

Emotions as Metarepresentational States of Mind

Rainer Reisenzein

Institute for Psychology, University of Greifswald
Franz-Mehringstr. 47, 17487 Greifswald, Germany
rainer.reisenzein@uni-greifswald.de

Abstract

I describe the outlines of a computational theory of emotions that views emotions as nonconceptual metarepresentations. According to this theory—which can be regarded as a computational explication of the belief-desire theory of emotion—at least a core subset of emotions including surprise are nonconceptual products of hardwired mechanisms whose primary function is to subserve the monitoring and updating of the central representational system of humans, the belief-desire system. The posited emotion-producing mechanisms are analogous to sensory transducers; however, instead of sensing the world, they sense the state of the person's belief-desire system and signal to the experiencer important changes in this system, in particular the fulfillment and frustration of desires and the confirmation and disconfirmation of beliefs.

1. The belief-desire theory of emotion

What are emotions, and what is their function in the economy of mind? I propose that at least for a core subset of emotions including surprise, these questions can be answered as follows: emotions are nonconceptual outputs (analog signals) of hardwired mechanisms whose primary function is to subserve the monitoring and updating of the central representational system of humans, the belief-desire system (see also Reisenzein, 1998; 1999; 2001). This particular view of emotions—which closely connects emotions to the updating of representations (see e.g., Paglieri, 2004) and assigns them important epistemic functions—may at first sight appear unusual. In fact, however, it is to a large degree already implicit in the currently dominant, cognitive theories of emotion (for an overview, see Scherer, Schorr, & Johnstone, 2001). For these theories assume implicitly that emotions are closely tied to changes in beliefs and desires; and at least some of them explicitly attribute to emotions an informational function (e.g., Ortony, Clore, & Collins, 1988).

In fact, the present theory can be viewed as an attempt to "naturalize" a particular brand of cognitive emotion theory, the belief-desire theory of emotions (*BDTE*), by sketching a computational model of this theory. For this reason, an excellent starting point for motivating and developing the current view of emotions is to begin with a summary of *BDTE*. Note, however, that my main aim is not to present a worked-out computational model of emotions, but to use computational thinking (thinking in terms of representational and computational mechanisms) as a tool to clarify *BDTE* and more generally, to become clearer about several unresolved issues of emotion psychology.

As mentioned, *BDTE* belongs to the broader class of cognitive emotion theories which have come to dominate the psychology of emotions during recent years (e.g., Arnold, 1960; Lazarus, 1991; Ortony, Clore, & Collins, 1988; Roseman, 1979; Smith & Kirby, 2001) and which form the basis of most existing computational models of emotion (e.g., Elliott, 1992; Gratch & Marsella, 2004; Staller & Petta, 2001).

As a distinct type of cognitive emotion theory within the cognitive approach to emotions, *BDTE* has however been primarily developed by philosophers (e.g., Davis, 1981; Green, 1992; Marks, 1982; Meinong, 1906; see also Reisenzein, 2006). The basic conceptual framework of *BDTE* is the same as that of the philosophical belief-desire theory of action that inspired the *BDI* (belief-desire-intention) approach to artificial agents (e.g., Bratman, Israel, & Pollack, 1988). Analogous to the belief-desire theory of actions, which assumes that actions are the product of cognitive or informational states (beliefs) and motivational states (desires), *BDTE* assumes that emotions are the product of cognitions (beliefs) and motives (desires). More precisely, emotions are viewed as reactions to "cognized" actual or potential fulfillments and frustrations of desires plus, in some cases (e.g., surprise, disappointment), confirmations or disconfirmations of beliefs.

To illustrate, according to *BDTE*, Mary is happy about Schroiber's election as chancellor if she (a) (firmly) believes that Schroiber was elected and (b) desires this state of affairs to obtain. In slightly more detail, the process of emotion generation (here illustrated for happiness) typically looks as shown in Figure 1. First, the person comes to desire some proposition (or state of affairs) *p*. (Note that this is the philosophical usage of "proposition"; psychologists typically use the term to denote sentences in a language-like mental representation system that *represent* states of affairs). For example, Mary acquires the desire that *Schroiber be elected as chancellor*. Some time later—as the result of new information acquired through the senses, communication from others, or inference from preexisting beliefs—the person acquires the belief that *p* obtains. For example, when watching the news on TV, Mary comes to believe that Schroiber was, indeed, elected as chancellor. Thereupon, the emotion occurs: Mary now feels pleased or happy that Schroiber was elected.




Figure 1. Basic belief-desire analysis of emotions

By amending and refining the just-described belief-desire analysis, it is possible to specify the cognitive and motivational preconditions of numerous emotions distinguished in ordinary language (e.g., Davis, 1981; Meinong, 1906; Searle, 1983). Indeed, there is reason to believe that *all emotions with propositional objects* (all emotions directed at states of affairs) are amenable to a belief-

desire analysis. These emotions cover the vast majority of the emotions distinguished in ordinary language (see also, Ortony et al., 1988). Here, a few examples must suffice: Mary is *happy* that *p* (e.g., that Schroiber is elected as chancellor) if she desires *p* and comes to firmly believe that *p* obtains; she is *unhappy* that *p* if she desires *not-p* (in the example: that Schroiber is not elected) and comes to firmly believe that *p* obtains. Mary *hopes* that *p* if she desires *p* but is uncertain whether or not *p* obtains; she *fears* *p* if she desires *not-p* but is uncertain about *p*. Mary is *surprised* that *p* if she up to now believed *not-p* and now comes to firmly believe *p*; she is *disappointed* that *p* if she desires *p* and up to now believed *p*, but now comes to firmly believe *not-p*; and she is *relieved* that *p* if she desires *not-p* and up to now believed *not-p*, but now comes to firmly believe *p*. Other-regarding emotions, such as *joy for another*, *pity*, *Schadenfreude* and *envy* can be analyzed as forms of happiness and unhappiness about a desired or undesired state of affairs *p* that concerns the positive or negative fate of another person. *Guilt*, *indignation* and related emotions can be analyzed by assuming that the object of the desire (the desired proposition) in these cases is the compliance of a person with a social or moral norm (e.g., Ortony et al., 1988; Staller & Petta, 2001).

As hinted by these examples, an important feature of the belief-desire analysis of emotion is the assumption that *all* emotions (with propositional objects), however complex, can be understood as reactions to "cognized" actual or possible fulfillments or frustrations of desires (plus confirmations or disconfirmations of beliefs). The complexity of humans' emotional life gets into the picture primarily, if not exclusively, via the *objects* of the beliefs or desires. Finally, the intensity aspect of emotions can be taken care of by considering *degrees* of belief and desire (e.g., Davis, 1981).

2. Naturalizing the Belief-Desire Theory of emotion

2.1 The representational system

To develop a computational model of *BDTE* requires to adopt the assumption that mental and especially cognitive processes are computations in an internal representation system (Reisenzein, 2001). I make the by now "classical" assumption of cognitive psychology that the central representation system of humans is symbolic and more precisely, that it is a language-like, or "propositional" system of representations, a language of thought (e.g., Fodor, 1987). A main reason for this assumption is that, in contrast to other proposed representation systems, a propositional system lends itself naturally to a transparent and plausible computational analysis

of the intentional mental states posited in folk psychology, such as beliefs and desires.

How can the naturalization of beliefs and desires be achieved, given a language-like system of mental representations? The answer, suggested by Fodor (1987) and others, is to treat *believing* and *desiring* as special modes of processing propositional representations (i.e., sentences in the language of thought). To use a frequent metaphor, let us assume that believing a proposition *p* consists, computationally, of having a sentence *s* representing *p* in a special memory store (which is accordingly called the "belief store"); and that desiring *p* consists of having a sentence *s* representing *p* in another memory store (the desire store). Note that this mode of speaking is meant to be shorthand for a functional description of beliefs and desires, that is, a description in terms of their causal roles in the system (Fodor, 1987). For example, to say that a sentence is the belief store is shorthand for saying that this sentence is "treated as true" by the system: it is used as a valid premise in inferences, is relied on when planning actions, elicits surprise when it turns out to be false, and so forth. The storage metaphor is thus not essential. What is essential, however, is that the system or agent is able to distinguish between propositions that it merely represents, without believing or desiring them; propositions that it believes; and proposition that it desires (see also, Paglieri, 2004).

To illustrate, consider a moment in Mary's belief-desire system (Fig. 2). At the moment considered, Mary has the indicated, and many more, "mentalese" sentences in her belief store. These sentences represent the states of affairs that she currently believes to obtain. For example, as can be seen, Mary currently believes that Schroiber will not win the election. Likewise, Mary currently has the indicated, and many more, sentences in her desire store. These represent the states of affairs that she currently desires. For example, Mary currently wishes that Schroiber wins the election.

2.2 The belief-belief comparator and the belief-desire comparator

According to *BDTE*, Mary experiences happiness about Schroiber's election as chancellor if she desires this state of affairs and comes to believe that it obtains. To model this process, let us begin by assuming that newly acquired beliefs are placed into a special memory store (a store reserved for newly acquired beliefs). We can then say that, computationally speaking, Mary feels happy that Schroiber won the election when (or soon after) a sentence representing this state of affairs is deposited in her store for new beliefs (Fig. 2).

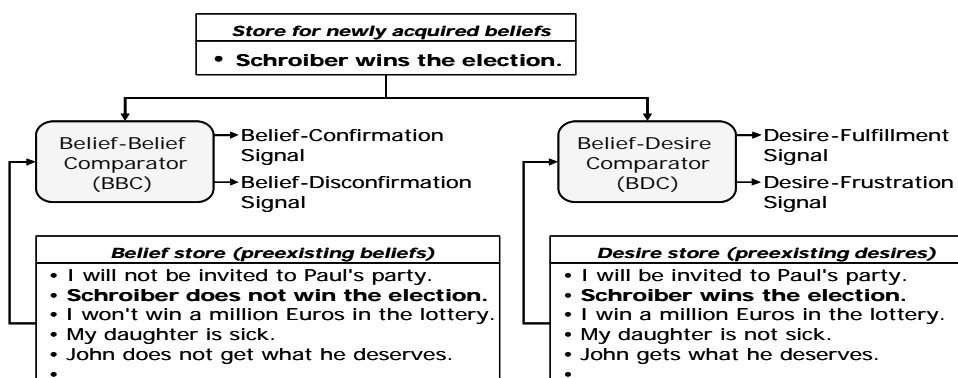


Figure 2. Illustration of the belief-belief and belief-desire comparators: A moment in Mary's belief-desire system

Figure 1 suggests that Mary's ensuing happiness about p is simply the causal effect of the joint presence of desiring p and believing p . However, when looked at from a computational perspective, it is clear that, for happiness about p to occur, it is not enough that the belief that p is present in the system simultaneously with the desire for p (i.e., it is not enough that Mary has a sentence representing p in her desire store, and another sentence representing p in her store for new beliefs). In addition, the system needs to *relate* these two facts: it has to *detect* that the newly believed proposition is identical to a desired proposition. To achieve this, a mechanism is needed that *compares* the newly acquired belief to the person's preexisting desires, looking for match and mismatch.

Similarly, according to *BDTE*, Mary is surprised about Schroiber's election as chancellor if she believed that this state of affairs would not occur and then comes to believe that it did, in fact, occur. Again, however, it is not sufficient for the occurrence of surprise that the newly acquired belief that p is present in the system simultaneously with the "old" belief that *not-p*. In addition, the system must recognize that the newly acquired belief conflicts with a preexisting belief. To achieve this, a mechanism is needed that compares the newly acquired belief to the preexisting beliefs for match versus mismatch.

The existence of such a *belief-desire comparator (BDC)* and *belief-belief-comparator (BBC)* can however also be motivated quite independently from any consideration of emotions (Reisenzein, 1998; 1999). For surely, if the belief-desire system of humans is to fulfill its major presumed function—to enable adaptive action in an imperfectly known and changing environment—it needs to be equipped with mechanisms that, if necessary, update the system in response to newly acquired information (beliefs). Updating means to add new beliefs and desires to the system, as well as—and more important in present context—to abandon old beliefs (if false) and old desires (if fulfilled). However, if this updating of beliefs and desires is to be adaptive, the need for and nature of the updating must first be *diagnosed*, by comparing the newly acquired information to the existing beliefs and desires. Accordingly, the updating mechanisms must contain appropriate comparator devices; and again, one can distinguish between two (at least analytically separable) comparators: one that compares newly acquired beliefs to existing beliefs (*BBC*), and another that compares newly acquired beliefs to existing desires (*BDC*).

To become clearer about how the two comparators work, let us again consider the moment when Mary comes to believe that Schroiber won the election. Computational speaking, a sentence representing this state of affairs is deposited in Mary's store for new beliefs (Fig. 2). This belief is now compared to Mary's preexisting beliefs and desires. The belief-belief comparator (*BBC*) compares the newly acquired belief to Mary's preexisting beliefs for match versus mismatch. A match means that a preexisting belief is confirmed, whereas a mismatch means that a preexisting belief is disconfirmed by the new information. Computationally speaking: the sentence currently in Mary's store for newly acquired beliefs, s_{new} , is compared to the sentences currently in her store for preexisting beliefs. If either a match (s_{new} is identical to a sentence s_{old} in the store for preexisting beliefs) or a mismatch (s_{new} is identical to the negation of a sentence in the belief store, $\neg s_{old}$) is detected, the *BBC* generates an output that signals the detection of the match or mismatch. In our example, Mary's *BBC* detects that the content of the newly acquired belief (Schroiber wins the election) is inconsistent with (is the negation of) the content

of a preexisting belief. Consequently, Mary's *BBC* outputs information about a mismatch—information that one of Mary's beliefs has been disconfirmed by the new information (see Fig. 2).

The belief-desire comparator (*BDC*) compares the newly acquired belief to preexisting desires for match versus mismatch. A match means that a desire has been (at least subjectively) fulfilled, whereas a mismatch means that a desire has been frustrated. Computationally speaking: the sentence s_{new} in Mary's store for newly acquired beliefs is compared with the sentences currently in her desire store. If either a match or a mismatch is detected, the *BDC* generates an output that signals the detection of the match or mismatch. In our example, Mary's *BDC* detects that the content of the newly acquired belief (Schroiber wins the election) is identical to the content of an existing desire. Consequently, Mary's *BDC* outputs information about a match—information that one of Mary's desires has been fulfilled (see Fig. 2).

Note that the foregoing description of the belief-belief and belief-desire comparison processes has been deliberately simplified. A more realistic model is sketched in Reisenzein (1998; 1999). This model (a) distinguishes between long-term memory (LTM) and working memory and assumes that the comparison of newly acquired and existing beliefs and desires always takes place within working memory (propositions in LTM therefore need to be retrieved into working memory before they can be compared to other propositions); (b) treats beliefs and desires as quantitative variables (i.e., it considers *degrees* of belief and desire); (c) assumes that the *BBC* and *BDC* compute not just the match or mismatch of newly acquired beliefs to existing beliefs and desires, but degrees of expectedness versus unexpectedness, and degrees of desiredness versus undesiredness; and (d) assumes that degrees of congruence or incongruence to existing beliefs and desires are also computed for newly acquired beliefs that are held with less than complete certainty (this assumption is necessary to model hope and fear).

2.3 Functional consequences of the detection of belief- and desire (in-)congruence

Having sketched how belief- and desire congruence and incongruence are detected, I turn to the functional consequences of a detected (in-)congruence. Of these consequences, three seem particularly important. First, attention is focused on the contents of the newly acquired beliefs that gave rise to a match or mismatch; that is, the (un-)expected or (un-)desired propositions. For example, in Mary's case, attention is directed to the unexpected proposition that Schroiber won the election. This focusing of attention is a precondition for the further conscious processing of the event in question, for example, an analysis of its causes (cf. Meyer, Reisenzein, & Schützwohl, 1997). Second, some minimal, immediate updating of the belief-desire system takes place automatically: Sentences representing state of affairs that are now believed to obtain are deleted from the desire store, and disconfirmed beliefs are deleted from the belief store. (Note that this does not mean that the belief or desire contents are *forgotten*; it only means that they are no longer believed or desired, respectively; cf. Paglieri, 2004). Third, and of particular importance in the present context, if the congruence or incongruence signals generated by the *BBC* and *BDC* exceed a certain threshold of intensity, they become conscious. Thereby, the information carried by these signals becomes

globally available for further processing (e.g., Baars, 1988; Block, 1995; see also, Reisenzein, 2000).

2.4 Implementation, mode of operation, and output format

So far, the *BBC* and the *BDC* have been described in terms of their inputs, outputs, and immediate functional effects. I now come to what are, in one sense, the most important assumptions of the present theory. These concern the implementation, mode of operation, and output format of the proposed comparator mechanisms.

First, the *BBC* and the *BDC* are not learned procedures, but are hardwired into the brain.

Second, partly as a consequence, these mechanisms (a) operate on a preconscious level, without and even against our intentions. Furthermore, (b) they operate continuously, that is, they work on every single, newly acquired belief; and (c) they compare each belief in parallel to the preexisting beliefs and desire contents (more precisely speaking, those that are currently in working memory). In short, every newly acquired belief is automatically—without intention, and preconsciously—compared simultaneously to all belief and desire contents currently in working memory.

Third, as already suggested by the forgoing description of the outputs of the *BBC* and *BDC*, I assume that their output is *not propositional*—it is not another sentence in the language of thought that represents the detection of a match or mismatch. Rather, I assume that the outputs of the *BBC* and *BDC* are nonpropositional and nonconceptual: they consist of *analog signals* that have no internal structure, but vary only in kind and intensity (cf. Oatley & Johnson-Laird, 1987; Picard, 1997). These signals carry information about the degree of expectedness versus unexpectedness, and the degree of desiredness versus undesiredness of the contents of newly acquired beliefs; but they do not represent these contents themselves.

In sum, I propose that the belief-desire system comes equipped with a set of basic, hardwired monitoring- and updating mechanisms, the *BBC* and the *BDC*. These mechanisms are, in a sense, similar to sensory transducers (i.e., sense organs for color, sound, touch, or bodily changes); in particular, their immediate outputs are analog signals. However, instead of sensing the world, these "internal transducers" sense the state of the belief-desire-system and signal important actual and impending changes in this system.

2.5 The belief-belief and belief-desire comparators and the emotions

According to *BDTE*, happiness about a proposition *p* occurs if one desires *p* and comes to believe *p*; whereas surprise about *p* is felt if one previously believed *not-p* and now comes to believe *p*. The computational analysis described in the last section suggested that the causal link between newly acquired beliefs and preexisting beliefs and desires on the one hand, and emotions on the other hand, is mediated by two comparison processes, the *BBC* and the *BDC*. Emotions result when the comparator mechanisms detect a match or mismatch of a newly acquired belief with preexisting beliefs (*BBC*) or desires (*BDC*). Thus, it turns out that the two basic, hardwired comparator mechanisms that service the belief-desire system are simultaneously the two basic, emotion-producing mechanisms. If this is accepted, the question regarding the nature and function of emotions raised in the introduction can be answered as follows:

1. Emotions are the nonpropositional signals of match versus mismatch (or considered quantitatively, the analog signals of desiredness, undesiredness, expectedness, and unexpectedness) produced by two basic, hardwired mechanisms that service the belief-desire system, the *BBC* and the *BDC*. At least, these signals form the nonconceptual core of emotions. They are consciously experienced as feelings of pleasure or displeasure (*BDC*), and as feelings of expectancy confirmation vs. surprise (*BBC*).

2. The function of the emotions (at least, one important function) is to assist in the updating of the fundamental action-guiding representational system of humans, the belief-desire system.

Finally, note that the outputs of the *BBC* and *BDC* carry information *about* beliefs and desires. As a consequence—if one accepts that nonconceptual mental states such as sensations can be called *representations* (e.g., Dretske, 1995; Tye, 1995)—emotions are *nonconceptual metarepresentations*. Specifically, emotional feelings represent to the experiencer, in a nonconceptual format, current and impending state changes of his or her belief-desire system (e.g., surprise: "a belief has been disconfirmed"; pleasure: "a desire has been fulfilled").

3. Discussion

As mentioned, the theory of emotions as metarepresentations can be regarded as a computational explication of the belief-desire theory of emotions. To the degree that the current theory is able to reproduce the predictions of *BDTE*, it is therefore supported by the same common-sense intuitions and systematic data that support *BDTE*. However, the explanatory capacity of the theory goes farther than this, for the theory also provides answers to several moot questions of emotion psychology, including the following:

1. The theory explains what is *distinctive* about emotional experiences: namely, that they are at core, unique (that is, specific to emotions) sensation-like experiences.

2. The theory provides for a *principled* demarcation of basic emotional feelings: The basic emotional feelings are exactly the consciously experienced outputs of the two comparator mechanisms (the *BBC* and the *BDC*).

3. The theory solves (or resolves) the "intentionality puzzle" (the puzzle of the object-directedness) of emotions. The puzzle is this: Emotions such as happiness, unhappiness, fear etc. seem to be normally directed at propositional objects (e.g., Mary is happy that Schroiber was elected), but sometimes they seem to lack objects (e.g., moods: Mary just feels happy). The present theory resolves this puzzle, somewhat radically, by assuming that the object-directedness of emotions is an *illusion*. Emotional feelings represent only the congruence or incongruence of newly acquired beliefs with existing beliefs or desires, but not the *content* of these beliefs and desires. For this reason alone, they cannot be directed at the objects of their causative beliefs and desires. The illusion of the directedness of emotions at propositional objects is due to the automatic, inescapable focusing of attention on (mis-)matching propositions. For example, when a belief is disconfirmed, the person near-simultaneously experiences a feeling of surprise and has her attention drawn to the proposition "at fault". It then seems to the person that she is surprised *about p*, that her feeling somehow *represents p*.

4. The theory (re-)solves the "cognition-emotion debate" in psychology, that is, the question of whether cognitions are necessary for emotions, and if yes in which sense. Specifically, the present theory suggests the following, differentiated answer to this question: (a) Cognitions

(beliefs) are necessary for (propositional) emotions: not only are emotions the result of operations on propositional representations; they are *themselves* representations (and hence cognitive in a wide sense of the term). (b) At the same time, however, emotions are not forms of propositional knowledge, and the mechanisms which produced them (the *BDC* and *BBC*) are not propositional inference procedures.

However, perhaps the most important consideration that speaks in favor of the present theory is that alternative interpretations of *BDTE* face difficulties. As noted, belief-belief and belief-desire comparators seem necessary to explain how emotions are generated; at issue is therefore only their special form. The only alternative to the present proposal that seems to have been entertained is that the *BBC* and the *BDC* are ordinary propositional inference processes: The person forms higher-order-beliefs (metacognitions) about the match or mismatch of a newly acquired belief with preexisting beliefs and desires. For example, in the case of surprise, the person reasons: "Up to now, I believed *not-p*; just now, I came to believe *p*; thus, my previous belief that *not-p* is disconfirmed." This metacognition then presumably causes the experience of surprise. This theory strikes me as extremely implausible. For example, it is doubtful whether small children, who already seem to experience happiness and surprise, are capable of making such complex inferences. In addition, from a phenomenological perspective, this theory simply gets the object-directedness of emotions wrong. For example, one is surprised that *p*, not that one's previous belief that *not-p* has been disconfirmed (or what comes down to the same, that *p* is unexpected). The present theory avoids these problems.

References

Arnold, M. B. (1960). *Emotion and personality (Vol. 1 & 2)*. New York: Columbia University Press.

Baars, B. (1988). *A cognitive theory of consciousness*. Cambridge: Cambridge University Press.

Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, 18, 227-287.

Bratman, M., Israel, D., & Pollack, M. (1988). Plans and resource-bounded practical reasoning. *Computational Intelligence*, 4, 349-355.

Davis, W. (1981). A theory of happiness. *Philosophical Studies*, 39, 305-317.

Dretske, F. (1995). *Naturalizing the mind*. Cambridge, MA: MIT Press.

Elliott, C. (1992). *The Affective Reasoner: A process model of emotions in a multi-agent system*. PhD Thesis. Evanston, IL: Northwestern University.

Fodor, J. A. (1987). *Psychosemantics: The problem of meaning in the philosophy of mind*. Cambridge, MA: MIT Press.

Gratch, J. and Marsella, S. (2004). A domain independent framework for modeling emotion. *Journal of Cognitive Systems Research*, 5, 269-306.

Green, O. H. (1992). *The emotions: A philosophical theory*. Dordrecht: Kluwer.

Lazarus, R. S. (1991). *Emotion and adaptation*. New York: Oxford University Press.

Marks, J. (1982). A theory of emotion. *Philosophical Studies*, 42, 227-242.

Meinong, A. (1906). Über Urteilsgefühle: Was sie sind und was sie nicht sind [On judgment-based feelings: What they are and what they are not]. *Archiv für die gesamte Psychologie*, 6, 22-58.

Meyer, W.-U., Reisenzein, R., & Schützwohl, A. (1997). Towards a process analysis of emotions: The case of surprise. *Motivation and Emotion*, 21, 251-274.

Oatley, K., & Johnson-Laird, P. N. (1987). Towards a cognitive theory of emotions. *Cognition and Emotion*, 1, 29-50.

Ortony, A., Clore, G. L. & Collins, A. (1988). *The cognitive structure of emotions*. Cambridge: Cambridge University Press.

Pagliari, F. (2004). Data-oriented belief revision: Towards a unified theory of epistemic processing. In E. Onaindia, & S. Staab (Eds.), *Proceedings of STAIRS 2004* (pp. 179-190). Amsterdam: IOS Press.

Picard, R. W. (1997). *Affective Computing*. Cambridge, MA: MIT Press.

Reisenzein, R. (1998). Outlines of a theory of emotions as metarepresentational states of mind. In A. H. Fischer (Ed.), *ISRE'98. Proceedings of the 10th Conference of the International Society for Research on Emotions, 1998* (pp. 181-191). Amsterdam: Faculty of Psychology.

Reisenzein, R. (1999). A theory of emotions as metarepresentational states of mind. *Unpublished manuscript* (available from the author on request).

Reisenzein, R. (2000). The subjective experience of surprise. In H. Bless & J. P. Forgas (Eds.), *The message within: The role of subjective experience in social cognition and behavior* (pp. 262-279). Philadelphia, PA: Psychology Press.

Reisenzein, R. (2001). Appraisal processes conceptualized from a schema-theoretic perspective: Contributions to a process analysis of emotions. In K. R. Scherer, A. Schorr & T. Johnstone (Eds.), *Appraisal processes in emotion: Theory, methods, research* (pp. 187-201). Oxford: Oxford University Press.

Reisenzein, R. (2006). Arnold's theory of emotion in historical perspective. *Cognition and Emotion*, in press.

Roseman, I. J. (1979, September). Cognitive aspects of emotions and emotional behavior. *Paper presented at the 87th Annual Convention of the APA*, New York City.

Searle, J. (1983). *Intentionality*. Cambridge: Cambridge University Press.

Scherer, K. R., Schorr, A. & Johnstone, T. (2001). *Appraisal processes in emotion: Theory, methods, research*. Oxford: Oxford University Press.

Smith, C.A. & Kirby, L.A. (2001). Toward delivering on the promise of appraisal theory. In K. R. Scherer, A. Schorr & T. Johnstone (Eds.), *Appraisal processes in emotion: Theory, methods, research* (pp. 121-138). Oxford: Oxford University Press.

Staller, A. & Petta, P. (2001). Introducing emotions into the computational study of social norms: A first evaluation. *Journal of Artificial Societies and Social Simulation*, 4(1). <<http://www.soc.surrey.ac.uk/JASSS/4/1/2.html>>

Tye, M. (1995). *Ten problems of consciousness. A representational theory of the phenomenal mind*. Cambridge, MA: MIT Press.