

A formal model of emotion triggers: an approach for BDI agents

Bas R. Steunebrink · Mehdi Dastani ·
John-Jules Ch. Meyer

Received: 4 November 2010 / Accepted: 23 August 2011 / Published online: 20 September 2011
© The Author(s) 2011. This article is published with open access at Springerlink.com

Abstract This paper formalizes part of a well-known psychological model of emotions. In particular, the logical structure underlying the conditions that trigger emotions are studied and then hierarchically organized. The insights gained therefrom are used to guide a formalization of emotion triggers, which proceeds in three stages. The first stage captures the conditions that trigger emotions in a semiformal way, i.e., without committing to an underlying formalism and semantics. The second stage captures the main psychological notions used in the emotion model in dynamic doxastic logic. The third stage introduces a BDI-based framework (belief–desire–intention) with achievement goals, which is used to firmly ground the preceding stages. The result is a formalization of emotion triggers for BDI agents with achievement goals. The idea of proceeding in these stages is to provide different levels of commitment to formalisms, so that it remains relatively easy to extend or replace the used formalisms without having to start from scratch. Finally, we show that the formalization renders properties of emotions that are in line with the psychological model on which it is based.

Keywords Logic of emotions · Intelligent agents · Cognitive modeling

B. R. Steunebrink
IDSIA & University of Lugano, Galleria 2, 6928 Manno-Lugano, Switzerland
e-mail: bas@idsia.ch

M. Dastani · J.-J. Ch. Meyer (✉)
Intelligent Systems Group, Institute of Information and Computing Sciences, Utrecht University,
P.O. Box 80089, 3508 TB, Utrecht, The Netherlands
e-mail: jj@cs.uu.nl

M. Dastani
e-mail: mehdi@cs.uu.nl

1 Introduction

Recently, there has been a lot of interest in bringing emotions to Artificial Intelligence, in particular to model rational agents (Adam et al. 2009; Breazeal 2002; Dastani and Meyer 2006; Gratch and Marsella 2004; Johns and Silverman 2001; Marinier and Laird 2004; Meyer 2006; Picard 1997; Sloman 2001; Steunebrink 2010; Steunebrink et al. 2007, 2008, 2009). There are (at least) three important reasons for this. First, an obvious application of emotions is to make artificial agents and robots more believable to human users. In particular, their behaviors are expected to appear increasingly convincing, social, and intuitive if they seem to have an emotional state matching that of a human in the same situation (Breazeal 2002; Gratch and Marsella 2004; Marinier and Laird 2004; Picard 1997). Second, from a more theoretical perspective, it is investigated what the role of emotions is in models of human decision-making and how they may be employed to make these models more accurate and effective (Coppin 2008; Elster 1996; Johns and Silverman 2001). Third, there exists psychological (Ekman and Davidson 1994; Frijda 1987; LeDoux 1996; Oatley and Jenkins 1996) and neurological (Damasio 1994) evidence that emotions are not only relevant but even necessary for rational behavior. Particularly, it has been shown that persons who do not experience emotions (e.g., due to specific brain damage) have trouble distinguishing between important and irrelevant details, consistently make bad decisions, and do not display adequate social behavior necessary to function normally in society. A related, more philosophical argument posits that emotions are an inevitable consequence of mechanisms that allow for intelligent and rational behavior in complex environments with limited resources (Elster 1996; Gordon 1987; Hoffman et al. 1991; Sloman 2001).

There is little consensus among psychologists as to what exactly constitutes an emotion and how it differs from related affective processes such as moods and impulses. However, this does not mean that making broad classifications is impossible or useless. According to a classification by Gross and Thompson (2007), *emotions* typically have specific objects and give rise to action tendencies relevant to these objects. Moreover, emotions can be both positive and negative. Emotions are often distinguished from *moods*, which are more diffuse and last longer than emotions. Other affective processes include *stress*, which arises in taxing circumstances and produces only negative responses; and *impulses*, which are related to hunger, sex, and pain and give rise to responses with limited flexibility. Of these four types of affective processes, we will focus on *emotions* in this paper.

With respect to emotions, usually three phases are distinguished. First, the perceived situation is *appraised* by an individual based on what he or she thinks is relevant and important. For example, Alice, who likes receiving presents, is given a necklace by Bob. Alice then judges receiving the necklace as desirable and Bob's action as praiseworthy. Consequently, the appraisal of this action and its outcome causes gratitude towards Bob to be triggered for Alice. Note that different types of emotions may be triggered simultaneously by the same situation, some of which may even be seen as conflicting. For example, Alice may at the same time be disappointed because it was not the necklace she had hoped to receive. Emotion theories dealing with appraisal are for example Frijda (1987), LeDoux (1996), Oatley and Jenkins (1996), Ortony et al.

(1988), and Scherer et al. (2001). Second, the appraisal of some situation can cause the triggered emotions, if exceeding some threshold, to create a conscious awareness of emotional feelings, leading to the *experience* of having emotions. For example, Alice's gratitude towards Bob will have a certain intensity and will probably decrease over a certain amount of time. All this may depend on, e.g., the degree of desirability of receiving a necklace and Alice's previous attitude towards Bob. Emotion theories dealing with these quantitative aspects of emotions are for example Ekman and Davidson (1994), Frijda (1987), and Ortony et al. (1988). Third, emotional feelings need to be *regulated*. For example, Alice may want to organize her behavior such that positive emotions are triggered as often as possible and negative emotions are avoided or drowned by positive ones. She could do this by being nice to Bob so that he will give her more presents, or avoiding him altogether so that she will never again be confronted with his bad taste in jewelry. In fact, some emotion theories posit that the main purpose of emotions is to function as a heuristical mechanism for selecting behaviors (Damasio 1994; LeDoux 1996; Oatley and Jenkins 1996). Emotion theories dealing with coping and behavioral consequences of emotions are for example Ekman and Davidson (1994), Frijda (1987), Lazarus (1994), and Oatley and Jenkins (1996).

With respect to models of rational agency, important topics include tracking of goal achievements (i.e. rate of success), revision of plans, and where to focus attention. Emotion theories can provide solutions to these issues by treating emotions as heuristics in the deliberation and decision making of agents. Given the aim to integrate emotions in the models of artificial rational agents on the one hand, and the existence of psychological theories of emotions on the other hand, a question that thus arises is how we can adopt, formalize, and use psychological theories of emotions in models of rational agency. Following psychological theories, at least three main issues (appraisal, experience, regulation) need to be addressed in this effort, but of course this cannot be done properly and comprehensively in one paper. The central question addressed in this paper is how to model and integrate the appraisal part of emotions in agent models. In particular, a formal framework must be built that is suitable for modeling agents and for investigating how appraisal can be integrated in this formal framework. In order to do this, we must map psychological concepts onto agent concepts so that they can be appropriately formalized.

In this paper, we will present a formalization of the eliciting conditions of emotions as described in the psychological model of Ortony et al. (1988).¹ We have chosen the OCC model because it provides a clear classification of a broad range of emotion types, it lists concise descriptions of the conditions that elicit emotions, and for this it uses concepts that are well studied and relatively straightforward to formalize. The presented formalization constitutes a formal model for the appraisal part of the OCC emotion theory. This formal model is an extension of an agent specification framework (i.e., KARO, Meyer et al. 1999; Meyer 2006) that specifies agents in terms of cognitive concepts that are also used in the OCC model. For each emotion type from the OCC model we translate the eliciting conditions, which define the appraisal process

¹ Henceforth to be referred to as “the OCC model” or simply “OCC.”

corresponding to that emotion type, into concepts from the formal agent framework. The contribution of this paper is to provide a formal agent model in which the appraisal part of the complete set of emotion types from the OCC model is integrated. We show that the formalization is adequate by providing a logical analysis and proving intuitive properties of the model.

This paper is organized as follows. In Sect. 2 we give an overview of the psychological OCC model of emotions. We will particularly study its logical structure in great detail, because it is on this structure that our formalization will be based. The formalization will proceed in three stages, spread over Sects. 3, 4, and 5, respectively. The first stage will be a semiformal specification of the logical structure of eliciting conditions. The second stage will capture the main notions used in the OCC model in dynamic doxastic logic. The third stage will formalize the main appraisal notions used in the OCC model in the KARO framework, which is an extension of dynamic doxastic logic, and thus firmly grounds the preceding stages. In Sect. 6 we will discuss and compare related work, and Sect. 7 will conclude this paper.

2 The OCC model

In their book “The Cognitive Structure of Emotions”, Ortony et al. (1988) have proposed a very interesting model of emotions that provides specifications of the eliciting conditions of emotions and the variables that affect their intensities. This psychological model of emotions is popular among computer scientists that are trying to build systems that reason about emotions or incorporate emotions in artificial characters. This popularity is due to the model’s clear and convincing structure.

2.1 Overview of the OCC model

The OCC model classifies 22 types of emotions. This is done by considering on which kinds of aspects of a situation one can focus his or her attention. OCC consider a human can either focus on consequences of events,² actions of agents, or aspects of objects. If one focuses on a consequence of an event, one can appraise this consequence as *desirable* or *undesirable* (or both, or neither) with respect to one’s *goals*. For example, joy about winning a lottery is an event-based emotion, because the satisfaction of the goal to become rich is a desirable consequence of the event of winning the lottery. If one focuses on an action of an agent, one can appraise this action as *praiseworthy* or *blameworthy* (or both, or neither) with respect to one’s *standards*. For example, pride about saving a child from drowning is an action-based emotion, because it is praiseworthy to perform an action which satisfies the standard that one should save a person’s life whenever (reasonably) possible. If one focuses on an aspect of an object, one can appraise this aspect as *appealing* or *unappealing* (or both, or neither) with respect to one’s *attitudes*. For example, love for an old car is an

² On p. 18 (Ortony et al. 1988), OCC say that “events are simply people’s construals about things that happen.” From a computational perspective, however, we would say that this is far from *simple*!

object-based emotion, because the car may have appealing aspects according to one's attitudes.

Within these three main categories of emotion types, the OCC model makes further differentiations based on, e.g., whether prospects are relevant (as in hope and fear), whether events apply to others (as in pity and gloating), or whether an action was performed by the self or someone else (to distinguish e.g. pride from admiration). Additionally, some event-based and action-based emotion types are combined to form a group of emotions concerning consequences of events *caused* by actions of agents. For example, anger can arise when one focuses on both the blameworthy action of another agent and an undesirable event which has been (presumed to be) caused by it. It should be emphasized that in the OCC model, emotions are never used to describe the entire cognitive state of an agent (as in "Alice is happy"); rather, emotions are always relative to individual events, actions, and objects. So Alice can be joyous about receiving her new furniture and at the same time be distressed about the height of the accompanying bill.

2.2 The logical structure of the OCC model

Although many ad hoc or simplified implementations of the OCC model have been made, there have been fewer attempts at formalizing the complete, logical structure of the proposed emotion model (e.g., [Adam et al. 2009](#); [Gratch and Marsella 2004](#); [Steunebrink 2010](#); [Steunebrink et al. 2007](#)). Here we will attempt to do so, using a formal logic containing constructs to reason about agents, their beliefs and actions, objects, and consequences of actions/events. We will be formalizing the eliciting conditions of emotions in this logic, trying to stay as close as possible to the OCC model.

On p. 19 of their book ([Ortony et al. 1988](#)), OCC present a diagram which structures their emotion types based on *focus of attention*. This diagram is often reproduced when an overview of the OCC model is to be given. In this section we give an overview of the OCC model, but we will illustrate the OCC model with a slightly different diagram (see Fig. 1). The purpose of this paper is to provide a logical account of emotion triggers. However, OCC's diagram as it appeared in [Ortony et al. \(1988\)](#), which is based on focus of attention, is not very well suited to guide our formalization, because it is not compositional (see [Steunebrink 2010](#) for a detailed discussion on this issue). Therefore, we have created the new diagram illustrating the structure of the emotion types based on their *eliciting conditions*. It should be noted that in personal communication, Ortony and Clore have confirmed Fig. 1 to be an accurate compositional illustration of the logical structure underlying the eliciting conditions of their emotion types ([Ortony and Clore 2009](#)). The following paragraphs serve to explain Fig. 1, which, importantly, serves as a guide for the formalization that will be presented in the next sections.

Figure 1 can be seen as an *inheritance* structure. This means that the depicted emotion types are *specializations* of those above them and *generalizations* of those below them. This inheritance-based perspective results in a compositional formulation of the eliciting conditions. At the most general level, all emotions are valenced

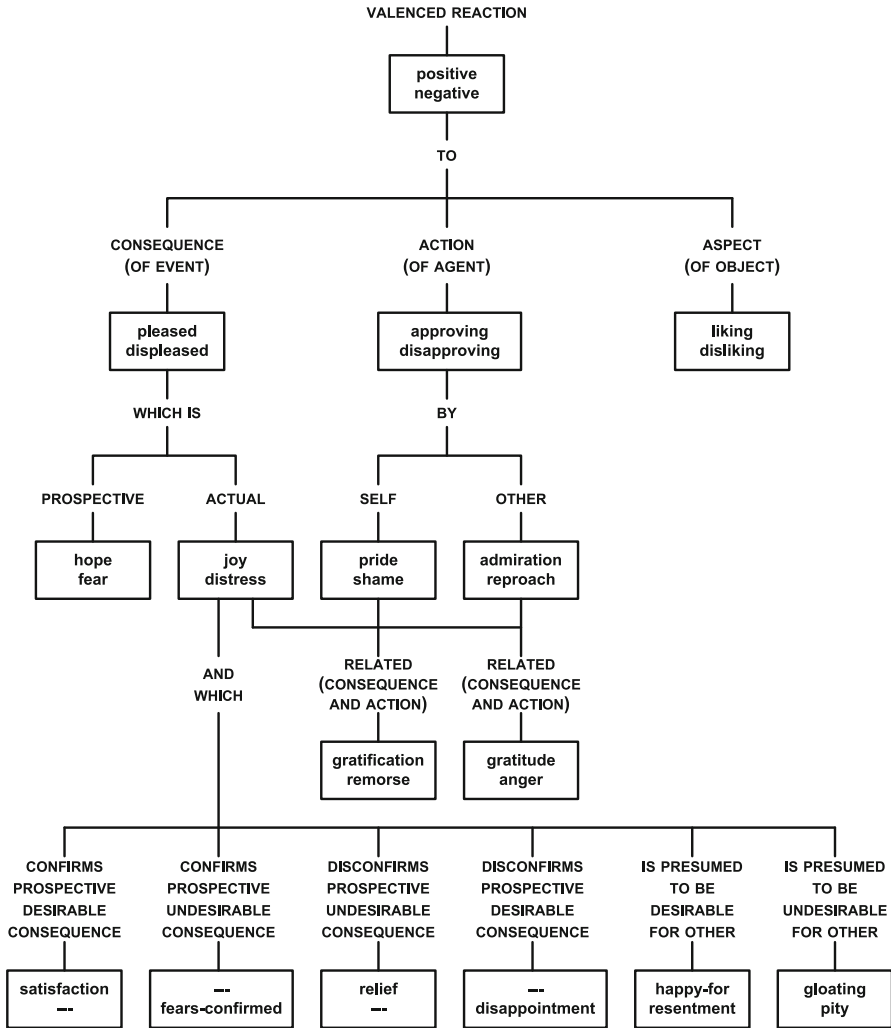


Fig. 1 An inheritance-based view of the eliciting conditions of emotions of the OCC model. Reading the labels on the lines from the top to any node provides a description of the emotion type represented by that node. Note though that these labels serve only for ease of reading; technical description are provided later on (see Table 1). The interested reader can find a more detailed analysis of the differences between OCC’s diagrams and descriptions and ours in [Steunebrink \(2010\)](#)

reactions (to something). Although valenced reactions can have different magnitudes, each one is at least either positive or negative. Therefore ‘positive’ and ‘negative’ have been placed at the top of the hierarchy. At the next level, the OCC model specifies that valenced reactions can be directed at either consequences of events,³ actions of

³ Although the phrase “(un)desirable event” is used many times by OCC, events are actually always appraised with respect to their consequences. For example, an earthquake in itself does not have a valence; only the consequences of this event (e.g., valuable lessons for seismologists, property damage, loss of life)

agents, or aspects of objects. ‘Pleased’ and ‘displeased’ have been chosen by OCC to function as labels for the most general type of valenced reactions to consequences of events, because they are very neutral sounding words with respect to intensity of experience, focus of attention, motivational and behavioral effects, etc. For the same reasons, ‘approving’ and ‘disapproving’ are used as labels for the most general type of valenced reactions to actions of agents, and ‘liking’ and ‘disliking’ are used as labels for the most general type of valenced reactions to aspects of objects.

With respect to valenced reactions to consequences of events, a distinction is made based on whether the consequence in question is prospective⁴ or not. For example, learning that tomorrow it will rain is an event, but it has an undesirable consequence (e.g., the undermining of the goal to have a dry picnic) that is prospective and not actual. But this event can also have a consequence which is actual (e.g., the achievement of the goal to know the weather forecast). This differentiation on prospects then results in distinguishing between the ‘hope’ and ‘fear’ types on the one hand (e.g., Alice fears tomorrow her picnic will get wet) and the ‘joy’ and ‘distress’ types on the other hand (e.g., Alice is joyous about having learned the weather forecast).

With respect to valenced reactions to actions of agents, a distinction is made on whether the action in question has been performed by the self or by someone else. The distinction between “one’s own action” and “someone else’s action” is, however, not as simple as it may seem. A mother may be proud of the achievements of her son, even though the actions of her son are, strictly speaking, not her own. To resolve this, the OCC model uses the concept of a *cognitive unit*: the mother can consider herself and her son as part of a single cognitive unit and then, when appraising her son’s actions as praiseworthy, feel proud of the actions performed by (an agent in) the cognitive unit. This differentiation on cognitive unit then results in distinguishing between the ‘pride’ and ‘shame’ types on the one hand, and the ‘admiration’ and ‘reproach’ types on the other hand.

At this point the reader may expect there to be a branch below liking and disliking, after seeing branches being added below pleased/displeased and approving/disapproving. Indeed, in the original diagram of the OCC model (Ortony et al. 1988, p. 19), a single branch appears below liking/disliking with the emotion types ‘love’ and ‘hate’. The idea of OCC was that ‘love’ and ‘hate’ are examples of emotions of the type ‘liking’ and ‘disliking’, respectively (Ortony and Clore 2009). However, this means that the distinction between love/hate and liking/disliking does not constitute a differentiation in terms of eliciting conditions, but merely that ‘love’ is a *token* for the

Footnote 3 continued

are appraised as being desirable or undesirable. Because desirability only applies to consequences of events, every instance of the phrase “(un)desirable event” should actually be read as a shorthand for “(un)desirable consequence of an event” (Ortony and Clore 2009).

⁴ The term “prospect” (used in, e.g., hope and fear) is intentionally ambiguous: it is used to refer to both *future* events and *uncertain* (past or current) events. For example, hoping that tomorrow will be a sunny day is future-directed, whereas hoping that a mailed package has safely reached its intended recipient is uncertainty-directed. Many formalizations appear to use OCC’s notion of prospect in only one of these senses. For example, Adam et al. (2009) and Gratch and Marsella (2004) only used uncertain prospects when formalizing hope and fear, whereas Steunebrink et al. (2007) only used future prospects.

type of emotions labeled ‘liking’ and that ‘hate’ is a *token* for the type of emotions labeled ‘disliking’. So in our inheritance-based perspective, no branch has to be added below liking/disliking, because ‘love’ and ‘hate’ are not specializations of ‘liking’ and ‘disliking’ with respect to eliciting conditions.

In addition to valenced reaction to either consequences of events or actions of agents, the OCC model also considers several types of emotions arising from observing relations between the two. Specifically, these emotion types correspond to valenced reactions to consequences of events *caused* by actions of agents. The eliciting conditions of these so-called compound emotion types are conjunctions of their constituent emotion types: ‘joy’ plus ‘pride’ is ‘gratification’, ‘joy’ plus ‘admiration’ is ‘gratitude’, ‘distress’ plus ‘shame’ is ‘remorse’, and ‘distress’ plus ‘reproach’ is ‘anger’. Note, however, that this “plus” contains an implicit assertion about their (presumably causal) relation. For example, ‘anger’ is specified as “(disapproving of) someone else’s blameworthy action and (being displeased about) the *related* undesirable event”.

At the bottom of the hierarchy (of Fig. 1) are several emotions that are special types of ‘joy’ and ‘distress’. Consider first ‘satisfaction’, ‘fears-confirmed’, ‘relief’, and ‘disappointment’. Together with ‘hope’ and ‘fear’, these six emotion types are called prospect-based in the OCC model. However, the eliciting conditions of the former four emotion types are not extensions of ‘hope’ and ‘fear’, but of ‘joy’ and ‘distress’. Specifically, ‘satisfaction’, ‘fears-confirmed’, ‘relief’, and ‘disappointment’ are emotions in response to *actual* consequences of events, namely consequences signaling the confirmation or disconfirmation of a *previously prospective* consequence. The relation between, e.g., hope and disappointment is thus more of a temporal kind. For example, first Bob hopes Alice will show up for their date, but when she does not, his hope turns into disappointment. Thus ‘satisfaction’, ‘fears-confirmed’, ‘relief’, and ‘disappointment’ are not special kinds of ‘hope’ or ‘fear’, but more like continuations of ‘hope’ or ‘fear’, counting from the point when an event has been perceived that signals the confirmation or disconfirmation of the thing hoped for or feared.⁵

Next consider ‘happy-for’, ‘resentment’, ‘gloating’, and ‘pity’, called the fortunes-of-others emotion types. These are valenced reactions arising from presuming that events have consequences for others. However, for there to be, e.g., a ‘happy-for’ emotion, the consequence that is desirable for the other must also be desirable for oneself to some degree (probably because it satisfies an interest goal to wish other people well, as suggested on p. 94; Ortony et al. 1988). But if a consequence of an event is appraised as being desirable for oneself, the conditions for ‘joy’ are satisfied. So logically speaking, the eliciting conditions of ‘happy-for’ and ‘gloating’ entail those of ‘joy’, and the eliciting conditions of ‘resentment’ and ‘pity’ entail those of ‘distress’. Therefore, in an inheritance-based hierarchy of eliciting conditions, the

⁵ The dashes accompanying ‘satisfaction’, ‘fears-confirmed’, ‘relief’, and ‘disappointment’ in Fig. 1 are intentional and do not indicate a problem. For example, ‘satisfaction’ is a label for a positively valenced reaction to the confirmation of a prospective desirable consequence, but the dashes below ‘satisfaction’ are a placeholder for a *negatively* valenced reaction to the confirmation of a prospective desirable consequence. Of course, in practice such a negative reaction to something positive never occurs, and thus does not have to be labeled.

Table 1 These emotion type specifications correspond directly to Fig. 1

<i>positive</i> and <i>negative</i> are valenced reactions (to “something”)
<i>pleased</i> is being <i>positive</i> about a consequence (of an event)
<i>displeased</i> is being <i>negative</i> about a consequence (of an event)
<i>hope</i> is being <i>pleased</i> about a prospective consequence (of an event)
<i>fear</i> is being <i>displeased</i> about a prospective consequence (of an event)
<i>joy</i> is being <i>pleased</i> about an actual consequence (of an event)
<i>distress</i> is being <i>displeased</i> about an actual consequence (of an event)
<i>satisfaction</i> is <i>joy</i> about the confirmation of a prospective desirable consequence
<i>fears-confirmed</i> is <i>distress</i> about the confirmation of a prospective undesirable consequence
<i>relief</i> is <i>joy</i> about the disconfirmation of a prospective undesirable consequence
<i>disappointment</i> is <i>distress</i> about the disconfirmation of a prospective desirable consequence
<i>happy-for</i> is <i>joy</i> about a consequence (of an event) presumed to be desirable for someone else
<i>resentment</i> is <i>distress</i> about a consequence (of an event) presumed to be desirable for someone else
<i>gloating</i> is <i>joy</i> about a consequence (of an event) presumed to be undesirable for someone else
<i>pity</i> is <i>distress</i> about a consequence (of an event) presumed to be undesirable for someone else
<i>approving</i> is being <i>positive</i> about an action (of an agent)
<i>disapproving</i> is being <i>negative</i> about an action (of an agent)
<i>pride</i> is <i>approving</i> of one’s own action
<i>shame</i> is <i>disapproving</i> of one’s own action
<i>admiration</i> is <i>approving</i> of someone else’s action
<i>reproach</i> is <i>disapproving</i> of someone else’s action
<i>gratification</i> is <i>pride</i> about an action and <i>joy</i> about a related consequence
<i>remorse</i> is <i>shame</i> about an action and <i>distress</i> about a related consequence
<i>gratitude</i> is <i>admiration</i> about an action and <i>joy</i> about a related consequence
<i>anger</i> is <i>reproach</i> about an action and <i>distress</i> about a related consequence
<i>liking</i> is being <i>positive</i> about an aspect (of an object)
<i>disliking</i> is being <i>negative</i> about an aspect (of an object)

fortunes-of-others emotion types must be placed below ‘joy’ and ‘distress’, because the latter two emotion types are generalizations of them. Note that “mixed feelings” are allowed here; for example, consider the good friends Alice and Bob vying for the same position, with Alice being awarded the job. Then Bob may be distressed about having lost out but at the same time be happy for his friend Alice. There is no contradiction here; the OCC model, and our formalization of it, allows for seemingly opposite emotions to be triggered simultaneously. But note that even this by no means implies that both will be *experienced* simultaneously or with equal intensity; we are just modeling emotion *elicitation* or triggering here (though see [Steunebrink 2010](#) for an additional quantitative model of emotional experience).

The specifications of eliciting conditions resulting from reading Fig. 1 are summed up in Table 1. It is crucial to take note of Table 1 because it will serve as our guide in the formalization of emotion triggers.

3 Capturing the logical structure of the OCC model

In this section, we will make a start with the formalization of the eliciting conditions of emotions according to the OCC model. It is important to note that a distinction is made between what triggers an emotion and how an emotion is experienced. The intensity at which an emotion is felt is influenced by many factors. Emotional experience is probably multidimensional, but it is also assumed that an estimate can be made of the “overall felt intensity” of any emotion (Frijda et al. 1992). For example, a questionnaire about emotional feelings may include a question like “indicate how angry you were when hearing about the political murder on a scale from 1 to 10” and such questions are usually not difficult to answer.

In the rest of this paper, it is assumed the following relation between emotion triggering and emotion experience exists:

emotion_type is experienced if and only if

- (1) *emotion_type* has been triggered sometime in the past and
- (2) overall felt intensity of *emotion_type* is positive

With “positive” we mean having a value strictly greater than zero. Emotional intensity does not take on negative values. An emotion type for which the triggering conditions hold is not necessarily experienced, because its intensity may be too low (i.e., zero). And an emotion type which is being experienced does not have to have its triggering conditions to hold, because it may have been triggered some time in the past. In this section and the next two, we will focus on the triggering conditions of the emotion types of the OCC model. The treatment of emotional experience, as expressed in terms of triggering and intensity as above, can be found in Steunebrink (2010).

Note that in this section, all formulas are only semiformal, since no semantics are yet given, only intuitive readings. We use logical connectives with their usual interpretation and some operators with suggestive names. The idea is that the presented formulas are formal enough to adequately capture the logical structure of the psychological model, while remaining free from a commitment to an underlying formalism. In later sections we will commit to one formalism and provide a firm grounding, but it is our intention that it remains possible to plug in another formalism if desired. This way one can provide a different interpretation of the operators, without having to start from scratch regarding the overall logical structure of the model.

The structure of our formalization of emotion triggers was illustrated in Fig. 1; we will be following this figure (from top to bottom) and its textual version Table 1 in the next subsections.

3.1 General emotion types

At the most abstract level, the OCC model considers an emotion as a valenced reaction, which can either be positive or negative. So ‘positive’ and ‘negative’ are regarded by OCC as the most general, undifferentiated emotion types. In order to know which one of ‘positive’ or ‘negative’ (or both, or neither) is triggered at some point, something must be perceived and valued, which is called appraisal. We should note that

the term ‘appraisal’ can be used to mean perception as well as valuation. In order to avoid confusion, we explicitly mention perception (of an event, action, or object) as a precondition for appraisal, and use the term appraisal strictly for valuation.

We can thus (trivially) specify the triggering conditions of ‘positive’ and ‘negative’ as the perception of something good and bad, respectively:

$$\mathbf{Positive}_i^T(X) \stackrel{\text{def}}{=} \mathbf{Perceive}_i(X) \wedge \mathbf{Good}_i(X) \tag{3.1}$$

$$\mathbf{Negative}_i^T(X) \stackrel{\text{def}}{=} \mathbf{Perceive}_i(X) \wedge \mathbf{Bad}_i(X) \tag{3.2}$$

Emotions are always relative to something, and here X stands for that “something.” $\mathbf{Positive}_i^T(X)$ is read as “a positively valenced reaction to X is triggered for agent i .” The superscript “**T**” (for trigger) indicates that we are talking about eliciting conditions, in order to avoid confusion with actual experience. It is crucial to note that $\mathbf{Positive}_i^T(X)$ is *not* the same as “agent i is positive about X .” The feeling of being positive about X may manifest itself gradually over time, if at all, and may not coincide with the satisfaction of its triggering conditions, which is what $\mathbf{Positive}_i^T(X)$ expresses.

$\mathbf{Perceive}_i(X)$ is read as “agent i perceives X .” $\mathbf{Good}_i(X)$ is read as “agent i appraises X as good,” and similarly for $\mathbf{Bad}_i(X)$.

With ‘positive’ and ‘negative’ at the top of the hierarchy, the first differentiation is with respect to the object of the emotion. As previously described, the OCC model considers three types: *consequences of events*, *actions of agents*, and *aspects of objects*. We can thus define ‘perceive’ as a disjunction of perceiving either of these three:

$$\mathbf{Perceive}_i(X) \stackrel{\text{def}}{=} \mathbf{PerceiveConseq}_i(X) \vee \mathbf{PerceiveAction}_i(X) \vee \mathbf{PerceiveObject}_i(X) \tag{3.3}$$

These three perception constructs will be clarified in the next three subsections, respectively.

If a consequence of an event is appraised as being good or bad, it is said to be ‘desirable’ or ‘undesirable’, respectively. If an action of an agent is appraised as being good or bad, it is said to be ‘praiseworthy’ or ‘blameworthy’, respectively. If an aspect of an object is appraised as being good or bad, it is said to be ‘appealing’ or ‘unappealing’, respectively. ‘Good’ and ‘bad’ can thus be defined in terms of these six notions as follows:

$$\mathbf{Good}_i(X) \stackrel{\text{def}}{=} \mathbf{Des}_i(X) \vee \mathbf{Praisew}_i(X) \vee \mathbf{Appeal}_i(X) \tag{3.4}$$

$$\mathbf{Bad}_i(X) \stackrel{\text{def}}{=} \mathbf{Undes}_i(X) \vee \mathbf{Blamew}_i(X) \vee \mathbf{Unappeal}_i(X) \tag{3.5}$$

A note about the types of arguments is in order here. As noted previously, desirability is only applicable to consequences of events, praiseworthiness is only applicable to actions of agents, etc. However, the current construction says that, e.g., $\mathbf{PerceiveConseq}_i(X) \wedge \mathbf{Praisew}_i(X) \rightarrow \mathbf{Positive}_i^T(X)$ is valid. Of course, in this example, either $\mathbf{PerceiveConseq}_i(X)$ or $\mathbf{Praisew}_i(X)$ must be applied to the wrong

type of X . Therefore it is assumed that all these constructs⁶ evaluate to false if they are applied to an argument of the wrong type. This way $\mathbf{PerceiveConseq}_i(X) \wedge \mathbf{Praisew}_i(X)$ is always false and the implication is still true.

It should also be noted that none of the desirable–undesirable, praiseworthy–blameworthy, and appealing–unappealing pairs are considered to be opposites, nor are they considered to be mutually exclusive. For example, a consequence which is not desirable is not necessarily undesirable; a lack of appeal does not make something unappealing; and the exact same action can be appraised as being both praiseworthy and blameworthy. Therefore, we really need six distinct appraisal constructs here. These three pairs of appraisal operators will be clarified in the next three subsections, respectively.

3.1.1 General event-based emotion types

Here we consider the emotion types concerning consequences of events. At the top of this branch are placed the labels ‘pleased’ and ‘displeased’ (see Fig. 1). OCC consider *desirability* as the central variable measuring how positive an event is for an individual. An event that is valued negatively is called *undesirable*. As noted previously, undesirability is not the same as the absence of desirability, nor are desirability and undesirability assumed to exclude each other; they are seen as separate variables.⁷ In the following, we will treat desirability in a qualitative manner, i.e., something is either desirable or not. Of course, there can be degrees of desirability, so when we say that something is desirable, this may be read as “having strictly positive desirability,” and when we say “not desirable,” this may be read as “having zero desirability.” Analogous readings apply to undesirability.

Below is then a (semiformal) logical description of the eliciting conditions of ‘pleased’ and ‘displeased’. $\mathbf{Pleased}_i^T(c)$ should be read as “pleased about consequence c of an event is triggered for agent i ,” and similarly for ‘displeased’.

$$\mathbf{Pleased}_i^T(c) \stackrel{\text{def}}{=} \mathbf{PerceiveConseq}_i(c) \wedge \mathbf{Des}_i(c) \tag{3.6}$$

$$\mathbf{Displeased}_i^T(c) \stackrel{\text{def}}{=} \mathbf{PerceiveConseq}_i(c) \wedge \mathbf{Undes}_i(c) \tag{3.7}$$

These formulas express that the eliciting conditions of being ‘pleased’ and being ‘displeased’ have two components; namely, the perception of a consequence of an event and the appraisal of that consequence as being (un)desirable. $\mathbf{PerceiveConseq}_i(c)$ is read as “agent i perceives consequence c (of an event)” and $\mathbf{Des}_i(c)$ is read as “agent i appraises consequence c as desirable (with respect to its goals),” or, less precisely, “ i desires c .” It will be clear that ‘pleased’ and ‘displeased’ are *undifferentiated* event-based emotions, because nothing is assumed about what kind of consequence we are

⁶ That is, **PerceiveConseq**, **PerceiveAction**, **PerceiveObject**, **Des**, **Undes**, **Praisew**, **Blamew**, **Appeal**, and **Unappeal**.

⁷ OCC use undesirability as a kind of negative desirability, i.e., a consequence is undesirable if its desirability is strictly less than zero. Here we prefer to keep desirability and undesirability as separate measures, each ranging over non-negative values. But it will be clear that the two approaches do not conflict.

dealing with, nor anything about who or what caused the event, nor to whom (other than the appraising agent) the consequence applies.

3.1.2 General action-based emotion types

At the top of the branch of emotion types concerning actions of agents are placed the labels ‘approving’ and ‘disapproving’, which are regarded by OCC as the most general action-based emotion types. The OCC model considers *praiseworthiness* and *blameworthiness* to be the central variables for valuating actions of agents. Analogously to ‘pleased’ and ‘displeased’ (see above), ‘approving’ and ‘disapproving’ can be specified as perceiving an action of an agent and appraising that action as praiseworthy or blameworthy (or both, if one has conflicting standards). **Approving**_{*i*}^T(*j* : *a*) should be read as “approving of action *a* by agent *j* is triggered for agent *i*,” and similarly for ‘disapproving’.

$$\mathbf{Approving}_i^T(j : a) \stackrel{\text{def}}{=} \mathbf{PerceiveAction}_i(j : a) \wedge \mathbf{Praisew}_i(j : a) \quad (3.8)$$

$$\mathbf{Disapproving}_i^T(j : a) \stackrel{\text{def}}{=} \mathbf{PerceiveAction}_i(j : a) \wedge \mathbf{Blamew}_i(j : a) \quad (3.9)$$

PerceiveAction_{*i*}(*j* : *a*) is read as “agent *i* perceives agent *j* has performed action *a*.” **Praisew**_{*i*}(*j* : *a*) is read as “agent *i* appraises action *a* by agent *j* as praiseworthy (with respect to its standards),” and similarly for **Blamew**_{*i*}(*j* : *a*) and ‘blameworthy’.

3.1.3 General object-based emotion types

At the top of the branch of emotion types concerning aspects of objects are placed the labels ‘liking’ and ‘disliking’, which are regarded by OCC as the most general object-based emotion types. The OCC model considers *appealingness* and *unappealingness* to be the central variables for valuating aspects of objects. Just as desirability only applies to *consequences* of events, the OCC model considers appealingness to apply to *aspects* of objects. In the rest of this paper, however, we will simplify slightly by not representing aspects explicitly. This is usually not problematic, because different aspects of an object can often be regarded as objects themselves. For example, different aspects of a car, (e.g., headlights, doors, wheels) are objects themselves that can be liked or disliked. The appraisal of aspects that are not objects (e.g., the car’s color) is simply assumed to be handled implicitly by the constructs for appealingness and unappealingness when applied to the object in question.

Analogously to the event-based and action-based emotion types above, ‘liking’ and ‘disliking’ can be specified as perceiving an object and appraising that object as appealing or unappealing. **Liking**_{*i*}^T(*x*) should be read as “liking of object *x* is triggered for agent *i*,” and similarly for ‘disliking’.

$$\mathbf{Liking}_i^T(x) \stackrel{\text{def}}{=} \mathbf{PerceiveObject}_i(x) \wedge \mathbf{Appeal}_i(x) \quad (3.10)$$

$$\mathbf{Disliking}_i^T(x) \stackrel{\text{def}}{=} \mathbf{PerceiveObject}_i(x) \wedge \mathbf{Unappeal}_i(x) \quad (3.11)$$

PerceiveObject_{*i*}(*x*) is read as “agent *i* perceives object *x*.” **Appeal**_{*i*}(*x*) is read as “agent *i* appraises object *x* as appealing (with respect to its attitudes),” and similarly for **Unappeal**_{*i*}(*x*) and ‘unappealing’.

3.2 Concrete emotion types

In this section we will provide semiformal descriptions of the eliciting conditions of the third layer (of Fig. 1) of emotion types.

3.2.1 Event-based emotion types

The first differentiation with respect to event-based emotion types is on whether the consequence in question is prospective or actual. First, we will treat the case of prospective consequences of events, leading to the emotion types labeled as ‘hope’ and ‘fear’.

$$\mathbf{Hope}_i^T(c) \stackrel{\text{def}}{=} \mathbf{Pleased}_i^T(c) \wedge \mathbf{Prospective}_i(c) \tag{3.12}$$

$$\mathbf{Fear}_i^T(c) \stackrel{\text{def}}{=} \mathbf{Displeased}_i^T(c) \wedge \mathbf{Prospective}_i(c) \tag{3.13}$$

Prospective_{*i*}(*c*) is read as “agent *i* considers *c* to be a prospective consequence (of an event).” **Hope**_{*i*}^T(*c*) is then read as “hope about consequence *c* (of an event) is triggered for agent *i*,” and similarly for ‘fear’.

Next, we will treat the case of actual consequences, leading to the emotion types labeled as ‘joy’ and ‘distress’.

$$\mathbf{Joy}_i^T(c) \stackrel{\text{def}}{=} \mathbf{Pleased}_i^T(c) \wedge \mathbf{Actual}_i(c) \tag{3.14}$$

$$\mathbf{Distress}_i^T(c) \stackrel{\text{def}}{=} \mathbf{Displeased}_i^T(c) \wedge \mathbf{Actual}_i(c) \tag{3.15}$$

Actual_{*i*}(*c*) is read as “agent *i* considers *c* to be an actual consequence (of an event).” **Joy**_{*i*}^T(*c*) is then read as “joy about consequence *c* (of an event) is triggered for agent *i*,” and similarly for ‘distress’. With the danger of getting ahead of ourselves, we remark that our grounding of the **Prospective** and **Actual** will make them mutually exclusive (see Sect. 3.5).

We emphasize again that ‘joy’ and ‘distress’ are considered as nothing more than convenient labels for these emotion types. Other labels are perfectly possible as well; for example, emotions of the type labeled as ‘joy’ include contentment, delight, being glad, happiness, cheerfulness, being ecstatic, and so on and so forth. Similarly, emotions of the type labeled as ‘distress’ include sadness, upset, being distraught, shock, etc. If one further differentiates the type of event towards which one is distressed, even more specific labels can be chosen. For example, being distressed about the loss of a loved one can be labeled as ‘grief’ and being distressed about the loss of an opportunity can be labeled as ‘regret’. The OCC model does not pursue further differentiation of ‘joy’ and ‘distress’ besides the emotion types shown at the bottom of Fig. 1, but this is certainly an interesting direction for future research.

3.2.2 Attribution emotion types

The OCC model considers one differentiation in the action-based emotion types, namely in the actor. By differentiating with respect to the concept of *cognitive unit* (see Sect. 2.2), the action-based emotion types can be captured as follows:

$$\mathbf{Pride}_i^T(j : a) \stackrel{\text{def}}{=} \mathbf{Approving}_i^T(j : a) \wedge \mathbf{CogUnit}_i(j) \quad (3.16)$$

$$\mathbf{Shame}_i^T(j : a) \stackrel{\text{def}}{=} \mathbf{Disapproving}_i^T(j : a) \wedge \mathbf{CogUnit}_i(j) \quad (3.17)$$

$$\mathbf{Admiration}_i^T(j : a) \stackrel{\text{def}}{=} \mathbf{Approving}_i^T(j : a) \wedge \neg \mathbf{CogUnit}_i(j) \quad (3.18)$$

$$\mathbf{Reproach}_i^T(j : a) \stackrel{\text{def}}{=} \mathbf{Disapproving}_i^T(j : a) \wedge \neg \mathbf{CogUnit}_i(j) \quad (3.19)$$

$\mathbf{CogUnit}_i(j)$ is read as “agent i views agent j as being in a cognitive unit with itself.” $\mathbf{Pride}_i^T(j : a)$ is then read as “pride about action a of agent j is triggered for agent i ,” and similarly for ‘shame’, ‘admiration’, and ‘reproach’.

A note on the usage of $\mathbf{CogUnit}_i(j)$ versus $\neg \mathbf{CogUnit}_i(j)$ is in order here. Consider Bob watching his country’s skier win a championship; now Bob may consider his country’s skier as being in his cognitive unit and feel pride, but he may also consider his country’s skier as being outside his cognitive unit and admire the skier for winning. At first sight this simultaneous occurrence of pride and admiration may seem to be excluded here, but this is in fact not the case. What our model says is that they cannot be *triggered* at the same moment, but they can still be triggered in succession and then be *experienced* simultaneously (experience is something that persists through time; triggering is a property of one instance in time). Because the set of agents contained in one’s cognitive unit is allowed to vary over time, Bob can indeed view his country’s skier as being in a cognitive unit with himself and not so, but just not at the exact same moment in time. $\mathbf{CogUnit}_i(j)$ is like a perspective of agent i on agent j , and, in line with OCC (Ortony 2009), we maintain that an agent cannot hold two different perspectives on another agent at the same time. In order to change its perspective, an agent must perform a (mental) action, which consumes (a small amount of) time. Since emotional experience can endure over time, a pride and an admiration emotion for one and the same action can be *experienced* simultaneously even though they may have been *triggered* at different times in the past. So mixed emotions are by no means excluded by requiring $\mathbf{CogUnit}_i(j)$ to be either true or false. The fact that **Des**, **Praisew**, and **Appeal** have negative counterparts (i.e., **Undes**, **Blamew**, and **Unappeal**) is due to the fact that they are valenced predicates, which **CogUnit** is not. For example, for desirability we can have positive (**Des**), negative (**Undes**), mixed feelings (**Des** and **Undes**), or indifferent (**Des** nor **Undes**); for **CogUnit** we have just yes or no, meaning that it suffices to have **CogUnit** and its logical negation.

At this point the reader may expect to find additional terms in the specifications of the action-based emotions, to deal with aspects such as effort and responsibility. Therefore we emphasize again that we have made a very consistent separation between what we treat qualitatively, and what quantitatively. In our approach, we treat qualitatively all concepts that are used by OCC to define the *eliciting conditions* of emotions, because these conditions state whether in a certain situation and point in time emotions

are triggered, yes or no. All concepts that do not appear in the eliciting conditions, but that do affect the *experience* of emotions, are treated quantitatively (so that we will be able to quantify emotional *intensity*).⁸ In the current paper, we only formalize the *eliciting conditions* of emotions, according to the OCC model. Since concepts such as effort and responsibility do not appear in the eliciting conditions as described by OCC, they also do not appear in our formalization thereof. (OCC do identify effort and responsibility, but only as a component of the quantitative strength of praiseworthiness and blameworthiness and not as a separate factor.)

3.2.3 Attraction emotion types

The OCC model does not structure the valenced reactions to aspects of objects, even though OCC admit that momentary reactions of liking and disliking are among the most salient experiences for humans. Interestingly, they do consider one variable affecting the intensity of liking and disliking reactions (besides appealingness), namely *familiarity*, but they have chosen not to differentiate based on this variable. Differentiating based on familiarity would not be correct because the relation between familiarity and overall liking or disliking is not monotonic (Ortony et al. 1988; Ortony 2009; Ortony and Clore 2009). As is also suggested by the proverb “familiarity breeds contempt,” liking of an object can decrease when one is very familiar with it, even though initially, liking usually increases with familiarity. Indeed, in the OCC model it is suggested that the relation between familiarity and overall liking probably follows a bell shape.

3.3 Compounds

Two branches of Fig. 1 combine to form the so-called compound emotion types. These emotions arise when one focuses on both the praiseworthiness of an action and the desirability of the *related* consequences. According to OCC, the eliciting conditions of the compound emotion types are a conjunction of the eliciting conditions of an event-based emotion (‘joy’ or ‘distress’) and an action-based emotion (‘pride’, ‘shame’, ‘admiration’, or ‘reproach’), together with an assertion about their relatedness. However, realizing that an action and a consequence of an event are related may come at a later time than perceiving either the action or the consequence. Therefore, we need to be able to look back in time when describing the eliciting conditions of the compound emotion types. To this end, we use the construct **Past** φ , which asserts that φ was true sometime in the past, where, importantly, the past is understood to include the present. The compound emotion types can then be captured as follows:

$$\mathbf{Gratification}_i^T(j : a, c) \stackrel{\text{def}}{=} \mathbf{Past Pride}_i^T(j : a) \wedge \mathbf{Past Joy}_i^T(c) \\ \wedge \mathbf{PerceiveRelated}_i(j : a, c) \quad (3.20)$$

⁸ The interested reader can find an additional quantitative model on top of the presented qualitative model in Steunebrink (2010).

$$\begin{aligned} \mathbf{Remorse}_i^T(j : a, c) &\stackrel{\text{def}}{=} \mathbf{Past Shame}_i^T(j : a) \wedge \mathbf{Past Distress}_i^T(c) \\ &\quad \wedge \mathbf{PerceiveRelated}_i(j : a, c) \end{aligned} \tag{3.21}$$

$$\begin{aligned} \mathbf{Gratitude}_i^T(j : a, c) &\stackrel{\text{def}}{=} \mathbf{Past Admiration}_i^T(j : a) \wedge \mathbf{Past Joy}_i^T(c) \\ &\quad \wedge \mathbf{PerceiveRelated}_i(j : a, c) \end{aligned} \tag{3.22}$$

$$\begin{aligned} \mathbf{Anger}_i^T(j : a, c) &\stackrel{\text{def}}{=} \mathbf{Past Reproach}_i^T(j : a) \wedge \mathbf{Past Distress}_i^T(c) \\ &\quad \wedge \mathbf{PerceiveRelated}_i(j : a, c) \end{aligned} \tag{3.23}$$

PerceiveRelated_{*i*}(*j* : *a*, *c*) is read as “agent *i* perceives action *a* of agent *j* as being related to consequence *c*.” **Gratification**_{*i*}^T(*j* : *a*, *c*) is then read as “gratification about action *a* of agent *j* and the related consequence *c* is triggered for agent *i*,” and similarly for ‘remorse’, ‘gratitude’, and ‘anger’.

In order to ensure that the action appearing twice in each of these definitions is really the same action, it is assumed that all actions are unique. This can be seen as each performed action being a unique *instance* of an action. For example, *j* : *a* in **Gratification**_{*i*}^T(*j* : *a*, *c*) appears in both **Pride**_{*i*}^T(*j* : *a*) and **PerceiveRelated**_{*i*}(*j* : *a*, *c*). With this assumption of uniqueness, the action that is the object of the perceived relation (i.e., *j* : *a* in **PerceiveRelated**_{*i*}(*j* : *a*, *c*)) must really be the same action as the one which is the object of the earlier action-based emotion (e.g., *j* : *a* in **Pride**_{*i*}^T(*j* : *a*)).

3.4 Derived emotion types

In this section we will provide semiformal descriptions of the eliciting conditions of the bottom layer (of Fig. 1) of emotion types.

3.4.1 Prospect-based emotion types

Whereas ‘hope’ and ‘fear’ concern *unconfirmed* prospects of events, the OCC model also distinguishes emotion types concerning *confirmed* and *disconfirmed* prospects, namely ‘satisfaction’, ‘fears-confirmed’, ‘relief’, and ‘disappointment’. As explained in Sect. 2.2, a confirmation or disconfirmation is regarded as an actual consequence of an event, and therefore these emotion types are specializations of ‘joy’ and ‘distress’. However, they do depend on an earlier instance of ‘hope’ or ‘fear’, so the formalizations below use the **Past** operator to capture this temporal link. These four prospect-based emotion types can then be captured as follows:

$$\begin{aligned} \mathbf{Satisfaction}_i^T(c, c') &\stackrel{\text{def}}{=} \mathbf{Joy}_i^T(c) \wedge \mathbf{Past Hope}_i^T(c') \\ &\quad \wedge \mathbf{Confirms}_i(c, c') \end{aligned} \tag{3.24}$$

$$\begin{aligned} \mathbf{Fears-confirmed}_i^T(c, c') &\stackrel{\text{def}}{=} \mathbf{Distress}_i^T(c) \wedge \mathbf{Past Fear}_i^T(c') \\ &\quad \wedge \mathbf{Confirms}_i(c, c') \end{aligned} \tag{3.25}$$

$$\begin{aligned} \mathbf{Relief}_i^T(c, c') &\stackrel{\text{def}}{=} \mathbf{Joy}_i^T(c) \wedge \mathbf{Past Fear}_i^T(c') \\ &\quad \wedge \mathbf{Disconfirms}_i(c, c') \end{aligned} \tag{3.26}$$

$$\mathbf{Disappointment}_i^T(c, c') \stackrel{\text{def}}{=} \mathbf{Distress}_i^T(c) \wedge \mathbf{Past Hope}_i^T(c') \wedge \mathbf{Disconfirms}_i(c, c') \quad (3.27)$$

$\mathbf{Confirms}_i(c, c')$ is read as “agent i considers consequence c as (partially) confirming consequence c' ”, and likewise for ‘disconfirm’. $\mathbf{Satisfaction}_i^T(c, c')$ is then read as “satisfaction about consequence c confirming consequence c' is triggered for agent i ,” and similarly for ‘fears-confirmed’, ‘relief’, and ‘disappointment’.

3.4.2 Fortunes-of-others emotion types

Finally, the four so-called fortunes-of-others emotion types are also specializations of ‘joy’ and ‘distress’, as explained in Sect. 2.2. These emotion types concern consequences of events *presumed* to be desirable or undesirable for someone else. In order to capture presumptions, we introduce the **Presume** operator; $\mathbf{Presume}_i\varphi$ is read as “agent i presumes φ (to be true).” When grounding these semiformal specifications in a BDI-based logic (as we will do later), the presume operator can easily be conflated with belief. However, in order to remain independent of any underlying formalism, we stick to OCC’s phrasing at this point and use ‘presume’ as the name for the operator. The fortunes-of-others emotion types can then be captured as follows:

$$\mathbf{Happy-for}_i^T(c, j) \stackrel{\text{def}}{=} \mathbf{Joy}_i^T(c) \wedge \mathbf{Presume}_i\mathbf{Des}_j(c) \quad (3.28)$$

$$\mathbf{Pity}_i^T(c, j) \stackrel{\text{def}}{=} \mathbf{Distress}_i^T(c) \wedge \mathbf{Presume}_i\mathbf{Undes}_j(c) \quad (3.29)$$

$$\mathbf{Gloating}_i^T(c, j) \stackrel{\text{def}}{=} \mathbf{Joy}_i^T(c) \wedge \mathbf{Presume}_i\mathbf{Undes}_j(c) \quad (3.30)$$

$$\mathbf{Resentment}_i^T(c, j) \stackrel{\text{def}}{=} \mathbf{Distress}_i^T(c) \wedge \mathbf{Presume}_i\mathbf{Des}_j(c) \quad (3.31)$$

$\mathbf{Happy-for}_i^T(c, j)$ is read as “happy-for about consequence c (of an event) for agent j is triggered for agent i ,” and similarly for ‘pity’, ‘gloating’, and ‘resentment’.

It may be interesting to note that it is not required for agent i to presume that agent j is aware of the event in question as well. For example, suppose Alice has just learned that she has won a magnificent cruise for two. She may feel very happy for her husband (who she intends to take the cruise with) without him being aware of the prize yet. Of course, Alice may feel inclined to tell her husband about the prize as soon as possible, but it would be unreasonable to argue that she cannot feel happy for him before having informed him.

3.5 Properties

To finish this section, we will show some properties of the (semiformal) specifications presented in this section. So far we have ‘reduced’ the eliciting conditions of the emotion types of the OCC model to formulas involving some standard logical connectives and the following seventeen constructs:

PerceiveConseq	Des	Prospective
PerceiveAction	Undes	Actual
PerceiveObject	Praisew	CogUnit
PerceiveRelated	Blamew	Confirms
Past	Appeal	Disconfirms
	Unappeal	Presume

If the specifications presented in this section are accurate, then the eliciting conditions of the emotion types of the OCC model are constructed around no more than seventeen⁹ notions, represented by the seventeen constructs above. About half of these constructs will be grounded in dynamic doxastic logic in the next section, whereas the remaining constructs will be grounded in KARO (which is a BDI-based extension of dynamic doxastic logic) in the section after that.

In the following properties, let \vdash_T (where T stands for trigger) be a classical propositional entailment relation with formulas (3.1)–(3.31) as axioms. Furthermore, $\Gamma \vdash_T \varphi$ denotes that φ is derivable assuming Γ . Although formal proofs of propositions appearing later in this paper are provided in Appendix A, the derivations of the propositions below only involve manipulation of regular propositional connectives and therefore we deem it as unnecessary to spell out these derivations.

The following properties read exactly as the type specifications for pleased, displeased, approving, disapproving, liking, and disliking given in Table 1.

$$\vdash_T \mathbf{Pleased}_i^T(c) \leftrightarrow \mathbf{Positive}_i^T(c) \tag{3.32}$$

$$\vdash_T \mathbf{Displeased}_i^T(c) \leftrightarrow \mathbf{Negative}_i^T(c) \tag{3.33}$$

$$\vdash_T \mathbf{Approving}_i^T(i : a) \leftrightarrow \mathbf{Positive}_i^T(i : a) \tag{3.34}$$

$$\vdash_T \mathbf{Disapproving}_i^T(i : a) \leftrightarrow \mathbf{Negative}_i^T(i : a) \tag{3.35}$$

$$\vdash_T \mathbf{Liking}_i^T(x) \leftrightarrow \mathbf{Positive}_i^T(x) \tag{3.36}$$

$$\vdash_T \mathbf{Disliking}_i^T(x) \leftrightarrow \mathbf{Negative}_i^T(x) \tag{3.37}$$

For example, Table 1 states that “*pleased* is being *positive* about a consequence (of an event).” So if we put a consequence c into ‘positive’ (i.e., $\mathbf{Positive}_i^T(c)$), we should get ‘pleased’. And indeed, formula (3.32) states that $\mathbf{Positive}_i^T(c)$ is equivalent to $\mathbf{Pleased}_i^T(c)$. The other properties follow the same pattern.

The following properties state that each pair of ‘siblings’ in the third layer of Fig. 1 completely subdivide their ‘parent’.

$$\Gamma \vdash_T \mathbf{Pleased}_i^T(c) \leftrightarrow (\mathbf{Hope}_i^T(c) \vee \mathbf{Joy}_i^T(c)) \tag{3.38}$$

$$\Gamma \vdash_T \mathbf{Displeased}_i^T(c) \leftrightarrow (\mathbf{Fear}_i^T(c) \vee \mathbf{Distress}_i^T(c)) \tag{3.39}$$

$$\vdash_T \mathbf{Approving}_i^T(i : a) \leftrightarrow (\mathbf{Pride}_i^T(j : a) \vee \mathbf{Admiration}_i^T(j : a)) \tag{3.40}$$

⁹ Not counting propositional connectives.

$$\vdash_{\text{T}} \mathbf{Disapproving}_i^{\text{T}}(i : a) \leftrightarrow (\mathbf{Shame}_i^{\text{T}}(j : a) \vee \mathbf{Reproach}_i^{\text{T}}(j : a)) \quad (3.41)$$

where $\Gamma = \mathbf{PerceiveConseq}_i(c) \rightarrow (\mathbf{Actual}_i(c) \vee \mathbf{Prospective}_i(c))$.

The following properties state that ‘siblings’ in the third layer of Fig. 1 exclude each other.

$$\Gamma \vdash_{\text{T}} \neg(\mathbf{Hope}_i^{\text{T}}(c) \wedge \mathbf{Joy}_i^{\text{T}}(c)) \quad (3.42)$$

$$\Gamma \vdash_{\text{T}} \neg(\mathbf{Fear}_i^{\text{T}}(c) \wedge \mathbf{Distress}_i^{\text{T}}(c)) \quad (3.43)$$

$$\vdash_{\text{T}} \neg(\mathbf{Pride}_i^{\text{T}}(j : a) \wedge \mathbf{Admiration}_i^{\text{T}}(j : a)) \quad (3.44)$$

$$\vdash_{\text{T}} \neg(\mathbf{Shame}_i^{\text{T}}(j : a) \wedge \mathbf{Reproach}_i^{\text{T}}(j : a)) \quad (3.45)$$

where $\Gamma = \neg(\mathbf{Actual}_i(c) \wedge \mathbf{Prospective}_i(c))$. Together with the previous set of properties, this means that the differentiations directly below pleased/displeased, approving/disapproving, and liking/disliking are strict and complete. It should be noted that these properties do *not* express that, e.g., an agent cannot experience fear and distress at the same time. The OCC model explicitly allows for mixed emotions, but the question is whether this means simultaneous *elicitation* of mixed emotions or simultaneous *experience* of mixed emotions, which are two different things. We note that the *eliciting conditions*, as described by OCC, of the pairs of emotion types above are disjunct, when both are applied to one and the same action or consequence. Therefore we conclude that OCC’s view about mixed emotions pertains only to the *experience* of mixed emotions (which is not treated in this paper). So we still consider our approach, including the four properties above, to be in line with the OCC model. For example, the fact that $\mathbf{Fear}_i^{\text{T}}(c) \wedge \mathbf{Distress}_i^{\text{T}}(c)$ is a contradiction means that the perception of *one* consequence c cannot *trigger* both fear and distress (because the triggering conditions for the ‘fear’ and ‘distress’ emotion types exclude each other), but it is still very possible that first fear for c is triggered and that shortly afterwards distress about c is triggered (e.g., because the fear has been confirmed), and that after this moment both the fear and the distress are *experienced* simultaneously for some time.¹⁰

The following properties state that ‘awareness’ of what one finds desirable and undesirable leads to ‘joy’ being equivalent to “happy-for-self” and ‘distress’ being equivalent to “self-pity.”

$$\Gamma_1 \vdash_{\text{T}} \mathbf{Joy}_i^{\text{T}}(c) \leftrightarrow \mathbf{Happy-for}_i^{\text{T}}(c, i) \quad (3.46)$$

$$\Gamma_2 \vdash_{\text{T}} \mathbf{Distress}_i^{\text{T}}(c) \leftrightarrow \mathbf{Pity}_i^{\text{T}}(c, i) \quad (3.47)$$

where $\Gamma_1 = \mathbf{Des}_i(c) \rightarrow \mathbf{Presume}_i \mathbf{Des}_i(c)$ and $\Gamma_2 = \mathbf{Undes}_i(c) \rightarrow \mathbf{Presume}_i \mathbf{Undes}_i(c)$.

The following properties state that proper ‘pride’ and ‘shame’ (in the sense that the agent of the praiseworthy/blameworthy action in question is exactly the self) are equivalent to “self-approving” and “self-disapproving,” respectively.

¹⁰ A suitable definition of emotional experience allowing such mixed emotions can be found in Steunebrink (2010).

$$\Gamma \vdash_T \mathbf{Pride}_i^T(i : a) \leftrightarrow \mathbf{Approving}_i^T(i : a) \tag{3.48}$$

$$\Gamma \vdash_T \mathbf{Shame}_i^T(i : a) \leftrightarrow \mathbf{Disapproving}_i^T(i : a) \tag{3.49}$$

where $\Gamma = \mathbf{CogUnit}_i(i)$. Note that these properties read exactly as the specifications for ‘pride’ and ‘shame’ in Table 1, e.g., “pride is approving of one’s own action.”

The inheritance-based view of the eliciting conditions of emotions, as illustrated in Fig. 1, raises the expectation that each depicted emotion type implies its parent. Indeed, chains of implications such as the one below can be made for all emotion types (except for the compound emotion types¹¹). For example:

$$\begin{aligned} \vdash_T \mathbf{Gloating}_i^T(c, j) &\rightarrow \mathbf{Joy}_i^T(c) \\ \vdash_T \mathbf{Joy}_i^T(c) &\rightarrow \mathbf{Pleased}_i^T(c) \\ \vdash_T \mathbf{Pleased}_i^T(c) &\rightarrow \mathbf{Positive}_i^T(c) \end{aligned}$$

Finally, it is worth emphasizing that the OCC model does not require appraisal to be consistent. Indeed, the following propositions are *not* derivable.

$$\not\vdash_T \neg(\mathbf{Des}_i(c) \wedge \mathbf{Undes}_i(c)) \tag{3.50}$$

$$\not\vdash_T \neg(\mathbf{Praisew}_i(j : a) \wedge \mathbf{Blamew}_i(j : a)) \tag{3.51}$$

$$\not\vdash_T \neg(\mathbf{Appeal}_i(x) \wedge \mathbf{Unappeal}_i(x)) \tag{3.52}$$

So it is *not* assumed that an agent’s goals, standards, and attitudes are consistent. This implies that ‘mixed feelings’ are possible, i.e., formulas such as $\mathbf{Admiration}_i^T(j : a) \wedge \mathbf{Reproach}_i^T(j : a)$ are satisfiable. In fact, for each pair of ‘opposing’ emotion types (i.e., those sharing a box in Fig. 1), we have that their eliciting conditions do not exclude each other. With slight abuse of notation, this can be expressed as follows.

$$\not\vdash_T \neg \left(\mathbf{Emotion}_i^{+T}(X) \wedge \mathbf{Emotion}_i^{-T}(X) \right) \tag{3.53}$$

where, e.g., $\mathbf{Emotion}_i^{+T}(X)$ stands for $\mathbf{Hope}_i^T(c)$ and $\mathbf{Emotion}_i^{-T}(X)$ stands for $\mathbf{Fear}_i^T(c)$.

4 Grounding in dynamic doxastic logic

We will now introduce a formalism that will ground many of the constructs used in the previous section to (semiformally) specify the eliciting conditions of the emotion types of the OCC model. We have chosen to use dynamic doxastic logic (van Ditmarsch et al. 2007; Segerberg 1995), because this is a well-understood formalism which readily provides ways for reasoning about agents and their actions (because it is dynamic)

¹¹ Each of the compound emotion types does inherit the eliciting conditions of its parents (e.g., ‘remorse’ inherits from ‘distress’ and ‘shame’), but they are preceded by a **Past** operator because the two inherited sets of conditions do not have to be satisfied at the same time. Still, a chain of implications can be made if one allows it to be “contaminated” by a **Past** operator.

and beliefs (because it is doxastic). Furthermore, this dynamic perspective allows for a straightforward representation of events and their consequences. Although dynamic doxastic logic is not concerned with objects, we will introduce a reasonable way of representing them as well.

We emphasize that it is perfectly possible to choose another formalism than we did, as long as it supports reasoning about events and desirability, actions and praiseworthiness, and objects and appealingness, as well as some temporal constructs for the prospect-based emotions. We do not define the appraisal constructs (desirability, praiseworthiness, appealingness) in pure dynamic doxastic logic in this section, because it lacks ways of representing goals, standards, and attitudes. In Sect. 5, then, we add BDI-based constructs and finish the grounding of the specification of eliciting conditions of emotions. Also, we will not be concerned with formal semantics until Sect. 5.

4.1 Basic operators

It is assumed there exists a set ATM of atomic propositions, with typical element p . Furthermore, we use the propositional connectives \neg , \wedge , \vee , \rightarrow , and \leftrightarrow with their usual interpretation, as well as \perp for falsum and \top for verum. We then typically use φ and ψ to denote arbitrary formulas.

In dynamic doxastic logic there are of course two modal operators, namely for belief and action. They are expressed and read as follows.

$\mathbf{B}_i\varphi$: Agent i believes φ (to be true).

$[i:\alpha]\varphi$: After the execution of action α by agent i , φ holds.

In the following we will also use converse actions (denoted as α^-), which are useful for expressing what was true before the execution of an action. For example, $[i:\alpha^-]\varphi$ expresses that, if it is the case that agent i has just performed action α , then φ was true before that action. Because we will often need to express that some agent has just performed some action, we define a convenient shorthand for this, as follows:

$$\mathbf{Done}(i:\alpha) \stackrel{\text{def}}{=} \langle i:\alpha^- \rangle \top \quad (4.1)$$

$\mathbf{Done}(i:\alpha)$ can thus be read as “agent i has just performed action α .” Note that angled brackets are used to denote the dual of the action modality, as usual.

In dynamic logic, actions are used as an abstraction of time, which means that temporal operators can be interpreted over actions. We will be using the following three basic temporal and action-based operators:

Prev φ : In the previous state, φ was true. We say “*the* previous state,” because we assume a linear history (and a branching future).

Past φ : Some time in the past (*including* the present), φ was/is true. Intuitively, this comes down to $\varphi \vee \mathbf{Prev} \varphi \vee \mathbf{Prev} \mathbf{Prev} \varphi \vee \dots$, but since this is an infinite formula, the past operator cannot be defined as an abbreviation.

Fut φ : Some time in the future (*including* the present), φ may be true. Intuitively, this can be seen as an existential quantification over agents and actions (cf. formula (4.29) in Sect. 4.5).

Although the **Past** and **Fut** operators are—as usual—understood to include the present, in the following we will mostly be using the future operator in situations that exclude the present. For convenience, then, we define a strict version of the future operator as follows.

$$\mathbf{Fut}^+ \varphi \stackrel{\text{def}}{=} \neg \varphi \wedge \mathbf{Fut} \varphi \tag{4.2}$$

Fut⁺ φ is then read as “some time in the future, but not presently, φ may be true.”

In the following subsections, we will show how we define many of the constructs used in Sect. 3 to model the emotion triggers. We will define these constructs as abbreviations (i.e. using $\stackrel{\text{def}}{=}$) of the operators just introduced. Some constructs, however, will be left undefined even here. These are the appraisal operators (i.e., **Des**, **Undes**, **Praisew**, **Blamew**, **Appeal**, and **Unappeal**), confirmation operators (**Confirms**, **Disconfirms**), and the cognitive unit operator (**CogUnit**). There are separate reasons for this, which will be explained in Sect. 4.5.


The next subsections are structured as follows. First we will define the constructs used for the event-based emotion types, then those for the action-based emotion types, and then those for the object-based emotion types. Finally, several properties of the presented definitions will be discussed in Sect. 4.5.

4.2 Events and their consequences

To start simple, we conflate presuming with believing:

$$\mathbf{Presume}_i \varphi \stackrel{\text{def}}{=} \mathbf{B}_i \varphi \tag{4.3}$$

The largest branch of the OCC model is concerned with valenced reactions to events; however, events are said to always be appraised with respect to their consequences. Now let us consider the distinction between consequences and events in more detail. The usual view in dynamic logic is that the execution of an action is regarded as an event. This makes sense because in dynamic logic, time passes only through the execution of actions, i.e., through a succession of events. Here we will follow this view and only regard executions of actions as events. We then consider a consequence of an event to be anything that was not true directly before the event, but is true directly after the event. This idea can be illustrated as follows.



This figure illustrates that state w_2 is the result of event $i : \alpha$, i.e., the execution of action α by agent i . Now any formula φ that is true in state w_2 (i.e., $w_2 \models \varphi$) but

was not true in the previous state called w_1 (i.e., $w_1 \models \neg\varphi$) is considered to be a *consequence* of the event $i : \alpha$. It will be clear that an event can also have multiple or no consequences.

For convenience, then, we introduce the following construct to capture consequences of events:


$$\mathbf{New} \varphi \stackrel{\text{def}}{=} \varphi \wedge \neg\mathbf{Prev} \varphi \tag{4.4}$$

New φ reads as: φ was not true in the previous state but φ is true in the current state. Note that **Prev** ψ expresses that ψ was true before the execution of the latest action, i.e., before the latest event. Therefore, if **New** φ holds for some formula φ , then φ can be regarded as a consequence of an event. Indeed, we have that $w_2 \models \mathbf{New} \varphi$ in the illustration above.

Even though **New** φ expresses that φ is a consequence of an event, it may very well be that no agent is aware of this consequence. There must be a change in an agent’s beliefs before we can say that it perceives a consequence of an event (or anything in general). Belief changes (updates) can easily be defined using the **New** operator, as follows:

$$\mathbf{BelUpd}_i(\varphi) \stackrel{\text{def}}{=} \mathbf{New} \mathbf{B}_i \varphi \tag{4.5}$$

BelUpd _{i} (φ) reads as: the beliefs of agent i have just been updated with φ . A situation where **BelUpd** _{i} (φ) holds can be illustrated as follows:



In the state before the event $j : \alpha$, agent i does not believe φ , i.e., it envisages worlds where $\neg\varphi$ holds. In the state after the event $j : \alpha$, agent i believes φ , i.e., in all worlds it holds as possible, φ is true.¹² In that state, then, **BelUpd** _{i} (φ) is true. It should be noted that **BelUpd** _{i} (φ) says nothing about the event (e.g., the action $j : \alpha$ in the illustration above) that actually ‘caused’ the belief update; all it expresses is that *something* happened and as a consequence, agent i believes φ to be true. In other words, from the perspective of agent i , φ is a consequence of an event.

With respect to emotions concerning consequences of events, the OCC model distinguishes between the types ‘hope’ and ‘fear’ on the one hand, and ‘joy’ and ‘distress’ on

¹² With the danger of getting ahead of ourselves, we use possible world semantics to illustrate these definitions. Indeed, we will ground the belief and action modalities using possible world semantics in Sect. 5. The definitions given in the present section do not really depend on such semantics; these illustration, then, only serve to get a feeling for what the defined constructs express.

the other hand, based on whether the consequence in question is prospective or actual, respectively. As noted in Sect. 2.2, the notion of “prospect” is intentionally ambiguous; it is used to describe both future consequences and uncertain consequences. In Sect. 3 we used **Prospective**_{*i*}(φ) and **Actual**_{*i*}(φ) to express that agent *i* considers φ to be a prospective or actual consequence of an event, respectively. Using definitions similar to **BelUpd** above, we define **Prospective** and **Actual** as follows:

$$\mathbf{Prospective}_i(\varphi) \stackrel{\text{def}}{=} \mathbf{FutUpd}_i(\varphi) \vee \mathbf{UncUpd}_i(\varphi) \tag{4.6}$$

$$\mathbf{Actual}_i(\varphi) \stackrel{\text{def}}{=} \mathbf{BelUpd}_i(\varphi) \tag{4.7}$$

where

$$\mathbf{FutUpd}_i(\varphi) \stackrel{\text{def}}{=} \mathbf{New B}_i \mathbf{Fut}^+ \varphi \tag{4.8}$$

$$\mathbf{UncUpd}_i(\varphi) \stackrel{\text{def}}{=} \mathbf{New} (\neg \mathbf{B}_i \varphi \wedge \neg \mathbf{B}_i \neg \varphi) \tag{4.9}$$

We thus split the definition of **Prospective** into two cases in order to capture future as well as uncertain prospects. The definition of **Actual** is the same as **BelUpd** (see also the illustration above). The definition of **FutUpd** is also like **BelUpd** but then with φ replaced by $\neg \varphi \wedge \mathbf{Fut} \varphi$, i.e., agent *i* comes to believe that φ is not true but that there exists a future in which φ will be true. The definition of **UncUpd** (“uncertainty update”) also resembles **BelUpd**. If, in the illustration for **BelUpd** above, the left ‘cloud’ would contain either only φ ’s or only $\neg \varphi$ ’s, and the right ‘cloud’ would contain a mixture of φ ’s and $\neg \varphi$ ’s, then **UncUpd**_{*i*}(φ) would be true in the bottom right state. **UncUpd**_{*i*}(φ) thus expresses that agent *i* has just become uncertain about whether or not φ holds.

Prospective_{*i*}(φ) and **Actual**_{*i*}(φ) are now defined such that they cover both cases for perceiving consequences of events that are distinguished in the OCC model. This means that **PerceiveConseq**_{*i*}(φ), which we used to express that agent *i* perceives consequence φ of an event, can be (trivially) defined as follows:

$$\mathbf{PerceiveConseq}_i(\varphi) \stackrel{\text{def}}{=} \mathbf{Prospective}_i(\varphi) \vee \mathbf{Actual}_i(\varphi) \tag{4.10}$$

In effect, this definition specifies that perceiving a consequence of an event means either perceiving a prospective consequence or an actual consequence. Indeed, these are exactly the two cases distinguished in the OCC model. If one wishes to distinguish more kinds of consequences of events, more disjuncts covering those cases could be added.

4.3 Agents and their actions

It is assumed there exists a set AGT of agent names with typical elements *i* and *j*, and a set ACT of actions with typical element α . Furthermore, it is assumed each action is unique, i.e., it can be performed only once. This can be seen as each performed action being a unique *instance* of an action.

In the previous section, we used $\mathbf{PerceiveAction}_i(j : \alpha)$ to express that agent i perceives that agent j has performed action α . Making use of the \mathbf{BelUpd} construct introduced above, we define $\mathbf{PerceiveAction}$ as follows:

$$\mathbf{PerceiveAction}_i(j : \alpha) \stackrel{\text{def}}{=} \mathbf{BelUpd}_i(\mathbf{Past Done}(j : \alpha)) \quad (4.11)$$

Note the use of the \mathbf{Past} operator here. Because $\mathbf{Done}(j : \alpha)$ only expresses that agent j has *just* performed action α , we need the \mathbf{Past} operator to express perceptions of actions that have been performed at some arbitrary time in the past. Thus $\mathbf{PerceiveAction}_i(j : \alpha)$ does not specify when exactly agent j performed action α , just that agent i now believes it did and that i did not believe so before.

For the compound emotion types (gratification, remorse, gratitude, anger) it was necessary to express a (presumed) relation between an action of an agent and a consequence. For this we used $\mathbf{PerceiveRelated}$, which we define here using \mathbf{BelUpd} again:

$$\mathbf{PerceiveRelated}_i(j : \alpha, \varphi) \stackrel{\text{def}}{=} \mathbf{BelUpd}_i(\mathbf{Related}(j : \alpha, \varphi)) \quad (4.12)$$

$$\mathbf{Related}(i : \alpha, \varphi) \stackrel{\text{def}}{=} \mathbf{Past}(\mathbf{Done}(i : \alpha) \wedge \mathbf{New} \varphi) \quad (4.13)$$

For convenience we define a construct for relatedness separately (as it will be useful later on). $\mathbf{Related}(i : \alpha, \varphi)$ expresses that some time in the past, φ became true just when agent i had performed action α . We do not suggest this establishes a causal relationship between the action and the formula; indeed, the relation merely exists in their co-occurrence. Note that this definition correctly expresses a relation because of the assumption of uniqueness of actions. It should also be noted that the figure in Sect. 4.2 illustrates this construct; in particular, $w_2 \models \mathbf{Related}(i : \alpha, \varphi)$. With these definitions, then, $\mathbf{PerceiveRelated}_i(j : \alpha, \varphi)$ expresses that agent i perceives action α of agent j to be related to consequence φ if and only if agent i comes to believe that φ became true exactly when action α was performed by agent j .

4.4 Objects and their aspects

It is assumed there exists a set OBJ of object names with typical element x . The OCC model requires us to be able to view agents as objects, so it is required that $\text{AGT} \subseteq \text{OBJ}$. Furthermore, it is assumed that there exists an atomic proposition for each object that identifies that object, i.e., $\{\text{object}_x \mid x \in \text{OBJ}\} \subseteq \text{ATM}$. The notation object_x is used to refer to the proposition identifying x as an object. For example, if $x = \text{mona_lisa}$ ($\in \text{OBJ}$), then object_x may be the proposition $\text{mona_lisa_is_an_object}$ ($\in \text{ATM}$). Using this notation, the construct $\mathbf{PerceiveObject}$ used to capture the perception of objects can simply be defined as follows:

$$\mathbf{PerceiveObject}_i(x) \stackrel{\text{def}}{=} \mathbf{BelUpd}_i(\text{object}_x) \quad (4.14)$$

So perceiving an object is equated with a description of the object being added to the agent’s beliefs.

4.5 Properties

We have so far ‘reduced’ the eliciting conditions of the emotion types of the OCC model to formulas involving propositional connectives and operators from dynamic doxastic logic (in particular, **B**, **Prev**, **Past**, and **Fut**). The following nine constructs are still undefined:

Des	Praisew	Appeal	Confirms	CogUnit
Undes	Blamew	Unappeal	Disconfirms	

These constructs will be grounded in KARO (which extends dynamic doxastic logic) in the next section. There are several reasons why they have not been defined in this section. First, the six appraisal constructs (**Des**, **Undes**, **Praisew**, **Blamew**, **Appeal**, **Unappeal**) require the notions of goals, standards, and attitudes, which are absent in pure dynamic doxastic logic. Second, for the confirmation constructs (**Confirms**, **Disconfirms**) we want to be more precise about what is being compared first. From definitions (3.24)–(3.27) it can be observed that the things being compared for (dis)confirmation are objects of event-based emotion types, which means that they must be related to goals. Although it may be possible to define what it means for one consequence to confirm or disconfirm another consequence just in propositional logic, we can be more precise if we know what goals look like. Therefore, we postpone defining **Confirms** and **Disconfirms** until we have formalized goals. Third, a proper formalization of the notion of cognitive unit (as expressed by **CogUnit**) would require substantially more (psychological) research, so we will leave this aspect mostly open in this paper.

To finish this section, we will show some properties of the formalization so far. In the following properties, let \vdash_{DD} (where DD stands for dynamic doxastic logic) be a normal modal entailment relation with formulas (3.1)–(3.31) and (4.1)–(4.14) as axioms. Then the following propositions are derivable.

$$\vdash_{DD} \neg(\mathbf{FutUpd}_i(\varphi) \wedge \mathbf{UncUpd}_i(\varphi)) \tag{4.15}$$

$$\vdash_{DD} \neg(\mathbf{BelUpd}_i(\varphi) \wedge \mathbf{UncUpd}_i(\varphi)) \tag{4.16}$$

$$\Gamma \vdash_{DD} \neg(\mathbf{BelUpd}_i(\varphi) \wedge \mathbf{FutUpd}_i(\varphi)) \tag{4.17}$$

$$\Gamma \vdash_{DD} \neg(\mathbf{Actual}_i(\varphi) \wedge \mathbf{Prospective}_i(\varphi)) \tag{4.18}$$

$$\Gamma \vdash_{DD} \neg(\mathbf{Hope}_i^T(\varphi) \wedge \mathbf{Joy}_i^T(\varphi)) \tag{4.19}$$

$$\Gamma \vdash_{DD} \neg(\mathbf{Fear}_i^T(\varphi) \wedge \mathbf{Distress}_i^T(\varphi)) \tag{4.20}$$

$$\vdash_{DD} \mathbf{Pleased}_i^T(c) \leftrightarrow (\mathbf{Hope}_i^T(c) \vee \mathbf{Joy}_i^T(c)) \tag{4.21}$$

$$\vdash_{DD} \mathbf{Displeased}_i^T(c) \leftrightarrow (\mathbf{Fear}_i^T(c) \vee \mathbf{Distress}_i^T(c)) \tag{4.22}$$

where $\Gamma = \neg(\mathbf{B}_i\varphi \wedge \mathbf{B}_i\neg\varphi)$. The first three properties state that **FutUpd**, **UncUpd**, and **BelUpd** are mutually exclusive. Because **Prospective** is defined in terms of **FutUpd** and **UncUpd**, this immediately results in the fourth property. The fourth property then leads to the fifth and sixth properties, because (4.18) was exactly the assumption needed for properties (3.42) and (3.43). The seventh and eighth properties are the same as properties (3.38) and (3.39), except without needing assumptions due to the way **PerceiveConseq** (4.10) has been defined.

At this point it may be instructive to write out the definition of the eliciting conditions of one emotion type and see our modular approach at work. Consider $\mathbf{Hope}_i^T(c)$, which has been defined as $\mathbf{Pleased}_i^T(c) \wedge \mathbf{Prospective}_i(c)$, which is equal to $(\mathbf{PerceiveConseq}_i(c) \wedge \mathbf{Des}_i(c)) \wedge \mathbf{Prospective}_i(c)$, which is the same as $((\mathbf{Prospective}_i(\varphi) \vee \mathbf{Actual}_i(\varphi)) \wedge \mathbf{Des}_i(c)) \wedge \mathbf{Prospective}_i(c)$. But this is equal to simply $\mathbf{Des}_i(c) \wedge \mathbf{Prospective}_i(c)$. So although at first there may appear to be some redundancy, our setup ensures that all inappropriate terms can be removed.

Below are several propositions showing how the dynamic and temporal operators introduced in this section interact. Because we have not yet introduced formal semantics, they cannot be called properties yet. The appendix provides formal proofs using the semantics introduced in Sect. 5.

$$[i : \alpha]\mathbf{Done}(i : \alpha) \tag{4.23}$$

$$\mathbf{Done}(i : \alpha) \rightarrow \mathbf{Prev} \top \tag{4.24}$$

$$\mathbf{Prev} \top \rightarrow (\mathbf{Prev} \varphi \leftrightarrow \neg\mathbf{Prev} \neg\varphi) \tag{4.25}$$

$$\neg\mathbf{Prev} \top \wedge \varphi \rightarrow \mathbf{New} \varphi \tag{4.26}$$

$$\mathbf{Done}(i : \alpha) \wedge \varphi \rightarrow \mathbf{Prev} \langle i : \alpha \rangle \varphi \tag{4.27}$$

$$\mathbf{Prev} [i : \alpha]\varphi \wedge \mathbf{Done}(i : \alpha) \rightarrow \varphi \tag{4.28}$$

$$\mathbf{Fut} \varphi \leftrightarrow \varphi \vee \langle i_1 : \alpha_1 \rangle \dots \langle i_n : \alpha_n \rangle \varphi \quad (\exists i_1, \dots, i_n, \exists \alpha_1, \dots, \alpha_n) \tag{4.29}$$

The first proposition reads rather tautologically: after the execution of action α by agent i , i has done α . The second proposition states that, if an action has just been done, there must exist a previous state. The third propositions states that **Prev** is its own dual, provided that there exists a previous state. The fourth proposition states that if no previous state exists, all formulas that are true now are also ‘new’. The fifth proposition states that everything that is true now must previously have been a possible result of the last performed action. As usual, $\langle \cdot \rangle$ abbreviates $\neg[\cdot]\neg$. The sixth proposition states that all necessary results of the last performed action must be true now. Finally, the seventh proposition states that **Fut** is like an existential quantification over agents and actions. Note that we slightly abuse notation here, because we do not have quantification in our object language.

5 Grounding in a BDI-based logic

In this section we use KARO (Meyer et al. 1999; Meyer 2006) as a framework for grounding the formalization of the emotions of the OCC model. The KARO framework is a mixture of dynamic logic, epistemic/doxastic logic, and several additional (modal)

operators for dealing with the motivational aspects of artificial agents. KARO was originally proposed as a specification logic for rational agents. It was thus designed to serve a purpose similar to that of the logics of Cohen and Levesque (1990) and Rao and Georgeff (1998). A crucial difference with these approaches is that KARO is primarily based on dynamic logic (Harel et al. 2000) rather than temporal logic (Emerson 1990). So one could view KARO as a kind of BDI (belief, desire, intention) logic based on dynamic logic. Although the specification of informational and motivational attitudes (such as knowledge and beliefs, and desires, goals, and commitments, respectively) had been the main aim for devising KARO (van der Hoek et al. 1999; Meyer et al. 1999), the logic has also proven to be applicable for the description of agent behavior, more in general. For example, in Meyer (2006) it has been used to specify four basic emotion types. Here we present a modest extension of KARO, including operators concerning past states, such that all emotion types of the OCC model can be specified in this framework. We emphasize again that our modular approach makes it possible to use other frameworks to ground our formalization as well.

5.1 Adding motivational constructs

In Sect. 4.3 we introduced the set ACT of actions. From ACT we derive the set PLANS, consisting of sequential compositions of actions, with typical element π . PLANS is the smallest set such that $ACT \subseteq PLANS$ and if $\alpha \in ACT$ and $\pi \in PLANS$, then $(\alpha; \pi) \in PLANS$.

The notation of the dynamic operator is extended to sequential compositions of actions and its dual, as follows:

$$[i : (\alpha; \pi)]\varphi \stackrel{\text{def}}{=} [i : \alpha][i : \pi]\varphi \tag{5.1}$$

$$\langle i : \pi \rangle\varphi \stackrel{\text{def}}{=} \neg[i : \pi]\neg\varphi \tag{5.2}$$

In KARO, there are quite a few operators for expressing motivational attitudes of agents (cf. van der Hoek et al. 1999; Meyer et al. 1999). Here we use three of them, namely for (achievement) goals, abilities, and commitments:

G_iφ: Agent *i* has (achievement) goal φ . Here φ represents a state of affairs which agent *i* wants to achieve. A goal φ is said to have been *achieved* when agent *i* believes φ holds, i.e., when **B_iφ** is true.

A_iπ: Agent *i* has the ability to do π .

Com_i(π): Agent *i* is committed to doing π .

Using these operators, several constructs are defined expressing (possible) motivations of an agent:

$$\mathbf{PracPoss}_i(\pi, \varphi) \stackrel{\text{def}}{=} \mathbf{A}_i\pi \wedge \langle i : \pi \rangle\varphi \tag{5.3}$$

$$\mathbf{Can}_i(\pi, \varphi) \stackrel{\text{def}}{=} \mathbf{B}_i\mathbf{PracPoss}_i(\pi, \varphi) \tag{5.4}$$

$$\mathbf{PossIntend}_i(\pi, \varphi) \stackrel{\text{def}}{=} \mathbf{Can}_i(\pi, \varphi) \wedge \mathbf{B}_i(\mathbf{G}_i\varphi \wedge \neg\varphi) \tag{5.5}$$

An agent has the *practical possibility* to perform an action/plan π to bring about φ iff it has the ability to perform π and doing so may bring about φ . An agent *can* perform π to bring about φ iff it believes it has the practical possibility to do so. An agent has the *possible intention* to perform π to accomplish φ iff it *can* do so and it believes φ is an unachieved goal.

For technical convenience, we do not require goals to be automatically dropped the moment they are achieved (cf. [Cohen and Levesque 1990](#)), assuming this can only be done using a special drop action (see also below formula (5.34)). The second part of the definition of *possible intention* above guards against already achieved goals.

5.2 Semantics

KARO is a dynamic doxastic logic, so we will introduce belief models and action models. Belief models are of the form $M = \langle S, R, V \rangle$, where

S is a non-empty set of states (or ‘possible worlds’).

$R = \{R_i \mid i \in \text{AGT}\}$ is a set of accessibility relations on S , one for each agent name, hence the notation R_i . So $R_i \subseteq S \times S$ for each $R_i \in R$.

$V : S \rightarrow 2^{\text{ATM}}$ is a valuation function, indicating which atomic propositions hold per state.

As is common in doxastic logic, each belief–accessibility relation R_i is required to be serial, transitive, and Euclidean, i.e., the modal logic KD45 is used for belief models.

The semantics of actions are defined *over* the Kripke models of belief, as actions may change the mental states of agents. Action models are of the form $\mathcal{M} = \langle S, \mathcal{R}, \text{Aux}, \text{Emo} \rangle$, where

S is a non-empty set of possible model–state pairs, where a model is of the form M as above and a state is from S therein. That is, if $(M, s) \in S$ and $M = \langle S, R, V \rangle$ then it must be that $s \in S$.

$\mathcal{R} = \{\mathcal{R}_{i:\alpha} \mid i \in \text{AGT}, \alpha \in \text{ACT}\}$ is a set of accessibility relations on S . Each transition is labeled with an agent name and an action, hence the notation $\mathcal{R}_{i:\alpha}$. So $\mathcal{R}_{i:\alpha} \subseteq S \times S$ for each $\mathcal{R}_{i:\alpha} \in \mathcal{R}$.

$\text{Aux} = \langle \text{Goals}, \text{Caps}, \text{Agd} \rangle$ is a structure of auxiliary functions, indicating per agent and model–state pair which goals (*Goals*), capabilities (*Caps*), and commitments (*Agd*) the agent has.

$\text{Emo} = \langle \text{Des}, \text{Undes}, \text{Praisew}, \text{Blamew}, \text{Appeal}, \text{Unappeal}, \text{CogUnit} \rangle$ is a structure of appraisal and judgment functions, indicating per agent and model–state pair how that agent appraises consequences (*Des*, *Undes*), actions of agents (*Praisew*, *Blamew*), and objects (*Appeal*, *Unappeal*), and how it judges cognitive units (*CogUnit*).

In order to have a branching future and a single history, it is required that $\bigcup \mathcal{R}$ is injective. This ensures that any model–state pair can be reached from at most one other model–state pair. Note however that this does not exclude parallel actions. For example, if $((M, s), (M', s')) \in \mathcal{R}_{i:\alpha}$ and $((M, s), (M', s')) \in \mathcal{R}_{j:\beta}$, then state (M', s') is a result of the parallel execution of $i:\alpha$ and $j:\beta$ in state (M, s) . With respect to the semantics of converse actions, let $\mathcal{R}_{i:\alpha^-} = (\mathcal{R}_{i:\alpha})^-$, as usual.

In the previous sections it was assumed that actions are unique. We are now in a position to formalize this assumption as a constraint on \mathcal{R} , as follows:

$$\forall \mathcal{R}_{i:\alpha} \in \mathcal{R} : \forall ((M, s), (M', s')) \in \mathcal{R}_{i:\alpha} : \mathcal{R}_{i:\alpha} \cap (\mathcal{S}' \times \mathcal{S}') = \emptyset \quad (5.6)$$

where $\mathcal{S}' = \{(M'', s'') \mid ((M', s'), (M'', s'')) \in (\bigcup \mathcal{R})^*\}$.¹³ This constraint can be read as follows. If state (M', s') is a result of action α by agent i ($((M, s), (M', s')) \in \mathcal{R}_{i:\alpha}$), then no possible future state after (M', s') (collected in \mathcal{S}') can be reachable by i doing α again ($\mathcal{R}_{i:\alpha} \cap (\mathcal{S}' \times \mathcal{S}') = \emptyset$). It may be interesting to note that this constraint implies that $\bigcup \mathcal{R}$ must be free of circles.

In line with previous work (Steunebrink et al. 2007, 2008), we define goal formulas as non-empty, consistent conjunctions of literals. This way goals can easily be broken up in subgoals; in particular, every non-empty ‘subconjunction’ of a goal formula is considered to be a subgoal. For example, if $p \wedge \neg q \wedge r$ is a goal, then $p, \neg q, r, p \wedge \neg q, p \wedge r, \neg q \wedge r$, and $p \wedge \neg q \wedge r$ are subgoals. Goal formulas are drawn from the set CCL which is defined as follows.

$$\text{LIT} = \text{ATM} \cup \{\neg p \mid p \in \text{ATM}\} \quad (5.7)$$

$$\text{CSL} = \{\Phi \mid \emptyset \subset \Phi \subseteq \text{LIT}, \Phi \not\vdash_{PC} \perp\} \quad (5.8)$$

$$\text{CCL} = \{\bigwedge \Phi \mid \Phi \in \text{CSL}\} \quad (5.9)$$

where PC stands for Propositional Calculus (so each conjunction in CCL is consistent). So LIT is the set of literals, CSL is the set of consistent sets of literals, and CCL is the set of consistent conjunctions of literals.

The mappings of the three auxiliary functions are now as follows. $Goals : \text{AGT} \times \mathcal{S} \rightarrow 2^{\text{CCL}}$ is a function returning the set of goals an agent has per model–state pair; $Caps : \text{AGT} \times \mathcal{S} \rightarrow 2^{\text{PLANS}}$ is a function that returns the set of actions that an agent is capable of performing per model–state pair; and $Agd : \text{AGT} \times \mathcal{S} \rightarrow 2^{\text{PLANS}}$ is a function that returns the set of actions that an agent is committed to (are on its ‘agenda’) per model–state pair. Note that it is *not* assumed that goals are mutually consistent.

Finally, the mappings of the appraisal and judgment functions are as follows. $Des, Undes : \text{AGT} \times \mathcal{S} \rightarrow 2^{\mathcal{L}}$ for desirability and undesirability (where \mathcal{L} is the set of all well-formed formulas); $Praisew, Blamew : \text{AGT} \times \mathcal{S} \rightarrow 2^{\text{AGT} \times \text{ACT}}$ for praiseworthiness and blameworthiness; $Appeal, Unappeal : \text{AGT} \times \mathcal{S} \rightarrow 2^{\text{OBJ}}$ for appealingness and unappealingness; and $CogUnit : \text{AGT} \times \mathcal{S} \rightarrow 2^{\text{AGT}}$ for cognitive unit.

5.3 Interpretation in KARO

We now have all ingredients necessary for the interpretation of formulas, presented below. Formulas are interpreted in state s of model M , where $(M, s) \in \mathcal{S}$. It should be noted that the pair (M, s) is itself a state of model \mathcal{M} , i.e., belief models (M) are nested in action models (\mathcal{M}) . Strictly speaking, we should write $(\mathcal{M}, (M, s)) \models \dots$, but we drop the \mathcal{M} for notational convenience.

¹³ A^* denotes the reflexive transitive closure of relation A .

Let $\mathcal{M} = \langle \mathcal{S}, \mathcal{R}, \text{Aux}, \text{Emo} \rangle$, $(M, s) \in \mathcal{S}$, and $M = \langle S, R, V \rangle$; formulas are then interpreted as follows.

Basic connectives:

$$\begin{aligned} M, s \models p & \quad \text{iff} & \quad p \in V(s) \quad \text{for } p \in \text{ATM} \\ M, s \models \neg\varphi & \quad \text{iff} & \quad \text{not } M, s \models \varphi \\ M, s \models \varphi \wedge \psi & \quad \text{iff} & \quad M, s \models \varphi \quad \text{and} \quad M, s \models \psi \end{aligned}$$

Mental attitudes:

$$\begin{aligned} M, s \models \mathbf{B}_i\varphi & \quad \text{iff} & \quad \forall s' \in S : (s, s') \in R_i \quad \text{implies} \quad M, s' \models \varphi \\ M, s \models \mathbf{G}_i\varphi & \quad \text{iff} & \quad \varphi \in \text{Goals}(i)(M, s) \\ M, s \models \mathbf{A}_i\pi & \quad \text{iff} & \quad \pi \in \text{Caps}(i)(M, s) \\ M, s \models \mathbf{Com}_i(\pi) & \quad \text{iff} & \quad \pi \in \text{Agd}(i)(M, s) \end{aligned}$$

Dynamic and temporal operators:

$$\begin{aligned} M, s \models \langle i : \pi \rangle \varphi & \quad \text{iff} & \quad \exists (M', s') \in \mathcal{S} : M', s' \models \varphi \\ & & \quad \text{and} \quad ((M, s), (M', s')) \in \mathcal{R}_i : \pi \\ M, s \models \mathbf{Fut} \varphi & \quad \text{iff} & \quad \exists (M', s') \in \mathcal{S} : M', s' \models \varphi \\ & & \quad \text{and} \quad ((M, s), (M', s')) \in (\bigcup \mathcal{R})^* \\ M, s \models \mathbf{Past} \varphi & \quad \text{iff} & \quad \exists (M', s') \in \mathcal{S} : M', s' \models \varphi \\ & & \quad \text{and} \quad ((M', s'), (M, s)) \in (\bigcup \mathcal{R})^* \\ M, s \models \mathbf{Prev} \varphi & \quad \text{iff} & \quad \exists (M', s') \in \mathcal{S} : M', s' \models \varphi \\ & & \quad \text{and} \quad ((M', s'), (M, s)) \in (\bigcup \mathcal{R}) \end{aligned}$$

Appraisal and judgment operators:

$$\begin{aligned} M, s \models \mathbf{Des}_i(\varphi) & \quad \text{iff} & \quad \varphi \in \text{Des}(i)(M, s) \\ M, s \models \mathbf{Undes}_i(\varphi) & \quad \text{iff} & \quad \varphi \in \text{Undes}(i)(M, s) \\ M, s \models \mathbf{Praisew}_i(j : \alpha) & \quad \text{iff} & \quad (j, \alpha) \in \text{Praisew}(i)(M, s) \\ M, s \models \mathbf{Blamew}_i(j : \alpha) & \quad \text{iff} & \quad (j, \alpha) \in \text{Blamew}(i)(M, s) \\ M, s \models \mathbf{Appeal}_i(x) & \quad \text{iff} & \quad x \in \text{Appeal}(i)(M, s) \\ M, s \models \mathbf{Unappeal}_i(x) & \quad \text{iff} & \quad x \in \text{Unappeal}(i)(M, s) \\ M, s \models \mathbf{CogUnit}_i(j) & \quad \text{iff} & \quad j \in \text{CogUnit}(i)(M, s) \end{aligned}$$

For clarity of presentation the interpretation of $\langle i : \pi \rangle \varphi$ is given. The future, past, and previous operators are interpreted over all action-accessibility relations in \mathcal{R} , which is done by taking the union $\bigcup \mathcal{R}$. $(\bigcup \mathcal{R})^*$ is then a relation connecting model–state pairs reachable in zero or more actions of agents. Notice that (M, s) and (M', s') are reversed for the future and past operators. As usual, $\models \varphi$ is used to denote that φ is valid, i.e., φ is satisfied in all possible model–state pairs.

5.4 Appraisal operators

Until now we have deferred the problem of specifying appraisal to the functions *Des*, *Undes*, *Praisew*, *Blamew*, *Appeal*, and *Unappeal*. In this section we have introduced (achievement) goals, which will allow us to give meaning to these functions, but only to a limited degree, because there are many other kinds of concerns, such as norms, interests, preservation, etc. Therefore, we will not *define* these appraisal functions; instead, we will *constrain* them such that they capture appraisal for agents with achievement goals only. The idea is that one can simply add more constraints to these appraisal function if the framework is enriched with more kinds of concerns.

Before introducing these constraints on the appraisal functions, we define two helper sets for matching subparts of goals and inverting goals, respectively:

$$\text{SUB} = \left\{ \left(\bigwedge \Phi_1, \bigwedge \Phi_2 \right) \left| \Phi_1, \Phi_2 \in \text{CSL}, \Phi_1 \subseteq \Phi_2 \right. \right\} \tag{5.10}$$

$$\text{INV} = \left\{ \left(\bigwedge \Phi_1, \bigwedge \Phi_2 \right) \left| \Phi_1 \in \text{CSL}, \Phi_2 = \{ \text{neg}(\varphi) \mid \varphi \in \Phi_1 \} \right. \right\} \tag{5.11}$$

where $\text{neg}(p) = \neg p$ and $\text{neg}(\neg p) = p$. Recall that we require achievement goals to be consistent conjunctions of literals. The set SUB will then be convenient for making subgoals, and the set INV will be convenient for inverting entire (sub)goals. For example, if $p_1 \wedge \neg p_2$ is a goal, then p_1 and $\neg p_2$ (and $p_1 \wedge \neg p_2$) are subgoals, and $\neg p_1 \wedge p_2$ is the inverted goal.

We will start now with the simplest case of desirability. Let *Des* be constrained such that every subgoal is desirable, and let *Undes* be constrained such that every inverted subgoal is undesirable:

$$\text{Des}(i)(M, s) \supseteq \text{SUB} \circ \text{Goals}(i)(M, s) \tag{5.12}$$

$$\text{Undes}(i)(M, s) \supseteq \text{INV} \circ \text{SUB} \circ \text{Goals}(i)(M, s) \tag{5.13}$$

where \circ is used to denote relation composition.¹⁴ These two constraints read just as they are written: *Des* contains all subgoals, and *Undes* contains all inverted subgoals. Let us illustrate the above constraint on *Undes*: if $\psi = p \wedge \neg q \wedge r$ is a goal, then the inverted subgoal $\varphi = \neg p \wedge q$ is undesirable, because $(\neg p \wedge q \wedge \neg r, p \wedge \neg q \wedge r) \in \text{INV}$ and $(\neg p \wedge q, \neg p \wedge q \wedge \neg r) \in \text{SUB}$, so $(\neg p \wedge q, p \wedge \neg q \wedge r) \in \text{SUB} \circ \text{INV}$.

The OCC model considers praiseworthiness (and its negative counterpart blame-worthiness) to be determined with respect to one’s standards. However, OCC note that the praiseworthiness of an action may be evaluated with respect to the desirability of the events caused by that action. Since we do not explicitly consider standards here, we will constrain *Praisew* and *Blamew* using *Des* and *Undes*, respectively. Of course, different or additional constraints may be studied if an explicit representation of standards were added to the logical framework.

¹⁴ In the case where R is binary and Y in unary, the composition $R \circ Y$ is defined as: $R \circ Y = \{x \mid \exists y : (x, y) \in R, y \in Y\}$.

We now constrain *Praisew* and *Blamew* as follows. An action of an agent is appraised as praiseworthy or blameworthy when the appraising agent believes that the action is related to a desirable or undesirable consequence, respectively.

$$\begin{aligned}
 &Praisew(i)(M, s) \supseteq \\
 &\quad \{(j, \alpha) \mid \exists \varphi \in Des(i)(M, s) : M, s \models \mathbf{B}_i \mathbf{Related}(j : \alpha, \varphi)\} \quad (5.14)
 \end{aligned}$$

$$\begin{aligned}
 &Blamew(i)(M, s) \supseteq \\
 &\quad \{(j, \alpha) \mid \exists \varphi \in Undes(i)(M, s) : M, s \models \mathbf{B}_i \mathbf{Related}(j : \alpha, \varphi)\} \quad (5.15)
 \end{aligned}$$

where **Related** is as defined in (4.13). We did not spell out the condition $M, s \models \mathbf{B}_i \mathbf{Related}(j : \alpha, \varphi)$ using the semantics because that would have made these constraints considerably more difficult to read without becoming more enlightening.

We emphasize that this constraint only makes an action praiseworthy to the agent who desires its consequence(s); indeed, other agents may well disagree and find the action in question blameworthy. Consider the following example from the OCC model (Ortony et al. 1988, p. 135, italics in original) to illustrate this point: “[T]he issue is not whether or not the blamed or praised person is *in fact* wholly or partially responsible for the situation, and hence subject to praise or blame for his or her action. For example, a subway vigilante who kills a teenager in the New York subway will not experience shame unless he believes that what he did was blameworthy. In other words, he has to believe that his action violated some standard or principle in which he believes. If in fact he views his action as upholding some principle (regardless of the opinions of others) he is more likely to feel pride than shame. On the other hand, we as observers can quite consistently view his action as blameworthy by our standards, (...)” In line with this view, the constraints above only specify praiseworthiness and blameworthiness in terms of an agent’s own desires, allowing other agents with other desires to disagree and assign praiseworthiness and blameworthiness differently.

According to the OCC model, appealingness and unappealingness are determined with respect to one’s attitudes. Here we will not constrain the functions *Appeal* and *Unappeal* with respect to objects that are not agents. Instead, we will only consider the appealingness of agents, as follows. An agent is appealing to the appraising agent if it has ever performed a praiseworthy action, and unappealing if it has ever performed a blameworthy action:

$$Appeal(i)(M, s) \supseteq \{j \mid \exists \alpha : M, s \models \mathbf{Past Approving}_i^T(j, \alpha)\} \quad (5.16)$$

$$Unappeal(i)(M, s) \supseteq \{j \mid \exists \alpha : M, s \models \mathbf{Past Disapproving}_i^T(j, \alpha)\} \quad (5.17)$$

Note that the definitions of ‘approving’ (3.8) and ‘disapproving’ (3.9) include the perception of the praiseworthy or blameworthy action. Again, we did not spell out these conditions using the semantics in order to keep the constraints concise and easy to read. Although these constraints specify that the sets of appealing and unappealing agents can only grow, it is important to keep in mind that this does not mean that every agent will constantly feel liking and disliking towards more and more agents over time. The eliciting conditions of liking and disliking also require the (mental) perception of another agent or object, limiting the (re)triggering of liking and disliking

emotions. Moreover, the actual *experience* of liking or disliking and their *intensities* depend on further factors such as the quantitative strengths of (un)appealingness of and familiarity with the other agent or object (Ortony et al. 1988; Steunebrink 2010).

It should be emphasized that we have started out with just achievement goals as the only concerns of agents. We have then defined (or rather, constrained) desirability and undesirability in terms of achievement goals, then defined praiseworthiness and blameworthiness in terms of desirability and undesirability, and then defined appealingness and unappealingness in terms of praiseworthiness and blameworthiness.

5.5 Cognitive unit

We do not constrain an agent’s judgment of when it considers itself to be in a cognitive unit with another agent, except that we require each agent to at least be in a cognitive unit with itself.¹⁵ This is expressed by the following constraint:

$$CogUnit(i)(M, s) \supseteq \{i\} \tag{5.18}$$

This constraint ensures that **Approving**_{*i*}^T(*i* : α) is equivalent to **Pride**_{*i*}^T(*i* : α) and that **Disapproving**_{*i*}^T(*i* : α) is equivalent to **Shame**_{*i*}^T(*i* : α), as anticipated in Sect. 3 (cf. formulas (3.48) and (3.49)).

5.6 Confirmation and disconfirmation

The only constructs yet undefined are **Confirms** and **Disconfirms**. Recall that **Confirms**_{*i*}(φ, ψ) expresses that agent *i* considers consequence φ to confirm consequence ψ , and likewise for disconfirmation. Now that we have restricted concerns of agents to achievement goals only and defined achievement goals as conjunctions of literals, representing (dis)confirmation has become quite straightforward.

For convenience, we first introduce some additional syntax. We will be using the operator \sqsubseteq as the syntactic counterpart of the set SUB. The interpretation of \sqsubseteq is thus as follows:

$$M, s \models \varphi \sqsubseteq \psi \quad \text{iff} \quad (\varphi, \psi) \in \text{SUB}$$

$\varphi \sqsubseteq \psi$ is then read as “ φ is a (logical) part of ψ .” Furthermore, we add a syntactical variant of INV. For a conjunction or literals φ , writing $\overline{\varphi}$ in the object language means the ‘inverse’ of φ in the sense of INV. In other words, for all $(\varphi, \psi) \in \text{INV}$, $\psi = \overline{\varphi}$. Note that INV is symmetric, so that $\overline{\overline{\varphi}} = \varphi$, as expected.

¹⁵ Even this constraint may be too strong in general, because it may preclude a kind of ‘insanity’ where one does not consider the self as the (cognitive) author of one’s own actions.

Using these new constructs, we define **Confirms** and **Disconfirms** as follows:

$$\mathbf{Confirms}_i(\varphi, \psi) \stackrel{\text{def}}{=} \mathbf{B}_i(\varphi \sqsubseteq \psi) \tag{5.19}$$

$$\mathbf{Disconfirms}_i(\varphi, \psi) \stackrel{\text{def}}{=} \mathbf{B}_i(\overline{\varphi} \sqsubseteq \psi) \tag{5.20}$$

These definitions express that a consequence φ confirms another consequence ψ when φ is a (logical) part of ψ . It is incorrect to require that φ be (logically) stronger than ψ , because the idea of ‘confirms’ is that it must also account for partial confirmations. For example, suppose Alice learns that a plane carrying four of her relatives has crashed; she will then fear they have all perished but hope for survivors. When later she learns that two of her relatives have survived the crash, this will both partially confirm her fear and partially confirm her hope. To account for partial confirmations, then, we use the construct $\varphi \sqsubseteq \psi$.

It should be noted that it is not impossible to define confirmation more generally. For example, “ φ (partially) confirms ψ ” can also be expressed as “ $\psi \models_{\text{CL}} \varphi$,” i.e., ψ logically entails φ . But then we would need a construct for representing entailment in the object language. In effect, \sqsubseteq is a kind of entailment relation ($\varphi \sqsubseteq \psi$ implies $\psi \models_{\text{CL}} \varphi$), but restricted to goal formulas (conjunctions of literals).

5.7 Properties

To finish this section, we will discuss some properties of the formalization in KARO of the eliciting conditions of emotions of the OCC model. As usual, $\models \varphi$ expresses that the formula φ is a validity, i.e., every state of every model satisfies φ . All definitions presented in Sects. 3, 4, and 5 are assumed to be in effect. Proofs of the properties below can be found in Appendix A.

The following properties show how the appraisal operators stem from just (achievement) goals and beliefs. Of course, concerns other than achievement goals can influence desirability, praiseworthiness, and appealingness, but we have restricted our study of appraisal in this paper.

$$\models \mathbf{G}_i\varphi \wedge \psi \sqsubseteq \varphi \rightarrow \mathbf{Des}_i(\psi) \wedge \mathbf{Undes}_i(\overline{\psi}) \tag{5.21}$$

$$\models \mathbf{Des}_i(\varphi) \wedge \mathbf{B}_i\mathbf{Related}(j : \alpha, \varphi) \rightarrow \mathbf{Praisew}_i(j : \alpha) \tag{5.22}$$

$$\models \mathbf{Undes}_i(\varphi) \wedge \mathbf{B}_i\mathbf{Related}(j : \alpha, \varphi) \rightarrow \mathbf{Blamew}_i(j : \alpha) \tag{5.23}$$

$$\models \mathbf{Past\ Approving}_i^T(j : \alpha) \rightarrow \mathbf{Appeal}_i(j) \tag{5.24}$$

$$\models \mathbf{Past\ Disapproving}_i^T(j : \alpha) \rightarrow \mathbf{Unappeal}_i(j) \tag{5.25}$$

Note that these properties correspond directly to the constraints specified in Sect. 5.4. The notation $\overline{\psi}$ was explained in Sect. 5.6. We emphasize again that desirability and undesirability, praiseworthiness and blameworthiness, and appealingness and unappealingness are not mutually exclusive; see properties (3.50), (3.51), and (3.52) in Sect. 3.5. Furthermore, desirability and undesirability are not assumed to be individually

consistent either. For example, $\mathbf{Des}_i(\varphi) \wedge \mathbf{Des}_i(\neg\varphi)$ and $\mathbf{Undes}_i(\varphi) \wedge \mathbf{Undes}_i(\neg\varphi)$ are not contradictions.

The following properties are restatements of (3.42), (3.43), (3.48), and (3.49), respectively.

$$\models \neg(\mathbf{Hope}_i^T(\varphi) \wedge \mathbf{Joy}_i^T(\varphi)) \quad (5.26)$$

$$\models \neg(\mathbf{Fear}_i^T(\varphi) \wedge \mathbf{Distress}_i^T(\varphi)) \quad (5.27)$$

$$\models \mathbf{Pride}_i^T(i : a) \leftrightarrow \mathbf{Approving}_i^T(i : a) \quad (5.28)$$

$$\models \mathbf{Shame}_i^T(i : a) \leftrightarrow \mathbf{Disapproving}_i^T(i : a) \quad (5.29)$$

Previously, additional assumptions were required to make these propositions derivable. In our formalization in KARO, we have turned these assumptions into constraints, making them truly properties.

In the following, we drop the agent subscripts (e.g., i and j) to ease notation; all terms requiring one are assumed to have the same agent subscript. The four properties below express the triggering conditions for the event-based (and self-based) emotion types in BDI-like terms, i.e., in terms of beliefs and goals.

$$\models \mathbf{G}\varphi \wedge \psi \sqsubseteq \varphi \wedge \mathbf{BelUpd}(\psi) \rightarrow \mathbf{Joy}^T(\psi) \quad (5.30)$$

$$\models \mathbf{G}\varphi \wedge \psi \sqsubseteq \varphi \wedge \mathbf{BelUpd}(\bar{\psi}) \rightarrow \mathbf{Distress}^T(\bar{\psi}) \quad (5.31)$$

$$\models \mathbf{G}\varphi \wedge \psi \sqsubseteq \varphi \wedge \mathbf{BelUpd}(\neg\psi \wedge \mathbf{Fut}\psi) \rightarrow \mathbf{Hope}^T(\psi) \quad (5.32)$$

$$\models \mathbf{G}\varphi \wedge \psi \sqsubseteq \varphi \wedge \mathbf{BelUpd}(\neg\bar{\psi} \wedge \mathbf{Fut}\bar{\psi}) \rightarrow \mathbf{Fear}^T(\bar{\psi}) \quad (5.33)$$

The first property states that ‘joy’ is triggered with respect to ψ if ψ is a subgoal of the agent and it has just updated its beliefs with ψ (i.e., subgoal ψ has just been achieved). Analogously, the second property states that ‘distress’ is triggered with respect to an inverted subgoal $\bar{\psi}$ if the subgoal ψ has just been undermined (i.e., part of ψ had previously been achieved but the agent now believes the inverse $\bar{\psi}$ to be true). The third and fourth properties have similar readings, save for being future-directed.

Of course it must be recognized that the current formalization also has its limitations. The OCC model describes emotions with respect to events, actions, and objects, and therefore the elicitation of emotions is naturally described in terms of the perception of events, actions, and objects. However, this emphasis on perception fails to incorporate changes in appraisal of goals, standards, and attitudes as triggers for emotions. For example, the perception of a desired consequence can trigger joy, but a new desire for a known consequence does not trigger joy in the current formalization. Formally, we now have:

$$\models \mathbf{New}\mathbf{B}\varphi \wedge \mathbf{Des}(\varphi) \rightarrow \mathbf{Joy}^T(\varphi)$$

$$\not\models \mathbf{B}\varphi \wedge \mathbf{New}\mathbf{Des}(\varphi) \rightarrow \mathbf{Joy}^T(\varphi)$$

but the latter should intuitively be valid as well. Of course it is possible to define the emotion triggers such that changes in appraisal are taken into account, but this will be a topic of future work.

The final properties that we will discuss relate intention with tracking of goal achievements and undermining. Given several (reasonable) assumptions, the notion of intention as used in KARO is related with a simultaneous elicitation of ‘pride’, ‘joy’, and ‘gratification’. Specifically, if an agent (possibly) intends to perform action α to achieve goal φ , then after actually performing α , ‘pride’ about having done so will be triggered, as well as ‘joy’ about the achievement, and ‘gratification’ about the action leading to the achievement.

$$\Gamma \models \mathbf{PossIntend}(\alpha, \varphi) \rightarrow [\alpha](\mathbf{Pride}^T(\alpha) \wedge \mathbf{Joy}^T(\varphi) \wedge \mathbf{Gratification}^T(\alpha, \varphi)) \tag{5.34}$$

where Γ is the following set of assumptions:

- $\mathbf{B}[\alpha]\psi \rightarrow [\alpha]\mathbf{B}\psi$: α is accordant, i.e., the agent does not forget the results of α .
- $\langle \alpha \rangle \psi \rightarrow [\alpha]\psi$: action α is deterministic.
- $\neg(\mathbf{Done}(\alpha) \wedge \mathbf{Done}(\beta))$ for $\alpha \neq \beta$: the agent cannot perform actions in parallel to α .
- $\mathbf{BG}\psi \rightarrow \mathbf{G}\psi$: believed goals are true goals.
- $\mathbf{Prev}\mathbf{G}\psi \wedge \neg\mathbf{G}\psi \rightarrow \mathbf{Done}(\mathbf{drop}(\psi))$: only ‘drop’ actions can remove goals.
- $\neg\mathbf{PossIntend}(\mathbf{drop}(\psi), \psi)$: the agent never intends to achieve a goal by dropping it.

Interestingly, in contrast to ‘pure’ KARO, the current framework allows us to reason about *subgoals*. This makes it possible to define a less strict version of **PossIntend**; namely, one expressing that an agent can achieve a subgoal with some action or plan (in contrast to a *complete* goal as required by **PossIntend**). Analogously, we can define a construct expressing that an agent can undermine a subgoal with some action or plan. These two constructs are defined below as **PossAch** and **PossUnd**, respectively (cf. formula (5.5) in Sect. 5.1).

$$\mathbf{PossAch}_i(\pi, \psi, \varphi) \stackrel{\text{def}}{=} \mathbf{Can}_i(\pi, \psi) \wedge \mathbf{B}_i(\mathbf{G}_i\varphi \wedge \psi \sqsubseteq \varphi \wedge \overline{\psi}) \tag{5.35}$$

$$\mathbf{PossUnd}_i(\pi, \psi, \varphi) \stackrel{\text{def}}{=} \mathbf{Can}_i(\pi, \psi) \wedge \mathbf{B}_i(\mathbf{G}_i\varphi \wedge \overline{\psi} \sqsubseteq \varphi \wedge \overline{\psi}) \tag{5.36}$$

PossAch_{*i*}(π, ψ, φ) is read as “agent *i* can possibly achieve subgoal ψ of goal φ with plan π ,” and **PossUnd**_{*i*}(π, ψ, φ) is read as “agent *i* can possibly undermine subgoal $\overline{\psi}$ of goal φ with plan π ” (here $\overline{\psi}$ is thus an inverted subgoal of φ). It may be interesting to note that **PossAch**_{*i*}(π, φ, φ) implies **PossIntend**_{*i*}(π, φ). Using **PossAch** we can strengthen property (5.34), and using **PossUnd** we can add an analogous case for the simultaneous elicitation of ‘shame’, ‘distress’, and ‘remorse’, as follows.

$$\Gamma \models \mathbf{PossAch}(\alpha, \psi, \varphi) \rightarrow [\alpha](\mathbf{Pride}^T(\alpha) \wedge \mathbf{Joy}^T(\psi) \wedge \mathbf{Gratification}^T(\alpha, \psi)) \tag{5.37}$$

$$\Gamma \models \mathbf{PossUnd}(\alpha, \psi, \varphi) \rightarrow [\alpha](\mathbf{Shame}^T(\alpha) \wedge \mathbf{Distress}^T(\psi) \wedge \mathbf{Remorse}^T(\alpha, \psi)) \tag{5.38}$$

where Γ is as in property (5.34).

6 Related work

In this section we discuss several related attempts at adopting psychological models of emotions for modeling artificial agents. We will discuss similarities and differences with the presented approach.

6.1 Previous work

In previous work, [Dastani and Meyer \(2006\)](#) and [Meyer \(2006\)](#) proposed a functional approach to describe the role of four basic emotions in practical reasoning. According to this functional approach, an agent is assumed to execute domain actions in order to reach its goals. The effects of these domain actions cause and/or influence the elicitation of emotions according to a human-inspired model. These emotions in turn influence the deliberation operations of the agent, functioning as heuristics for determining which domain actions have to be chosen next, which completes the circle.

The specification and implementation of emotions carried out by [Dastani and Meyer \(2006\)](#) and [Meyer \(2006\)](#) follows Oatley and Jenkins' (1996) model of emotions. In contrast to our approach of capturing a broad and complete¹⁶ range of emotion types, they consider only four emotions: *happy*, *sad*, *angry*, and *fearful*. Each emotion functions as a *label* of an aspect of an agent's cognitive state. The deliberation of an agent then behaves in accordance with heuristics associated with these four emotions. Later we have extended this approach by showing how interaction between *hope* and *fear* can influence an agent's deliberation ([Steunebrink et al. 2007](#)).

In other previous work, we have shown how emotional experience can be modeled on top of a formalization of emotion triggers ([Steunebrink 2010](#); [Steunebrink et al. 2008](#)). We have then also shown how emotion regulation can be modeled on top of a formalization of emotional experience ([Steunebrink 2010](#); [Steunebrink et al. 2009](#)). This was done by introducing a notion of *action tendency* into the formalization, which indicates which action(s) an agent tends to perform to mitigate negative emotional experience. The present paper, however, is our first complete presentation of our formalization of the eliciting conditions of the emotion types of the OCC model.

6.2 Another formalization of the OCC model

The construction of a complete formalization of the OCC model in agent logic has previously been attempted by [Adam et al. \(2009\)](#).¹⁷ Our approach is similar to Adam's formalization in the sense that both use BDI-like logics (belief, desire, intention) to formalize the emotions of the OCC model and that both approaches are based on modal logic. Below we will briefly discuss major differences between the presented formalization of the OCC model and the one by Adam.

¹⁶ That is, complete with respect to one psychological model of emotions, namely the OCC model.

¹⁷ In the following, we simply use "Adam" to refer to [Adam et al. \(2009\)](#).

Just like us, Adam aims to be “as faithful as possible” to the OCC model. However, Adam’s formalization of OCC’s emotion types has been tailored to their BDI-based logical framework. In contrast, our formalization proceeds in three stages, where the first stage captures the logical structure of the OCC model and only the last stage commits to BDI. Furthermore, Adam’s logical framework incorporates several very strong assumptions. For example, desires and ideals are assumed never to change and to be free of contradictions (thus excluding many forms of ‘mixed feelings’); agents are assumed to have complete introspection with respect to their desires; and all actions are assumed to be deterministic, public, and accordant (i.e. no forgetting of effects). By refraining from making such assumptions, we believe our formalization is able to account for more situations in which emotions can arise (according to psychology).

Some of Adam’s definitions of emotions do not capture all aspects of what is supposed to be formalized. For example, Adam’s formalization of hope and fear does not account for future-directed prospects (only current uncertainty); ‘easy’ actions preclude pride and shame; and partial (dis)confirmations cannot trigger satisfaction, fears-confirmed, relief, or disappointment. Admittedly, the OCC model may be implicit or ambiguous with respect to these and other aspects, but ideally, the process of formalization should explicate such issues and offer clarifications.

There is some confusion in Adam’s formalization between emotion elicitation and experience. Adam claims to formalize the eliciting conditions of emotions (as do we), and the action-based emotions indeed appear to incorporate a trigger, namely in the form of the perception of an action. However, Adam’s formalizations of the event-based emotions do not incorporate any triggers. For example, joy is defined as $Joy_i\varphi \stackrel{\text{def}}{=} Bel_i\varphi \wedge Des_i\varphi$, but this expresses a ‘state of joy’ more than a trigger for joy. (Also note that this definition does not force φ to represent a consequence of an event.) Indeed, in the text Adam often identifies the satisfaction of an emotion formula with feeling the emotion in question. When Adam defines the compound emotions simply as conjunctions (e.g., $Gratification_i(i : \alpha, \varphi) \stackrel{\text{def}}{=} Pride_i(i : \alpha, \varphi) \wedge Joy_i\varphi$), it is then unclear what $Gratification_i(i : \alpha, \varphi)$ actually represents because it mixes triggering ($Pride_i(i : \alpha, \varphi)$) and experience ($Joy_i\varphi$). In our approach, we have made a clear distinction between emotion elicitation (treated in this paper) and experience (treated in [Steunebrink 2010](#)) in order to avoid such confusion.

Finally, Adam’s formalization renders a number of properties of emotions that we find too strong. For example, Adam proves that $\vdash \neg(Joy_i\varphi \wedge Distress_i\varphi)$ and similarly for all pairs of opposing emotions applied to the same argument(s). Such formulas are not valid in our formalization because we allow goals, standards, and attitudes to be inconsistent. However, if their consistency would be adopted as a constraint, it would indeed be provable in our framework that opposing emotion triggers contradict. Furthermore, Adam derives complete introspection of emotions (i.e. $\vdash Emotion_i\varphi \leftrightarrow Bel_iEmotion_i\varphi$ and $\vdash \neg Emotion_i\varphi \leftrightarrow Bel_i\neg Emotion_i\varphi$). However, if $Emotion_i\varphi$ is supposed to be a formalization of the eliciting conditions of $Emotion$, as Adam intends, then we find this unintuitive; one does not have to be aware of what triggered an emotion. On the other hand, it is intuitive to suppose that one is aware of what one does and does not feel, i.e., if $Emotion_i\varphi$ were to represent the subjective experience of $Emotion$. Again, because of confusion between elicitation and experi-

ence in Adam's formalization, it is difficult to judge the status of these introspection properties.

6.3 A computational model of emotions

Gratch and Marsella (2004) have been working on a computational framework for modeling emotions. The framework is claimed to be domain-independent and they have implemented a process model, called EMA after the title of Lazarus (1994), for social training applications. The appraisal process used in EMA is inspired by the OCC model. As with our approach, the cognitive reasoning aspects of EMA are represented using BDI concepts and the emphasis of appraisal is on goal attainment.

In contrast to our approach, Gratch and Marsella take a computational, quantitative approach towards modeling appraisal. Specifically, the eliciting conditions of emotions modeled in EMA are based on quantitative measures of, e.g., desirability and likelihood. The calculation of these quantitative measures is facilitated by the usage of subjective probabilities for beliefs and assignment of utilities to states. However, precise triggering conditions for all emotions are not provided, so it is hard to judge how strictly Gratch and Marsella follow psychological models of emotions and how they deviate from or extend these. For example, in line with the OCC model, likelihood of a desirable event is given as a precondition for hope. However, in EMA likelihood of an event is equated with the believed probability of the event, such that likelihood can also be used as a precondition for joy (in particular, if the likelihood of the event equals one). But like Adam's notion of expectation, such a definition of likelihood models uncertainty about the current state but not prospects about *future* events.

Concerning the three main topics of modeling emotions discussed in the introduction (i.e. appraisal, experience, and regulation), Gratch and Marsella distinguish only between appraisal and regulation. Probably owing to their computational, quantitative approach, emotional experience appears to be merged with appraisal in EMA. Their main focus is on modeling how emotions influence behavior, emphasizing modeling of coping strategies for artificial agents. Behavioral effects of emotions have not been treated in this paper, but a detailed analysis using the presented framework can be found in Steunebrink (2010).

7 Conclusions

In this paper we have studied the OCC model (Ortony et al. 1988) of emotions and proceeded with the formalization of the eliciting conditions of emotions according to this psychological model. We have carried out this formalization in three stages. First, we have captured OCC's specifications of eliciting conditions in a semiformal manner, thus without committing to a particular formalism and semantics. Second, we have shown how OCC's main notions of events, consequences, actions, agents, and beliefs can be captured in dynamic doxastic logic. Third, we have represented OCC's main appraisal notions in the KARO framework, which is a BDI-based extension of dynamic doxastic logic, thereby firmly grounding the preceding two stages. It should be noted that the collection of emotion triggers that are satisfied in a certain state

should *not* be regarded as a representation of an agent's full emotional state. Rather, it only represents which *new* emotions are triggered in that state, without specifying whether these newly triggered emotions are (or will ever be) experienced.

The idea is that the emotion triggers function as input for an additional quantitative model of emotions. Such a quantitative model should specify how triggered emotions are *experienced*. Elsewhere (Steunebrink (2010) and Steunebrink et al. (2008)), we have explained how emotional experience can be modeled on top of qualitative models of emotion triggers similar to the one presented in this paper. This was done by introducing functions representing different parameters of emotional experience. By setting thresholds for these parameters, different emotion words of the same emotion type can be modeled. For example, 'annoyed', 'livid', and 'outraged' can each be represented in the logic as different emotions of the type 'anger'.

Finally, it should be specified *what to do* with experienced emotions. Previous work on formalizing the behavioral effects of emotions either used a subset of the emotions as described in psychological model of Oatley and Jenkins (Dastani and Meyer 2006; Meyer 2006), or a subset of the OCC model (Steunebrink et al. 2007, 2009). We are currently continuing in the line of Steunebrink et al. (2009) to formally specify emotion-based action tendency for all emotion types presented in this paper. Unfortunately, psychological literature on emotion regulation (i.e. the effect of elicited emotions on behavior) does not provide classifications and schemes as clear as those on appraisal, such as the OCC model. (For an overview of psychological research on the subject, we refer the reader to Gross and Thompson (2007).) Thus to proceed with formalizing in the direction of specifying the effects of emotions on the behavior and decision making of agents will require more creativity on the part of logicians and computer scientists. It should be noted that we have not included details on how to integrate experience and regulation in the presented formalization of emotions in order to limit the length of this paper. However, research in this direction can be found in other work (Steunebrink 2010; Steunebrink et al. 2009, 2010).

Our future work on emotions is twofold. On the one hand, we are continuing our work on refining the formalization of quantitative aspects of emotions and specifying their effects on behavior and decision making. On the other hand, we are working on an implementation of this formal model of emotions on top of the interpreter of an agent programming language. This way we expect to validate the added value of emotions on the decision making and believability of artificial agents.

Appendix A: Proofs of propositions

All proofs use the semantics given in Sect. 5.2.

Proposition (4.23) Take an arbitrary agent i and action α . Now for all $((M, s), (M', s')) \in \mathcal{R}_{i:\alpha}$ we obviously have that $((M', s'), (M, s)) \in \mathcal{R}_{i:\alpha^-}$ and $M, s \models \top$, i.e., $M', s' \models \langle i : \alpha^- \rangle \top$, which is the same as $M', s' \models \mathbf{Done}(i : \alpha)$. Because (M', s') was an arbitrary $\mathcal{R}_{i:\alpha}$ -successor of (M, s) , we have that $M, s \models [i : \alpha] \mathbf{Done}(i : \alpha)$. Because (M, s) was arbitrary, $[i : \alpha] \mathbf{Done}(i : \alpha)$ is valid. \square

Proposition (4.24) Assume $M, s \models \mathbf{Done}(i : \alpha)$ for arbitrary $(M, s), i, \alpha$. Then there exists a model–state pair (M', s') such that $((M', s'), (M, s)) \in \mathcal{R}_{i:\alpha}$. Because $\mathcal{R}_{i:\alpha} \subseteq \bigcup \mathcal{R}$, $((M', s'), (M, s)) \in (\bigcup \mathcal{R})$. Because $M', s' \models \top$, we have that $M, s \models \mathbf{Prev} \top$. Because (M, s) was arbitrary, $\mathbf{Done}(i : \alpha) \rightarrow \mathbf{Prev} \top$ is valid. \square

Proposition (4.25) Assume $M, s \models \mathbf{Prev} \top$ for arbitrary (M, s) . This implies that $\exists(M', s') : ((M', s'), (M, s)) \in (\bigcup \mathcal{R})$. But $\bigcup \mathcal{R}$ is constrained to be injective, so $\exists(M', s') : ((M', s'), (M, s)) \in (\bigcup \mathcal{R})$ and $M', s' \models \varphi$ is true if and only if $\forall(M', s') : ((M', s'), (M, s)) \in (\bigcup \mathcal{R})$ implies $M', s' \models \varphi$ is true, i.e., $M, s \models \mathbf{Prev} \varphi \leftrightarrow \neg \mathbf{Prev} \neg \varphi$. Because (M, s) was arbitrary, $\mathbf{Prev} \top \rightarrow (\mathbf{Prev} \varphi \leftrightarrow \neg \mathbf{Prev} \neg \varphi)$ is valid. And because $\mathbf{Prev} \varphi$ implies $\mathbf{Prev} \top$, this proposition can be rewritten as $\mathbf{Prev} \varphi \leftrightarrow \neg \mathbf{Prev} \neg \varphi \wedge \mathbf{Prev} \top$. \square

Proposition (4.26) If $\neg \mathbf{Prev} \top$ holds, then for any φ , $\neg \mathbf{Prev} \varphi$ holds. $\mathbf{New} \varphi$ is defined as $\varphi \wedge \neg \mathbf{Prev} \varphi$, so $\varphi \wedge \neg \mathbf{Prev} \top \rightarrow \mathbf{New} \varphi$ is valid. \square

Proposition (4.27) Assume $M, s \models \mathbf{Done}(i : \alpha)$ and $M, s \models \varphi$ for arbitrary $(M, s), i, \alpha, \varphi$. Let (M', s') be the model–state pair such that $((M', s'), (M, s)) \in \mathcal{R}_{i:\alpha}$ (there can only be one such (M', s') because $\bigcup \mathcal{R}$ is injective). Now $M', s' \models \langle i : \alpha \rangle \varphi$ and therefore $M, s \models \mathbf{Prev} \langle i : \alpha \rangle \varphi$. Because (M, s) was arbitrary, $\mathbf{Done}(i : \alpha) \wedge \varphi \rightarrow \mathbf{Prev} \langle i : \alpha \rangle \varphi$ is valid. \square

Proposition (4.28) Assume $M, s \models \mathbf{Prev} [i : \alpha] \varphi$ and $M, s \models \mathbf{Done}(i : \alpha)$ for arbitrary $(M, s), i, \alpha, \varphi$. Let (M', s') be the model–state pair such that $((M', s'), (M, s)) \in \mathcal{R}_{i:\alpha}$ (there can only be one such (M', s') because $\bigcup \mathcal{R}$ is injective). Now $M', s' \models [i : \alpha] \varphi$ and therefore $M, s \models \varphi$. Because (M, s) was arbitrary, we have that $\mathbf{Prev} [i : \alpha] \varphi \wedge \mathbf{Done}(i : \alpha) \rightarrow \varphi$ is valid. \square

Proposition (4.29) The expression $M, s \models \mathbf{Fut} \varphi$ is interpreted as $\exists(M', s') \in \mathcal{S} : ((M, s), (M', s')) \in (\bigcup \mathcal{R})^*$ and $M', s' \models \varphi$. This is the same as $M, s \models \varphi$ or $\exists n \in \mathbb{N} : \exists i_0, \dots, i_n \in \text{AGT} : \exists \alpha_0, \dots, \alpha_n \in \text{ACT} : \exists((M, s), (M', s')) \in \mathcal{R}_{i_0:\alpha_0} \circ \dots \circ \mathcal{R}_{i_n:\alpha_n} : M', s' \models \varphi$, i.e., $M, s \models \varphi \vee \langle i_1 : \alpha_1 \rangle \dots \langle i_n : \alpha_n \rangle \varphi$. So $\mathbf{Fut} \varphi \leftrightarrow \varphi \vee \exists i_0, \dots, i_n \exists \alpha_0, \dots, \alpha_n (\langle i_1 : \alpha_1 \rangle \dots \langle i_n : \alpha_n \rangle \varphi)$ is valid (although strictly speaking we do not have quantification in our object language). \square

Proposition (5.21)–(5.25) These propositions follow immediately from constraints (5.12)–(5.17). \square

Proposition (5.26) and (5.27) These propositions follow immediately from propositions (4.19) and (4.20) and the fact that $\neg(\mathbf{B}_i \varphi \wedge \mathbf{B}_i \neg \varphi)$ follows from the seriality of R_i . \square

Proposition (5.28) and (5.29) These propositions follow immediately from propositions (3.48) and (3.49) and constraint (5.18). \square

Proposition (5.30)–(5.33) $\mathbf{Joy}_i^T(\psi)$ is equivalent to $\mathbf{Des}_i(\psi) \wedge \mathbf{Actual}_i(\psi)$ and $\mathbf{Distress}_i^T(\bar{\psi})$ is equivalent to $\mathbf{Undes}_i(\bar{\psi}) \wedge \mathbf{Actual}_i(\bar{\psi})$. By Proposition (5.21), $\mathbf{G}_i \varphi \wedge \psi \sqsubseteq \varphi$ implies $\mathbf{Des}_i(\psi)$ and $\mathbf{Undes}_i(\bar{\psi})$. Because $\mathbf{Actual}_i(\psi)$ is equivalent to

$\mathbf{BelUpd}_i(\psi), \mathbf{G}_i\varphi \wedge \psi \sqsubseteq \varphi \wedge \mathbf{BelUpd}_i(\psi)$ implies $\mathbf{Joy}_i^T(\psi)$ and $\mathbf{G}_i\varphi \wedge \psi \sqsubseteq \varphi \wedge \mathbf{BelUpd}_i(\bar{\psi})$ implies $\mathbf{Distress}_i^T(\bar{\psi})$. Propositions (5.32) and (5.33) follow analogously, by noting that $\mathbf{BelUpd}_i(\mathbf{Fut}^+\psi)$ is equivalent to $\mathbf{FutUpd}_i(\psi)$, which implies $\mathbf{Prospective}_i(\psi)$. \square

Proposition (5.34) Assume $M, s \models \mathbf{PossIntend}_i(\alpha, \varphi)$ for arbitrary $(M, s), i, \alpha, \varphi$. Take an arbitrary model–state pair (M', s') such that $((M, s), (M', s')) \in \mathcal{R}_{i:\alpha}$. To prove: $M', s' \models \mathbf{Pride}_i^T(i : \alpha) \wedge \mathbf{Joy}_i^T(\varphi) \wedge \mathbf{Gratification}_i^T(i : \alpha, \varphi)$, i.e., $M', s' \models \mathbf{PerceiveAction}_i(i : \alpha) \wedge \mathbf{Praisew}_i(i : \alpha) \wedge \mathbf{CogUnit}_i(i) \wedge \mathbf{PerceiveConseq}_i(\varphi) \wedge \mathbf{Actual}_i(\varphi) \wedge \mathbf{Des}_i(\varphi) \wedge \mathbf{PastPride}_i^T(i : \alpha) \wedge \mathbf{PastJoy}_i^T(\varphi) \wedge \mathbf{PerceiveRelated}_i(i : \alpha, \varphi)$. We can immediately cross out several of the conjuncts: $\mathbf{CogUnit}_i(i)$ follows directly from constraint (5.18); by definition (4.10), $\mathbf{PerceiveConseq}_i(\varphi)$ follows from $\mathbf{Actual}_i(\varphi)$; and $\mathbf{PastPride}_i^T(i : \alpha)$ and $\mathbf{PastJoy}_i^T(\varphi)$ will follow automatically because $\psi \rightarrow \mathbf{Past}\psi$ is a validity. Writing out definitions, we have to prove $M', s' \models \mathbf{NewB}_i\mathbf{PastDone}(i : \alpha) \wedge \mathbf{NewB}_i\varphi \wedge \mathbf{NewB}_i\mathbf{Related}(i : \alpha, \varphi) \wedge \mathbf{Des}_i(\varphi) \wedge \mathbf{Praisew}_i(i : \alpha)$. If $M, s \models \neg\psi$ and $M', s' \models \psi$, then $M', s' \models \mathbf{New}\psi$. So for state (M, s) we have to show:

- (i) $M, s \models \neg\mathbf{B}_i\mathbf{PastDone}(i : \alpha)$
- (ii) $M, s \models \neg\mathbf{B}_i\varphi$
- (iii) $M, s \models \neg\mathbf{B}_i\mathbf{Related}(i : \alpha, \varphi)$

and for state (M', s') we have to show:

- (a) $M', s' \models \mathbf{B}_i\mathbf{PastDone}(i : \alpha)$
- (b) $M', s' \models \mathbf{B}_i\varphi$
- (c) $M', s' \models \mathbf{B}_i\mathbf{Related}(i : \alpha, \varphi)$
- (d) $M', s' \models \mathbf{Des}_i(\varphi)$
- (e) $M', s' \models \mathbf{Praisew}_i(i : \alpha)$

By Proposition (5.22), (e) is implied by (d) and (c). Because $\psi \rightarrow \mathbf{Past}\psi$ is valid, so is $\mathbf{B}_i\psi \rightarrow \mathbf{B}_i\mathbf{Past}\psi$, which means that it suffices to show for (M', s') that:

- (A) $M', s' \models \mathbf{B}_i\mathbf{Done}(i : \alpha)$
- (B) $M', s' \models \mathbf{B}_i\varphi$
- (C) $M', s' \models \mathbf{B}_i\neg\mathbf{Prev}\varphi$
- (D) $M', s' \models \mathbf{Des}_i(\varphi)$

From Proposition (4.23), we have that $[i : \alpha]\mathbf{Done}(i : \alpha)$ is valid. Then by necessitation $\mathbf{B}_i[i : \alpha]\mathbf{Done}(i : \alpha)$ is also valid, and by the assumption that $\mathbf{B}_i[i : \alpha]\psi \rightarrow [i : \alpha]\mathbf{B}_i\psi$, $[i : \alpha]\mathbf{B}_i\mathbf{Done}(i : \alpha)$ is also valid. But if $M, s \models [i : \alpha]\mathbf{B}_i\mathbf{Done}(i : \alpha)$, then $M', s' \models \mathbf{B}_i\mathbf{Done}(i : \alpha)$, which proves (A). $M, s \models \mathbf{PossIntend}_i(\alpha, \varphi)$ implies $M, s \models \mathbf{B}_i(i : \alpha)\varphi$, which implies $M, s \models \mathbf{B}_i[i : \alpha]\varphi$ (because α was assumed to be deterministic), which implies $M, s \models [i : \alpha]\mathbf{B}_i\varphi$ (because α was assumed to be accordant). But then $M', s' \models \mathbf{B}_i\varphi$, which proves (B). It is easy to verify that $\neg\varphi \rightarrow [i : \alpha]\neg\mathbf{Prev}\varphi$ is a validity; then $\mathbf{B}_i\neg\varphi \rightarrow \mathbf{B}_i[i : \alpha]\neg\mathbf{Prev}\varphi$ is also valid. Because $\mathbf{PossIntend}_i(\alpha, \varphi)$ implies $\mathbf{B}_i\neg\varphi$, $M, s \models \mathbf{B}_i[i : \alpha]\neg\mathbf{Prev}\varphi$. But then $M, s \models [i : \alpha]\mathbf{B}_i\neg\mathbf{Prev}\varphi$ and $M', s' \models \mathbf{B}_i\neg\mathbf{Prev}\varphi$, which proves (C). $\mathbf{PossIntend}_i(\alpha, \varphi)$ implies $\mathbf{B}_i\mathbf{G}_i\varphi$, which implies $\mathbf{G}_i\varphi$ (because it was assumed that $\mathbf{B}_i\mathbf{G}_i\varphi \rightarrow \mathbf{G}_i\varphi$).

So because $M, s \models \mathbf{G}_i\varphi$, $M', s' \models \mathbf{Prev} \mathbf{G}_i\varphi$. Furthermore, because it was assumed that $\neg\mathbf{PossIntend}_i(\mathbf{drop}(\varphi), \varphi)$ is valid, $\alpha \neq \mathbf{drop}(\varphi)$. It was assumed that $\neg(\mathbf{Done}(i : \alpha) \wedge \mathbf{Done}(i : \mathbf{drop}(\varphi)))$, so the fact that $M', s' \models \mathbf{Done}(i : \alpha)$ implies $M', s' \models \neg\mathbf{Done}(i : \mathbf{drop}(\varphi))$. But it was also assumed that $\mathbf{Prev} \mathbf{G}_i\varphi \wedge \neg\mathbf{G}_i\varphi \rightarrow \mathbf{Done}(i : \mathbf{drop}(\varphi))$, so it must be that $M', s' \models \mathbf{G}_i\varphi$, which implies $M', s' \models \mathbf{Des}_i(\varphi)$, which proves (D). It is easy to verify that the requirement of unique actions (see constraint (5.6)) validates $\mathbf{Past} \mathbf{Done}(i : \alpha) \rightarrow [i : \alpha]\perp$, i.e., $\langle i : \alpha \rangle \top \rightarrow \neg\mathbf{Past} \mathbf{Done}(i : \alpha)$. Then $\mathbf{B}_i \langle i : \alpha \rangle \top \rightarrow \mathbf{B}_i \neg\mathbf{Past} \mathbf{Done}(i : \alpha)$ is also valid. The antecedent is implied by $\mathbf{PossIntend}_i(\alpha, \varphi)$, so $M, s \models \mathbf{B}_i \neg\mathbf{Past} \mathbf{Done}(i : \alpha)$. But by seriality of R_i , $M, s \models \neg\mathbf{B}_i \mathbf{Past} \mathbf{Done}(i : \alpha)$, which proves (i). $\mathbf{PossIntend}_i(\alpha, \varphi)$ implies $\mathbf{B}_i \neg\varphi$, so by seriality of R_i , $M, s \models \neg\mathbf{B}_i\varphi$, which proves (ii). It is easy to verify that $\mathbf{Past}(\mathbf{Done}(i : \alpha) \wedge \mathbf{New} \varphi) \rightarrow \mathbf{Past} \mathbf{Done}(i : \alpha)$ is a validity. But then $\neg\mathbf{B}_i \mathbf{Past} \mathbf{Done}(i : \alpha) \rightarrow \neg\mathbf{B}_i \mathbf{Past}(\mathbf{Done}(i : \alpha) \wedge \mathbf{New} \varphi)$ is also a validity. So (i) implies (iii), which proves (iii). We can now conclude that indeed $M', s' \models \mathbf{Pride}_i^T(i : \alpha) \wedge \mathbf{Joy}_i^T(\varphi) \wedge \mathbf{Gratification}_i^T(i : \alpha, \varphi)$. Because (M', s') and (M, s) were arbitrary, $\mathbf{PossIntend}_i(\alpha, \varphi) \rightarrow [i : \alpha](\mathbf{Pride}_i^T(i : \alpha) \wedge \mathbf{Joy}_i^T(\varphi) \wedge \mathbf{Gratification}_i^T(i : \alpha, \varphi))$ is valid. \square

Proposition (5.37) and (5.38) The proofs of these propositions are largely the same as the proof of Proposition (5.34) above. For example, assuming $\mathbf{PossAch}_i(\alpha, \psi, \varphi)$ or $\mathbf{PossUnd}_i(\alpha, \psi, \varphi)$ still implies point (ii), because we then have that $M, s \models \mathbf{B}_i \overline{\psi}$, that $\overline{\psi} \rightarrow \neg\psi$ is valid, that R_i is serial, and thus that $M, s \models \neg\mathbf{B}_i\psi$. To prove Proposition (5.38), points (d) and (e) become $M', s' \models \mathbf{Undes}_i(\psi)$ and $M', s' \models \mathbf{Blamew}_i(i : \alpha)$, respectively. $\mathbf{PossUnd}_i(\alpha, \psi, \varphi)$ implies $\mathbf{G}_i\varphi \wedge \overline{\psi} \sqsubseteq \varphi$, which by Proposition (3.50) implies $\mathbf{Undes}_i(\overline{\psi}) = \mathbf{Undes}_i(\psi)$ and then by Proposition (5.23) implies $\mathbf{Blamew}_i(i : \alpha)$. \square

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

- Adam, C., Herzig, A., & Longin, D. (2009). A logical formalization of the OCC theory of emotions. *Synthese*, 168(2), 201–248.
- Breazeal, C. L. (2002). *Designing sociable robots*. Cambridge, MA: MIT Press.
- Cohen, P. R., & Levesque, H. J. (1990). Intention is choice with commitment. *Artificial Intelligence*, 42, 213–261.
- Coppin, G. (2008). Emotion, personality and decision-making: Relying on the observables. In *Proceedings of the third international conference in human centered processes (HCP-2008)*, Delft, The Netherlands.
- Damasio, A. R. (1994). *Descartes' Error: Emotion, Reason and the Human Brain*. New York: Grosset/Putnam.
- Dastani, M., & Meyer, J.-J. Ch. (2006). Programming agents with emotions. In *Proceedings of the 17th European conference on artificial intelligence (ECAI'06)* (pp. 215–219). Riva del Garda, Trentino, Italy.
- Ekman, P., & Davidson, R. J. (Eds.). (1994). *The nature of emotion: Fundamental questions*. Series in affective science. Oxford: Oxford University Press.
- Elster, J. (1996). Rationality and the emotions. *Economic Journal*, 106(438), 1386–1397.

- Emerson, E. A. (1990). Temporal and modal logic. In J. van Leeuwen (Ed.), *Handbook of theoretical computer science* (Chap. 16, pp. 995–1072). Elsevier/MIT Press, Amsterdam/Cambridge, MA.
- Frijda, N. H. (1987). *The emotions*. Studies in emotion and social interaction. Cambridge, MA: Cambridge University Press.
- Frijda, N. H., Ortony, A., Sonnemans, J., & Clore, G. L. (1992). The complexity of intensity: Issues concerning the structure of emotion intensity. In M. S. Clark (Ed.), *Emotion, review of personality and social psychology* (Vol. 13, pp. 60–89). London: SAGE Publications.
- Gordon, R. M. (1987). *The structure of emotions: Investigations in cognitive philosophy*. Cambridge, MA: Cambridge University Press.
- Gratch, J., & Marsella, S. (2004). A domain-independent framework for modeling emotions. *Journal of Cognitive Systems Research*, 5(4), 269–306.
- Gross, J. J., & Thompson, R. A. (2007). Emotion regulation: Conceptual foundations. In J. J. Gross (Ed.), *Handbook of emotion regulation*. New York: Guilford Press.
- Harel, D., Kozen, D., & Tiuryn, J. (2000). *Dynamic logic*. Cambridge, MA: MIT Press.
- Hoffman, R. R., Waggoner, J. E., & Palermo, D. S. (1991). Metaphor and context in the language of emotion. In R. R. Hoffman & D. S. Palermo (Eds.), *Cognition and the symbolic processes: Applied and ecological perspectives* (pp. 163–185). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Johns, M., & Silverman, B. G. (2001). How emotion and personality effect the utility of alternative decisions: A terrorist target selection case study. In *10th Conference on computer generated forces and behavioral representation*. SISO, May 2001.
- Lazarus, R. S. (1994). *Emotion and adaptation*. Oxford: Oxford University Press.
- LeDoux, J. E. (1996). *The Emotional brain: Mysterious underpinnings of emotional life*. New York: Simon & Schuster.
- Marinier, R. P., & Laird, J. E. (2004). Toward a comprehensive computational model of emotions and feelings. In *Proceedings of the international conference on cognitive modeling (ICCM'04)* (pp. 172–177). Pittsburgh, PA.
- Meyer, J.-J. Ch. (2006). Reasoning about emotional agents. *International Journal of Intelligent Systems*, 21(6), 601–619.
- Meyer, J.-J. Ch., van der Hoek, W., & van Linder, B. (1999). A logical approach to the dynamics of commitments. *Artificial Intelligence*, 113, 1–40.
- Oatley, K., & Jenkins, J. M. (1996). *Understanding emotions*. Oxford, UK: Blackwell Publishing.
- Ortony, A. (2009, October). Personal communication.
- Ortony, A., & Clore, G. L. (2009, April–June). Personal communication.
- Ortony, A., Clore, G. L., & Collins, A. (1988). *The cognitive structure of emotions*. Cambridge, UK: Cambridge University Press.
- Picard, R. W. (1997). *Affective computing*. Cambridge, MA: MIT Press.
- Rao, A. S., & Georgeff, M. P. (1998). Decision procedures for BDI logics. *Journal of Logic and Computation*, 8(3), 293–344.
- Scherer, K. R., Schorr, A., & Johnstone, T. (Eds.). (2001). *Appraisal processes in emotion: Theory, methods, research*. Series in affective science. Oxford: Oxford University Press.
- Seegerberg, K. (1995). Belief revision from the point of view of doxastic logic. *Bulletin of the IGPL*, 3, 535–553.
- Slovan, A. (2001). Beyond shallow models of emotion. *Cognitive Processing*, 2(1), 177–198.
- Steunebrink, B. R. (2010) *The Logical structure of emotions*. PhD thesis, Utrecht University, Utrecht, The Netherlands.
- Steunebrink, B. R., Dastani, M., & Meyer, J.-J. Ch. (2007). A logic of emotions for intelligent agents. In *Proceedings of the 22nd conference on artificial intelligence (AAAI'07)*. Vancouver: AAAI Press.
- Steunebrink, B. R., Dastani, M., & Meyer, J.-J. Ch. (2008). A formal model of emotions: Integrating qualitative and quantitative aspects. In G. Mali, C. D. Spyropoulos, N. Fakotakis, & N. Avouris (Eds.), *Proceedings of the 18th European conference on artificial intelligence (ECAI'08)* (pp. 256–260). Greece: IOS Press.
- Steunebrink, B. R., Dastani, M., & Meyer, J.-J. Ch. (2009). A formal model of emotion-based action tendency for intelligent agents. In *Proceedings of the 14th Portuguese conference on artificial intelligence (EPIA'09)*. Berlin: Springer.
- Steunebrink, B. R., Dastani, M., & Meyer, J.-J. Ch. (2010). Emotions to control agent deliberation. In W. van der Hoek, G. A. Kaminka, Y. Lesperance, & M. Luck (Eds.), *Proceedings of the ninth*

- international conference on autonomous agents and multiagent systems (AAMAS 2010)*, Toronto, Canada.
- van der Hoek, W., van Linder, B., & Meyer, J.-J. Ch. (1999). An integrated modal approach to rational agents. In M. J. Wooldridge & A. S. Rao (Eds.), *Foundations of rational agency*, Vol. 14 of *Applied logic series* (pp. 133–168). Dordrecht: Kluwer.
- van Ditmarsch, H., van der Hoek, W., & Kooi, B. (2007). *Dynamic epistemic logic*. Berlin: Springer.