional states are shed in the same sorts of ways that ordinary beliefs or instrumental desires are shed.

At the other extreme is the *extreme modes of cognition proposal*, which holds that emotions are best understood as states that arise externally to ratiocination, and that

⁹ The representations in question may include ratiocinative beliefs and desires, but they may also include non-ratiocinative elements that impinge on ratiocination, like perceptual representations or intuitions.

modulate the ways in which cognition (including ratiocination) proceeds. On this extreme proposal, what distinguishes emotional states from non-emotional states is not their representational content, but instead it is their kinematics – the various sorts of causes and effects that emotional states normally have within cognition. In an extreme form, this proposal holds that this kinematics is very much different from the kinematics of the paradigmatic players in ratiocination. For example, it may be essential to certain emotions that they normally draw our attention to certain things, set our priorities, introduce compelling urges in us to do various things, affect which sorts of memories are most easily accessible, affect which rules for inference, learning, or decision we employ, affect which sorts of features our perceptual systems will be geared towards picking out, and/or alter which sorts of scenarios we may easily imagine.

Between these two extremes there is a wide range of possible intermediate positions. A more moderate version of the purely ratiocinative proposal might admit that it is essential to some emotional states that they have a kinematics that, in fact, differs in at least some ways from the kinematics of ordinary beliefs and desires. A more moderate version of the modes of cognition proposal might admit that the representational content of some emotional states plays a role in making them be the emotional states that they are, and/or admit that some emotional states either *themselves* play a belief-like or desire-like role in ratiocination, or else generate beliefs or desires which themselves play a belief-like or desire-like role in ratiocination. A *very* moderate version of one of these extreme proposals might be indistinguishable from a very moderate version of the other.

Above I suggested that one might find intuitive motivation for a good theory of the emotions by attending to common metaphors that depict emotions as spirits or demons that possess an agent and alter the ways in which her cognition proceeds. This metaphor clearly points us towards some form of the modes of cognition proposal. But in particular, it points us towards a moderate, rather than an extreme, version of this proposal. For demons clearly might generate beliefs or desires which might themselves play a belief-like or desire-like role in ratiocination, and there is also an intuitive sense in which a demonic possession might be *about* some object (e.g., an object that caused a demonic possession to begin, or an object that the demon draws attention to or causes one to imagine things about). Furthermore, there is an intuitive sense in which a demonic possession couldn't be the sort of demonic possession that it is, if it weren't of a sort that normally generates certain sorts of beliefs or desires, or if it weren't about the object that it is about. Hence, the demons-metaphor provides intuitive motivation for a *slightly moderate form* of the modes of cognition proposal.

As will become clear in what follows, I think that the best theory of emotions will indeed be a slightly moderate form of the modes of cognition proposal – one which takes on much more of the modes of cognition proposal than have most existing theories of the emotions. I propose that emotions may be usefully categorized on the basis of (1) the clusters of things that normally trigger them, (2) the clusters of modulating effects that they normally have upon other cognitive processes, and (3) the sorts of intentional content that they normally have.¹⁰ As is the case with most categorizations in the life

¹⁰ Again, I think the appropriate reading for “normally” is a teleo-functionalist reading.

sciences, we should expect this categorization to yield classes with somewhat fuzzy boundaries.

VIII. Why think Emotions are Beliefs?

Before I argue in defense of a (slightly moderate) modes of cognition proposal, it will be good to look at some reasons that might lead people to embrace a purely-ratiocinative proposal instead. This will serve not only to allay suspicions that I am arguing against a straw man, but also to broach some important topics that any theory of the emotions must address. In later sections I will argue that a modes of cognition proposal can satisfactorily address each of these topics.

For the sake of brevity, I will concentrate our attention upon just the version of the purely ratiocinative proposal that has been most strongly defended by philosophers, the version that holds that emotions are a sort of belief (Solomon 1999, Neu 2000, Nussbaum 2001). I will not explicitly consider other purely ratiocinative views that instead hold emotions to be a sort of desire or a sort of complex combination of beliefs and desires (e.g., Marks 1982; Robinson 1983; Roberts 1988; Oakley 1992). However, many (although perhaps not all) of the considerations that arise regarding the view that emotions are a sort of belief would arise in a similar form regarding these alternative views.

A good starting point is the apparent *intentionality* of emotions. We commonly, and seemingly quite successfully, attribute various sorts of *intentionality* or *aboutness* to our emotions. We are jealous *of* people; we are angry *at* or *with* them; and we feel love *for* them. We are afraid *that* certain events might happen. We are disgusted *by* various objects. We are elated *about* events that go well, and grieve *over* our losses. A purely ratiocinative proposal has a ready explanation for this apparent intentionality: emotions are a sort of belief, and beliefs are, of course, paradigmatic bearers of intentionality.

A second point that might seem to weigh in favor of a purely ratiocinative view is the fact that, many emotions seem to have a kinematics that is, at least at first blush, quite like that of beliefs. For example, all parties will agree that in most cases, someone who is afraid of a particular object will behave in roughly the same ways as someone who believes that object is dangerous, and someone who loves a person will behave in much the same ways as someone who believes that the well being of that person is of crucial importance to her valued projects. These observations leave open the possibility that emotions are unlike ordinary beliefs in many other important ways (as I will soon argue that they are), but these observations do lend at least some initial intuitive credibility to the claim that the emotions in these cases *just are* beliefs.

Similarly, all parties may agree that, at least in many cases, a subject enters into an emotional state soon after she gets what would be good evidence for a corresponding belief. For example, a subject comes to fear, say, poly-unsaturated fats soon after she gets good evidence that poly-unsaturated fats pose a danger to her. Again, such observations leave open the possibility that many emotions come about in ways that are

importantly different from the way that beliefs come about, but still these observations do lend some initial credibility to the view that emotions are beliefs.

A third point that might seem to favor a purely ratiocinative proposal is that it can tell a coherent story regarding the ways in which our emotional reactions change as we gain experience and wisdom. New experiences offer us new evidence regarding the sorts of roles that various things may play in our lives. When emotions are taken to be a sort of belief, it makes sense that this changing base of evidence may lead to changes in emotions. And insofar as our growing base of evidence tends to give us increasingly accurate beliefs about the world, this view may tell a story about how the emotions may be 'educated' to better fit the world.

Similarly, a modes of cognition view has a story to tell about cross-cultural differences in emotions. Different sorts of evidence are available in different cultures, and hence, by the kinematics by which people form beliefs of the basis of evidence, people in different cultures form different emotions.

A fourth point that might attract people to a purely-ratiocinative proposal is that it promises to offer a unified account of both the brief, 'hot' emotions that arise only for the duration of the heat of the moment, and 'cooler' longer-lasting emotions that may stay with us for months or years. For there seems to be some important sense in which there is a unified account of both those 'occurrent' beliefs that briefly play a prominent and active role in our thinking, and those 'non-occurrent' or 'dispositional' beliefs that we hold for very long periods of time and that only subtly or occasionally have significant influences on our thoughts. A purely ratiocinative proposal promises to spell out a unified account of emotions as a corollary of a unified account of beliefs.

A fifth point is that the purely ratiocinative proposal offers to make sense of the intuitive idea that some emotions are rational and others are irrational. If emotions are a sort of belief, then they may depict the world as it is or they may fail to do so, and they may be well-founded or ill-founded. Hence, the rationality of the emotions may be construed as a special case of the rationality of beliefs.

In the coming sections, I will return to each of the above issues, and argue that a modes of cognition proposal gives a treatment of these issues that is at least as good as (and often much better than) the one given by a purely ratiocinative proposal.

IX. The Intentionality of Emotions.

A (slightly moderate) modes of cognition view allows that there are several different senses in which emotions have intentionality.

For one thing, this view allows that it may be essential to emotions that they interact in various ways with other representations in a cognitive system. When an emotion has essential links to representations that are about a particular object, there is then a derivative sense in which that emotion is about that object too. For example, if an emotion is normally caused by beliefs about some object, and/or if this emotion normally

tends to shift cognitive resources towards doing processing that is itself about that object, then there is a sense in which that emotion is about that object too.

It may be worth noting that this sort of intentionality does not involve any *direction of fit* between emotions and the world. This sort of intentionality is more like the intentionality of a name or a concept than the intentionality of a belief or desire. Things that have this sort of intentionality may be *about* objects in the world, but they needn't represent those objects as being a certain way, nor need they tend to bring it about that those objects become a certain way.

A modes of cognition view also allows that many emotions *do* have a direction of fit. A token has a *belief-like* or *mind-to-world* direction of fit if it is a normal job of the mechanisms that produce that token to produce tokens in a way that tracks relevant aspects of how the world is in order to enable further mechanisms to successfully perform their normal jobs. E.g., so long as belief-fixing mechanisms are doing their normal jobs, one will normally form the belief that the cat is on the mat only when the cat is on the mat, and this belief may then lead to the success of further reasoning and behavior. A token has a *desire-like* or *world-to-mind* direction of fit if a normal job of that token in cognition is to tend to bring it about that the world is a certain way.¹¹ E.g., in normal cognition, a desire that the cat be on the mat has it as a normal job to tend to bring it about that the cat is on the mat.

It is quite clear that many emotions have a belief-like direction of fit. E.g., fears are normally generated in response to danger and bring about appropriate responses to danger; and instances of grief are normally generated in occasions where one has suffered a significant loss, and they bring it about that one thinks through many of the significant consequences that such a change in circumstances has on one's projects.

In other instances, it seems most plausible to attribute only a *desire-like* direction of fit to emotion states. For example, it is implausible that it is a normal job of rage-states (or, more properly speaking, of the mechanisms that generate rage-states) to track some aspect of the world. Instead it seems that rage states may normally be brought about in a fair number of different ways in a fair number of different circumstances. It is plausible to suppose that a normal job of a rage state is to bring it about that harm comes to the 'object' of one's rage. This certainly seems to be a consequence that rage is often geared towards, and further there are good design reasons to have a cognitive system sometimes go into states that are geared towards doing this – such states may do well to fend off direct threats and/or to deter future aggressions, and the mere fact that an agent is disposed to become enraged makes that agent a less attractive target for exploitation by others (Pinker 1997). In light of these observations, it is more plausible to attribute to rage a desire-like direction of fit than a belief-like one. Similar considerations suggest that *love* primarily has a desire-like direction of fit – the normal job of love is (at least plausibly) to bring about and preserve the well-being of its object.

¹¹ These proposed definitions owe much to Ruth Millikan's (1984) work.

One common response to suggestions along these lines is to say that these emotions do have a belief-like direction of fit: love depicts an object as being such that its well-being is central to one's valued projects, while rage depicts an object as being such that its being damaged would be central to one's valued projects. If such a move is acceptable, we might reduce *all* desires to beliefs: any desire for X is tantamount to a belief that X is 'good' or that X would fit nicely in the scheme of things that one values. There may be good reasons to resist this sort of move.¹² However, even if we allow this move, we should note that there is still an important distinction to be made between tokens that attempt to reflect facts about *how certain value-free aspects of the world are* (e.g., how the cat is positioned with respect to the mat) and tokens that attempt to reflect facts about *how it would be good for the world to be*. Hence, even if we allow this move, we still must end up agreeing that rage and love normally work only in the latter way, while many beliefs and many other emotions work in the former.

Interestingly, there are some emotional states which have *both* a belief-like and a desire-like direction of fit – they say that the world is one way, and at the same time they try to bring it about that the world is some other way. E.g., it is plausible that immediate fear states have it as a normal job *both* (1) to indicate to other cognitive mechanisms that some object *is* a pressing danger, *and* (2) to bring it about that that object no longer poses a pressing danger. Hence, there is a sense in which an immediate fear state represents an object as being dangerous; and there is another sense in which the very same fear state represents a desirable state of affairs, one in which the object no longer poses such a danger. Ruth Millikan (1996) calls representations like these 'pushmi-pullyu' representations, and suggests that such representations are the more primitive precursors of more specialized representations that have only one direction of fit. On a modes of cognition view, it is not at all surprising that many emotional states would have this dual-nature. As we will now see, this dual nature is more problematic for a purely-ratiocinative view.

X. The Kinematics of Emotions.

One primary difference between the two extreme proposals stated above involves the kinematics of emotions, the normal patterns of interaction between emotions and other cognitive states. The purely ratiocinative proposal held that the kinematics of emotions just is the kinematics of paradigmatic players in ratiocination, beliefs and instrumental

¹² For one set of reasons, see Lewis (1996), who argues that the rational kinematics of beliefs differs from that of desires, and hence concludes that no single state may display both sorts of kinematics in ratiocination. Another reason to resist this move may be drawn from the work of various authors (Millikan 1984, Dennett 1975, Davidson 1973) who suggest that attributions of representational content must be guided by considerations that mirror those that constrain good explanations; there is a real concern that attributing a belief-like direction of fit to a desire does not deliver any explanatory mileage beyond what is delivered by attributing to it a desire-like direction of fit, and hence, for reasons of parsimony, we should attribute only a desire-like direction of fit to such states. Notice that the case under consideration (where we're asking whether a desire that P might also be well-construed as a belief that *it would be good that P*) differs from the pushmi-pullyu cases discussed in moment (where we're noting that a single state can serve both as a desire-like representation of P and, at the same time, a belief-like representation of Q, where both P and Q may be expressed in value-free terms).

desires. A modes of cognition proposal denies this, and holds that the kinematics of emotions differs significantly from that of beliefs and instrumental desires.

Let us now take up the question of which of cognitive proposal does best to capture the kinematics of emotions. My arguments in this section will resonate with those of a number of other authors who have argued that the rational kinematics of emotions is not reducible to the rational kinematics of beliefs. (De Sousa 2003 attributes such arguments to de Sousa 1987; Ben-Ze'ev 2000; Goldie 2000; and Elster 2003.)

As a starting point, we may consider the relation between kinematics and direction of fit. We have seen that many emotions have a belief-like direction of fit – they are produced by mechanisms whose normal job involves being sensitive to information about how the world is and producing emotional states that will then convey that information to systems that will make good use of it. A purely ratiocinative view that takes emotions to be a sort of belief will likely at least approximate the kinematics of these emotions.

However, we have also seen that some emotions have a desire-like direction of fit – they have it as a normal job to bring about changes in how the world is, and not necessarily to track how the world in fact is. The rational kinematics of desire-like states is often taken to be different from that of belief-like states (Lewis 1996). Hence, it seems, the purely ratiocinative view that emotions are beliefs will probably fail to accommodate the kinematics of emotions that have a desire-like direction of fit.

More complicated, still, are the pushmi-pullyu emotional states which have both a belief-like and a desire-like direction of fit. It would be quite surprising if such states have just the kinematics of ratiocinative beliefs. It might be more promising to construe these as purely ratiocinative complexes involving both beliefs and desires as components (as is proposed by Marks 1982 and Oakley 1992). However, there will still be many hard challenges for such a view, including Lewis' (1996) argument that it is impossible that a single state have *both* the ratiocinative kinematics of a belief *and* the ratiocinative kinematics of a desire. I think the most plausible conclusion is to deny that the kinematics of emotions reduces to the kinematics of beliefs and/or desires, and to instead embrace the flexibility in describing the kinematics of emotions that is afforded by a modes-of-cognition view.

This conclusion may also be supported in other ways. We may begin by considering our catalog of modulating effects that emotions might have upon cognition. There is a fair amount of empirical evidence supporting the claim that certain emotions normally have effects of these sorts (see, e.g., Cosmides & Tooby, 2000). Furthermore, reflection upon the usefulness of emotions as an adaptive 'design feature' in cognition makes plausible the claim that emotions would have many effects of these sorts; and introspective reflection upon our own emotional experiences seems to reveal that emotions do have many of these effects on our cognition.

Take, for example, a fear of a coiled snake. Introspectively, this fear certainly seems not to be a mere premise that might be combined with other premises to yield further conclusions or combined with beliefs and desires to yield further instrumental desires.

Instead this fear seems to be a state that (at least normally) demands that the snake be attended to; that makes memories of snakes readily available while making memories of unrelated matters very difficult to recall; that allocates various cognitive resources towards this snake and away from distant or unrelated matters; and that floods one's imagination with images of what the snake might do. Furthermore, it seems to be *essential* to the fear that it is the sort of state that is at least normally poised to grip one's cognition in these ways. A representation that is poised *only* to serve as a premise in ratiocination and that is not poised to grip one's cognition in these other ways seems *therefore* probably not to be an instance of fear at all.¹³

Putting the same point more generally (and borrowing a phrase from Proust), there is a strong intuition that emotions are not thoughts; instead they are *upheavals in thought*.¹⁴ The characteristics that make emotions *upheavals* are characteristics that make the kinematics of emotions significantly different from the kinematics that is characteristic of beliefs and desires in ratiocination. If emotions are a sort of belief, then they must be a sort of belief whose kinematics is quite different from the characteristic kinematics of beliefs in ratiocination.¹⁵

The problems for a purely ratiocinative view continue to mount as we compare the ways that ratiocinative beliefs are normally generated and abandoned to the ways in which emotional states are normally generated and abandoned.

The kinematics of ratiocination is such that having in mind a strong line of reasoning supporting a conclusion (and having an interest in determining whether or not this conclusion holds) normally generates a belief in this conclusion. It certainly *seems* to be the case that one does not come to have an emotion in the same way. I can *ratiocinatively* recognize how very important someone or some goal is towards projects that I hold dear – at least I can be fully poised to use such claims as premises in further ratiocination – and yet I can fail to be emotionally engaged. A classic example is reasoning to the conclusion that someone would fit wonderfully into all the projects that you value, and yet not falling in love. A good argument is normally sufficient to get one to believe its conclusion, but there is *no* sort of good argument that is normally sufficient to get one to fall in love.¹⁶ Hence, the kinematics of love must be different from the

¹³ At best, this state might be an abnormal, pathological instance of fear – the sort of fear that one imagines vicious sociopaths to have. A teleofunctionalist view does well to accommodate the possibility of such pathological cases. See also Lewis (1983).

¹⁴ One might find it ironic that Martha Nussbaum – a leading defender of the view that emotions are a sort of thought – chose to entitle her (2001) book “Upheavals of Thought”.

¹⁵ I am remaining neutral regarding how many representations with belief-like direction of fit we want to count as ‘beliefs’. Nussbaum (2001) sometimes states one of the central tenets of her view as being that it is both necessary and sufficient for having a particular emotion that one have a belief (or judgment) with an appropriate content. If we construe the term ‘belief’ very broadly, then the necessity-half of this claim is plausibly true, at least for many emotions. However, I think the sufficiency-half of this claim is false – for a state to be an emotion, it is not sufficient that it be a belief with the right content; for, additionally, it must be a state which, at least normally, modulates cognition in various upheaval-like ways, and this is something that a belief with that content needn't normally do. (This is intuitively illustrated by the snake-fear case described above, and by the arguing-someone-into-love case below.)

¹⁶ One potential objection might go as follows: “Sure, you can hear a compelling argument for P (e.g., that someone would fit wonderfully into all the projects you value) without thereby falling in love. That just

kinematics of ratiocinative beliefs. And hence, if love is a ‘belief’ at all, it must be a different sort of belief from the beliefs that play in ordinary ratiocination.

Similar observations apply to the kinematics surrounding the discovery of contradictory beliefs. If one discovers oneself to have contradictory beliefs, then (barring cases of extreme insanity) one abandons one belief or the other, or both.¹⁷ In contrast, it certainly seems that discovering that a ratiocinative belief runs counter to an emotional state is often not enough to get either that belief or that emotion to go away. For example, when I stand in high places, my ratiocinative processes can (seemingly) be univocally confident about my complete safety, and yet my dizzying fear of falling may persist unabated, and prevent me from thinking about other things in the comfortable way that I ordinarily can. I can attend quite carefully both to my belief that I am safe and to my fear of falling, without either of these seeming to be weakened by the other. This suggests that the kinematics of emotions must be different from the kinematics of ratiocinative beliefs. When one ratiocinative belief runs counter to another, and both are attended to, at least one of these must be defeated (on pain of great irrationality); but not so when an emotion runs counter to a ratiocinative belief.

More generally, I cannot remember any experience in which it seemed that I did some reasoning and arrived at or abandoned *an emotion* as a conclusion. Instead, it seems, my emotions consistently (sometimes even obstinately) lag behind my reasoning. I maintain my attention on certain topics – the ones which I’ve found to do pretty well at enticing my emotions – and then, if I’m lucky, my emotions eventually follow suit.¹⁸ Insofar as it is likely that this introspective report is correct and representative of others’ emotions,

shows that P wasn’t the relevant content for falling in love. To fall in love, one must instead come to believe some other content Q.” One plausible reaction at this point is to demand a specification of *which* content Q is supposed to be sufficient – a demand that I doubt will ever be met. But even without waiting for a particular specification of Q, we may note that there is a strong intuition that *whatever* content Q the objector settles upon, a good argument for Q may compel a rational subject to believe Q, but needn’t compel that subject to fall in love. It often takes more than just a compelling argument (even one for a very fancy conclusion) to get someone to fall in love.

A possible rejoinder would be to deny that the characteristic content of emotions is such that one could come to believe it in more prosaic ways. De Sousa (2003) attributes versions of this rejoinder to Goldie (2000), Wollheim (2000), Charland (2002) and Tappolet (forthcoming). Whatever merits this rejoinder might have, it is worth stressing that this rejoinder involves a significant departure from the purely ratiocinative proposal described above. For this rejoinder holds that emotions have content quite different from the sorts of content that play into ratiocinative relations, and hence this rejoinder must deny the purely ratiocinative proposal that the kinematics of emotions just is the kinematics of ratiocination.

¹⁷ It may often be, however, that a defeated belief is not completely expunged from one’s long-term memory. Hence, one may repeatedly find oneself believing something that has been defeated many times, as Nussbaum (2001, pp. 35-6) reports that she repeatedly finds herself believing that the U.S. Supreme Court is in California. What one cannot do, however, on pain of great irrationality, is retain a pair of contradictory beliefs in ratiocination while carefully attending to each of them. I am confident, for example, that whenever Nussbaum attends to her contradictory beliefs about the location of the Supreme Court, her belief that it is in California is quickly defeated, and prevented (at least for the time being) from playing a further role in her ratiocination.

¹⁸ My introspective report is picturesquely echoed in Plato: “Your passion boils and rages, and fights for what you regard as right... [O]nly success or death can stop it fighting the good fight, unless it is recalled by your rational mind and calmed down, as a dog is by a shepherd” (Plato, Republic 440c).

this is an intuitive reason to conclude that emotions are not ratiocinative beliefs, and that their kinematics differs from that of ratiocinative beliefs.

Thus far my argument in this section has been largely negative. I have offered a number of compelling reasons for thinking that the kinematics of emotions must be quite different from the kinematics of ratiocination. Since different emotions have different directions of fit, and some even have both directions of fit at the same time, it is unlikely that emotions have just the kinematics of ratiocinative beliefs. It seems to be essential to emotions that they are *upheavals in thought* – that they have effects quite different from the effects of ordinary beliefs. Furthermore, it seems that emotions are not generated in the same way that ratiocinative beliefs are generated; and it seems that they cannot always be defeated by countervailing beliefs in the way that ratiocinative beliefs almost always can. All of these considerations weigh in favor of a modes of cognition proposal over a purely ratiocinative proposal.

Eventually, it would be good to have a positive account of the kinematics of emotions, rather than just the negative conclusion that this kinematics differs significantly from the kinematics of ratiocination. We shouldn't expect to be able to generate this account from the comfort of the philosopher's armchair; instead it will probably take a lot of hard empirical work. I mentioned above [in the preceding paper] a catalog of potential impacts that emotions might have on cognition, together with intuitive, empirical and theoretical reasons for thinking that many emotions do have such impacts. Future empirical work may do much to extend, modify and/or refine this catalog.

A full positive account of the kinematics of emotions must also include an understanding of when emotions are generated, and when they are abandoned. A full list of the various ways in which emotions are generated must contain, at least, the following:

- by ratiocinative events (e.g., your putting the clues together to reach the conclusion that your spouse is cheating may trigger anger or jealousy)
- by other cognitive events (e.g., your sub-conscious recognition of a threatening posture may trigger fear)
- by not-very-cognitive events (e.g., changes in blood sugar level¹⁹).

There are various ways you might consciously attempt to influence your emotions, either to change them or to cease them. These include:

¹⁹ One might be inclined to include in this category the classic experiments in which Stanley Schacter and Jerome Singer (1962) found that subjects who had been injected with epinephrine were more likely than controls to report having emotional experiences. I think it is *possible* that epinephrine injections might bring about real changes in an agent's cognitive/emotional state. However, I think the most plausible initial hypothesis is just that the epinephrine brings about a cluster of bodily changes, thereby setting off a cluster of physiological sensors that would normally go off in such circumstances only if the subject were in an emotional state; hence subjects who are used to using these sensors as indicators of emotional states tend to conclude that they are in emotional states. (And then subjects attempt to come up with – or 'confabulate' – the most plausible explanation for their being in some sort of emotional state.) So, I'm inclined to say that these probably aren't real cases of emotion – instead they're just cases where subjects are tricked by a cluster of bodily changes into thinking that they are in emotional states.

- Alter your ratiocination. (Concentrate on something else; silently count to 10; try to bring to mind occasions where you felt other emotions)
- Alter your surroundings. (Get some fresh air, quiet, food, exercise, sex or sleep).
- In learning to ride a bicycle, you consciously force yourself through the motions several times, and pretty soon sub-conscious systems pick up the skill. Similarly, it seems likely that sub-conscious skills may be developed for diminishing and/or profitably utilizing at least some emotions.

Since different cultural practices lead members of different cultures to do different ones of these things in response to certain circumstances, any of the above sorts of effects might help bring about cross-cultural differences in when people experience emotions, and in how emotions affect them.

These comments are meant only as a perfunctory attempt to suggest the sorts of issues that would arise in describing a full kinematics of the emotions, and to suggest that such a full kinematics might have the means to give a good account of both the ‘educability’ of emotions, and their variation across cultures. Unfortunately, it is beyond the scope of this paper to pursue these issues in more detail.

XI. A Unified Account of the Emotions.

Common usage categorizes some strikingly different phenomena under the same emotional headings. For example, consider the typical response to seeing a coiled snake – a very gripping upheaval in thought, though one with a very short time span – and the typical response to finding that one’s employer must fire a significant number of employees – a state that may persist for months or even years, with varying effects on one’s cognition. Common usage counts both these responses as instances of fear.

One potential response – the one given by Griffiths (1997) – to such cases would be to say that common usage has mistakenly lumped together disparate phenomena. There is a natural class of phenomena including the response to the snake and many other ‘heat of the moment’ responses that generate similar physiological and behavioral effects. As for the ‘fear’ of being fired, this is some quite different phenomenon, one which a mature psychology would not lump in with the (properly so called) fear of a coiled snake. I grant that it might turn out that the ‘emotions’ are a motley lot that defy a unified understanding, but I think we should be reluctant to settle upon this conclusion.

This is a good point to return to the demons-metaphor which has, thus far, served us well as an intuitive motivation. Intuitively, demonic possessions may last for different lengths of time and may vary in their effects during their duration. Despite all these differences, there may be many traits that these demonic possessions have in common – they all involve modulations in cognition. Further, there intuitively may be useful ways of categorizing demonic possessions that lump together short term possessions and longer term ones. E.g., these possessions may be similar in how they are caused, in what they are ‘about’, or in the sorts of effects that they have on cognition.

Similarly, if we conceive of emotions as essentially being states that normally modulate one's mode of cognition, then we may give a quite unified account of emotions that produce such changes over shorter or longer periods of time, and of emotions that produce different sorts of changes at different points during their duration. Furthermore, this approach to understanding emotions allows that useful categorizations of emotions might lump together certain 'heat of the moment' emotions with other 'cooler' emotions that have similar causes, effects, or representational contents.

For example, the 'hot' and immediate fear of a snake is like the 'cooler' fear of being fired in that both are normally triggered by indicators of a potential danger, in that both serve to represent the presence of such a danger, and in that both tend to bring about many modulations in cognition that serve to mitigate this danger – e.g., drawing attention and cognitive resources to this danger, making more easily accessible those memories that are relevant to this danger, and forcing the imagination of scenarios involving this danger. It is at least plausible that there is an interesting natural grouping of emotional responses that have this sort of cognitive profile, and hence it is at least plausible that common usage really has captured an interesting natural grouping in categorizing these instances together.

Hence, I conclude that a modes of cognition view offers good prospects for offering a unified account of both 'heat of the moment' and 'cooler' emotions.

XII. The Rationality of Emotions.

In this final section, I will consider the question of whether emotions are essentially *irrational* upheavals in our thought, or whether there might be a sense in which emotions may be rational or intelligent.

My discussion above suggests a number of reasons to think that emotions are *not* rational. It is common to take as a paradigm of rationality the kinematics of ratiocination – the ways in which we form ratiocinative beliefs and desires on the basis of other beliefs and desires. I have argued that the kinematics of emotions is probably quite different from the kinematics of ratiocination. Hence, in this sense, emotions fail to be 'rational'. Furthermore, I have argued that our emotions often stubbornly persist in the face of contrary evidence and our ratiocinative attempts to change them. As much as our evidence and our ratiocination sometimes recommend falling in love, our hearts often fail to follow suit; as much as our evidence and our ratiocination maintain that a glass elevator is perfectly safe, our fear persists undiminished. Hence, there is a sense in which emotions often come into conflict with the paradigm of 'rationality' that is ratiocination.

Despite these observations, there are also a number of interesting senses in which many emotions are 'rational'. We have noted that many emotional states have (belief-like and/or desire-like) representational content. Hence, there is a sense in which an emotion may be 'rational' or 'irrational' depending on how well or poorly it coheres with one's beliefs and desires. We have also noted that emotions are often triggered by states that themselves have representational content. Hence, there is a sense in which an emotion may be 'rational' or 'irrational' depending upon whether or not it was caused by

representations whose contents make reasonable the embracing of the content of that emotion. In both of these senses, many ordinary fears would be deemed ‘rational’, while many compulsive phobias would be deemed ‘irrational’ – a categorization that fits well with common intuitions.

Additionally, a modes of cognition proposal holds that it is essential to emotions that they normally have many striking effects on cognition. In many cases, such modulations may be deemed *useful* or *detrimental*, *apt* or *inapt*, *rational* or *irrational*.

This point should not be taken lightly. One lesson of a half-century of Artificial Intelligence research is that, for many interesting domains of problems, it is relatively easy to describe an algorithmic procedure that would *eventually* find a solution for a given problem, but it often is quite challenging to find an algorithm that will do so *quickly* and with *limited computational resources*. What are efficient computational strategies in some domains are quite poorly suited to many other domains. Hence, one of the keys to developing anything like general artificial intelligence seems to be developing ways of shifting and combining computational strategies – or ‘modes of computation’, if you will – in ways that allow a system to deal efficiently with whatever problems are at hand. Insofar as these findings are indicative of the sorts of design challenges that evolutionary design processes have faced in building effective cognitive systems, the lesson is that general intelligence probably must consist in often shifting and combining cognitive strategies – or *modes of cognition* – in order to deal effectively with the various sorts of problems that animals face on different occasions. Hence, one might conclude, emotions (construed as modes of cognition) are probably of very central importance in making us be the sorts of generally rational and intelligent agents that we are.²⁰

The preceding considerations give us theoretical reasons to think that our emotions – construed as states that modulate our cognition – play an essential role in (our peculiarly human form of) rationality. We also have strong empirical reasons to accept this conclusion. Damasio (1994) has quite convincingly argued on the basis of empirical studies that normal emotional capacities are essential to ordinary reasoning and planning – both in providing necessary modulations in reasoning strategy, and in generating states (of brain and body) whose existence may then serve to support further reasoning. When our emotions do their normal jobs well, they play a central role in our rationality, and when they fail to do their jobs well, the result is often significant irrationality.

Hence, a modes of cognition view does well to capture the senses in which our emotions are always *non-rational* (in that they work differently from, and sometimes counter to, our ratiocination), and the undeniable senses in which they are of central importance to our rationality and intelligence.

²⁰ For a critical discussion of some reasoning along these lines, see Evans (2002).

References.

- Ben-Ze'ev, Aaron. 2000. *The Subtlety of Emotions*. Cambridge, MA: MIT Press.
- Charland, Louis. 2002. "The Natural Kind Status of Emotion." *British Journal for the Philosophy of Science* 53:511-37.
- Cosmides, Leda, and John Tooby. 2000. "Evolutionary Psychology and the Emotions." In *Handbook of Emotions*, Michael Lewis and Jeannette M. Haviland-Jones, 91-115. New York: Guilford Press.
- Damasio, Antonio. 1994. *Descartes' Error: Emotion, Reason, and the Human Brain*. New York: G.P. Putnam's Sons.
- Davidson, Donald. 1973, 'Radical Interpretation', *Dialectica*, 27.
- de Sousa, Ronald. 1987. *The Rationality of Emotion*. Cambridge, MA: MIT Press.
- _____. 2003 "Emotion". *Stanford Encyclopedia of Philosophy*.
<http://plato.stanford.edu/entries/emotion/>
- Dennett, Daniel. 1975. "True Believers: The Intentional Strategy and Why it Works" reprinted in D. Rosenthal ed. *The Nature of Mind*. Oxford: UP. pp. 339-53.
- Elster, Jon. 2003. "Emotion and Action." In *Thinking About Feeling: Contemporary Philosophers on Emotions*, ed. Robert C. Solomon. Oxford; New York: Oxford UP.
- Evans, Dylan. 2002. "The search hypothesis of emotions". *British Journal for the Philosophy of Science* 53:497-509.
- Frank, Robert. 1988. *Passions Within Reason: The Strategic Role of Emotions*. New York: Norton.
- Goldie, Peter. 2000. *The Emotions: A Philosophical Exploration*. Oxford: Oxford University Press.
- Gordon, R. 1987. *The Structure of Emotions*. Cambridge: Cambridge UP.
- Griffiths, Paul. 1997. *What Emotions Really Are: The Problem of Psychological Categories*. Chicago: University of Chicago Press.
- James, William. 1884. "What is an Emotion?" *Mind*, 9, 188-205.

- Lange, Carl. 1885. *Om Sindsbevaegelser: Et psyko-fysiologisk studie* in C.G.Lange & W.James (Eds.), I.A.Haupt (Trans.) *The Emotions*, 1922, Baltimore: Williams and Wilkins.
- Lewis, David. 1983. "Mad Pain and Martian Pain" reprinted in D. Rosenthal ed. *The Nature of Mind*. Oxford: UP. pp. 229-35.
- _____. "Desire as Belief II", *Mind* 105: 303-313.
- Marks, Joel. 1982. "A Theory of Emotion." *Philosophical Studies* 42: 227-242.
- Millikan, Ruth. 1984. *Language, Thought, and Other Biological Categories*. MIT Press.
- _____. 1996. "Pushmi-pullyu Representations", in James Tomberlin, ed., *Philosophical Perspectives* vol. IX, Atascadero CA: Ridgeview Publishing, pp. 185-200.
- Neu, Jerome. 2000. *A Tear is an Intellectual Thing: the Meanings of Emotion*. Oxford, New York: Oxford UP.
- Nussbaum, Martha. 2001. *Upheavals of Thought: The Intelligence of Emotions*. Cambridge: Cambridge UP.
- Pinker, Stephen. 1997. *How The Mind Works*. Norton.
- Prinz, Jesse. 2004. *Gut Reactions*. New York: Oxford UP.
- Oakley, Justin. 1992. *Morality and the Emotions*. London: Routledge and Kegan Paul.
- Oatley, K. & Johnson-Laird, P.N. 1987. "Toward a Cognitive Theory of the Emotions." *Cognition and Emotion*, 1, 29-50.
- Panseppe, Jaak. 2000. "Emotions as Natural Kinds With the Mammalian Brain." In M. Lewis & J. Haviland-Jones (Eds.), *Handbook of Emotions* (2ns ed., pp. 137-56). New York: Guilford Press.
- Roberts, R.C. 1988. "What an Emotion Is: A Sketch". *Philosophical Review*, 97, 183-209.
- Robinson, J. 1983. "Emotion, Judgment, and Desire". *Journal of Philosophy*, 80, 731-41.
- Ryle, Gilbert. 1949. *The Concept of Mind*. Chicago: University of Chicago Press.
- Solomon, Robert. 1999. "The Philosophy of Emotions." In *Handbook of Emotions*, ed. Mark Lewis and Jeannette Haviland-Jones, 3-15. New York: Guilford Press.

Sizer, Laura. 2000. "Towards a Computational Theory of Mood." *British Journal for the Philosophy of Science*: 51. 743-69.

Skinner, B.F. 1962. *Science and Human Behavior*. New York: Free Press.

Stocker, Michael, with Elizabeth Hegeman. 1992. *Valuing Emotions*. Cambridge: Cambridge UP.

Tappolet, Christine. Forthcoming. "Irrational Emotions". In *Perspectives on Irrationality*, ed. Daniel Weinstock. Toronto: University of Toronto Press.

Watson, J.B. 1919. *Psychology from the Standpoint of a Behaviorist*. Philadelphia: Lippincott.

Wollheim, Richard. 1999. *On the Emotions*. New Haven: Yale University Press.

Wright, Ian, Aaron Sloman, and Luc Beaudoin. 1996. "Towards a Design-Based Analysis of Emotional Episodes." *Philosophy, Psychiatry, and Psychology* 3:101-126.