Chapter 6

The Avatar Will See You Now: Emotional and Social Responses to Support from a Virtual Human

Summary

When in emotional distress, people often seek social support. For support to benefit emotional recovery, the listener should stimulate adaptive reappraisals of the event. Unfortunately, however, support is not always available, and when available, not always effective. Furthermore, sharers typically show a strong preference for socio-affective support (comforting, validation), and are less receptive to cognitive support (reappraisal). In the present study, we examined virtual humans' potential as an alternative route through which effective support may be provided. To assess people's willingness to accept cognitive support from a virtual human, we therefore compared the evaluation of socio-affective versus cognitive support. Participants (N = 115) shared two personal emotional experiences with a virtual human, who provided either socio-affective or cognitive support via the Wizard-of-Oz method (i.e., a human-operated avatar). We found that the provision of cognitive support was experienced as equally helpful in coping with the emotional experience as socio-affective support, and led to a similar level of experienced interpersonal closeness and desire to interact with the virtual human again. Furthermore, both support types made sharers feel better and reduced experienced emotional intensity relating to the upsetting event. These findings suggest that interacting with a virtual human may be a way to overcome sharers' resistance to cognitive support, which should be conducive to emotional recovery. Different explanations of these findings are discussed.

This chapter is based on:

Pauw, L. S., Sauter, D. A., Van Kleef, G. A., Gratch, J., & Fischer, A. H. (in prep). The avatar will see you now: Emotional and social responses to support from a virtual human. Manuscript in preparation.

Introduction

When people are in emotional distress, they often have the urge to share their emotions with others, a phenomenon that has been termed social sharing (Rimé, 2009; Rimé et al., 1998). Indeed, social sharing is a frequently employed means of regulating one's emotions: By talking to others, sharers can elicit support that might facilitate their coping (Rimé, 2009). However, sometimes people may not feel comfortable sharing their personal issues with others, for example when it concerns highly stigmatized issues, such as HIV, psychiatric problems, or physical disabilities (Smart & Wegner, 2000). Others lack friends, or their friends may have grown tired of listening to the same issues over and over again (Forest, Kille, Wood, & Holmes, 2014). In these cases, it can be useful to talk to a therapist who is uninvolved in the person's life and who can provide professional help. Unfortunately, however, even if people can overcome the stigma sometimes associated with talking to a therapist and are willing to discuss their personal issues with them, therapists are not always accessible. Therapists can be expensive, sometimes located far away, and often have waiting lists of several months, sometimes even up to years (Hunt & Eisenberg, 2010; Johnson et al., 2007). In light of these problems, both public and scientific interest in building virtual humans is growing (Markoff & Mozur, 2015; Romeo, 2016). Virtual humans are computer-generated characters that respond to words and non-verbal behavior in human-like ways, using emotional expressions, language, and body language (Gratch et al., 2002), and can provide an alternative or complementary online form of support that is available at all times.

In order to effectively reduce people's experience of negative emotions, three conditions need to be met (Burleson & Goldsmith, 1998; Greenberg, 2004; Rimé, 2009): First, those sharing their emotions – henceforth 'sharers' – need to be willing to disclose an emotionally upsetting event. Second, when doing so, they should focus on their thoughts and feelings regarding the emotionally upsetting event, rather than, for example, merely the facts of the situation. And third, the listener, or support provider, should facilitate adaptive ways of reinterpreting (i.e., reappraising) this event. As outlined below, prior work suggests that these first two boundary conditions can be fulfilled when sharing with virtual humans – sometimes even to a greater extent than when sharing among humans. Our study was designed to examine whether virtual humans can also fulfill the third boundary condition: We tested how sharers respond to virtual humans' attempts to help the sharer reappraise the situation, and whether they perceive this to be helpful.

In examining this question, we draw on theories of social sharing, which outline how people can help reduce others' emotional distress. Generally, two primary

types of support have been distinguished: socio-affective and cognitive support (Rimé, 2009). Socio-affective support includes comfort, validation, and understanding, while cognitive support is directed at changing the way the other thinks about the situation by recreating meaning and reappraisal (Rimé, 2009). As such, cognitive support provision constitutes a way to fulfill the third condition for effective sharing: facilitating adaptive ways of reinterpreting the emotional event. Yet, prior research has shown that sharers generally have a strong and consistent preference for socio-affective support (Duprez, Christophe, Rimé, Congard, & Antoine, 2014; Pauw, Sauter, Van Kleef, & Fischer, 2019; Pauw, Sauter, Van Kleef, & Fischer, 2018), which might render them unreceptive to cognitive support (Marigold et al., 2014). In the current study, we investigated virtual humans' potential as an alternative route through which cognitive support may be effectively provided. People may have different expectations of virtual humans when receiving different types of support, for example because they do not ascribe them the mental capacity to experience empathy (Gray et al., 2007). We therefore examined whether sharers give virtual humans more leeway in providing cognitive support, which has been typically shown to be perceived as less desirable and less normative than socio-affective support in interactions between humans. First, we describe previous research illustrating how sharing with virtual humans fulfills the first two conditions of effective social sharing, that is, people are willing to disclose an upsetting situation and their associated thoughts and feelings to virtual humans. Second, we discuss why sharing with a virtual human may be a promising way to fulfill the third condition, that of facilitating reappraisal.

Sharing with a Virtual Human Fosters Willingness to Disclose Emotions

In order for people to be willing to share and discuss sensitive issues, a safe and non-threatening conversational environment is required (Burleson & Goldsmith, 1998; Greenberg, 2004). By creating anonymity and rapport, virtual humans can provide a safe environment in which people feel comfortable sharing their negative emotions. When interacting with another human, sharers may be faced with self-presentational concerns: by sharing their emotions, they may come across as weak, and thereby risk being ridiculed, rejected, or stigmatized (Burleson & Goldsmith, 1998; Derks, Fischer, & Bos, 2008). When talking to a computer, however, these risks are nullified (Caplan & Turner, 2007). Several studies have shown that, compared to those who believed they were interacting with a human-operated computer, participants who believed they were interacting with a computer reported lower fear of self-disclosure, as well as lower impression management, displayed their sadness more intensely, and were rated by observers as more willing to disclose (Gratch, Lucas, King, & Morency, 2014; Lucas et al.,

2014, 2017). These benefits of talking to a virtual human seemed to be driven by sharers feeling anonymous and not judged. For example, Pickard and colleagues (Pickard, Roster, & Chen, 2016) found that people preferred to talk to a virtual human (rather than to a human) when discussing highly sensitive information, but not when discussing information rated low in sensitivity – the reason for their differential preference being that they did not feel judged by a virtual human. In a similar vein, Kang and Gratch (2010) found that especially those who were socially anxious benefitted from talking to a virtual human, again underlining the potential of virtual humans to buffer self-presentational concerns and thereby create a safe environment.

Complementing this line of work, recent research suggests that virtual humans are also able to create a feeling of rapport – another key ingredient contributing to people's willingness to disclose negative emotions (Burgoon & Le Poire, 1999; Burleson & Goldsmith, 1998; Clark & Finkel, 2005). By using non-verbal behaviors such as positive facial expressions, encouraging head nods, and posture changes that are contingent upon the human sharer, virtual humans bring about subjective feelings and behavioral consequences associated with rapport, such as less tension and embarrassment, and greater feelings of closeness and trust (Gratch, Kang, & Wang, 2013; Kang, Gratch, Wang, & Watt, 2008; Wang & Gratch, 2010). Furthermore, by using verbal techniques such as social dialogue, empathic responding, and reciprocal self-disclosure, virtual humans may similarly foster feelings of rapport and trust (e.g., Bickmore & Picard, 2005; Birnbaum et al., 2016b, 2016a; Hoffman, Birnbaum, Vanunu, Sass, & Reis, 2014; Liu & Sundar, 2018). Together, these studies thus suggest that virtual humans are capable of fulfilling the first two conditions for effective social support: People are willing to disclose personally upsetting events and feel comfortable talking about their negative thoughts and feelings with virtual humans.

Can Sharing with a Virtual Human Foster Reappraisal?

A third condition for effective emotional recovery to occur concerns the listener's facilitation of more adaptive appraisals of the situation (i.e., reappraisals; Burleson & Goldsmith, 1998; Rimé, 2009). By helping the sharer take a different perspective on the situation – for example, by encouraging a more positive view on the situation – sharers can be helped to change the way they think about the situation, which changes the emotions that the situation elicits (Dobkin et al., 2004; Nils & Rimé, 2012; Panzarella et al., 2006; Rimé, 2009). Speaking to its potential for bringing about emotional recovery, reappraisal has generally been found to be the most adaptive way of regulating one's own emotions (Webb et al., 2012). Furthermore, reappraisal lies at the heart of cognitive behavioral therapy (CBT), which is successful in treating a wide variety of disorders characterized by

emotion regulation impairments (for a review of meta-analyses, Butler, Chapman, Forman, & Beck, 2006; Hofmann, Asnaani, Vonk, Sawyer, & Fang, 2012). Furthermore, research on social sharing has shown that whereas socio-affective support has been found to increase temporary feelings of relief, perceived understanding, and closeness to the listener, cognitive support is generally more conducive to bringing about long-term emotional recovery (Batenburg & Das, 2014; Brans, Van Mechelen, Rimé, & Verduyn, 2013; Lepore, Fernandez-Berrocal, Ragan, & Ramos, 2004; Morelli, Lee, Arnn, & Zaki, 2015; Nils & Rimé, 2012; Pauw et al., 2018; Rimé, 2009).

While cognitive support provision may thus be a way to fulfill the third boundary condition for effective sharing, prior research has shown that sharers generally have a strong and consistent preference for socio-affective support (Duprez et al., 2014; Pauw et al., 2019, 2018). Furthermore, socio-affective support provision is highly normative (Brans et al., 2013; Pauw, Sauter, Van Kleef, & Fischer, in press). Listeners who provide socio-affective support are also more popular than those who tend to provide cognitive support – an effect that may be explained by sharers appraising cognitive support as a challenge to their existing views of the emotional situation (Niven, Garcia, van der Löwe, Holman, & Mansell, 2015). Similarly, there is evidence suggesting that sharers seek socio-affective support also when sharing with a virtual human. For example, chatbots that displayed sympathy and empathy when giving advice about a sensitive personal issue (e.g., by validating the negativity of the situation) were perceived as more supportive compared to those that provided merely unemotional advice (Liu & Sundar, 2018). Furthermore, a set of studies in which participants interacted with a social robot showed that participants preferred interacting with a robot that reacted responsively (non-verbal signs of attentive listening and verbal validation of participants' feelings) rather than unresponsively (neither non-verbal nor verbal feedback; Birnbaum et al., 2016b, 2016a; Hoffman, Birnbaum, Vanunu, Sass, & Reis, 2014): Participants felt more understood, validated and cared for, engaged in greater non-verbal approach behaviors (e.g., leaning forward, eye contact, smiling), and were more motivated to have the robot as a companion in stressful periods. In sum, these findings show that sharers prefer emotional and empathic support over non-responsive, unemotional support, even from virtual humans.

The question, then, is whether sharers are open to cognitive support when provided by a virtual human, or whether they would, like with human interaction partners, prefer socio-affective support. The findings outlined above indeed suggest that people appreciate socio-affective support from virtual humans. However, people do make important distinctions between humans and virtual humans (see De Melo & Gratch, 2015), and consequently may have different (support) expectations when sharing with a virtual human. Given that people

generally only ascribe the ability to feel and experience to humans but hardly to machines (Gray et al., 2007), people may not expect virtual humans to be highly empathic and to provide support the way they expect from human listeners. This may lead to a less pronounced preference for socio-affective support over cognitive support. Indeed, while not directly studying people's support expectations of virtual humans, there is work suggesting that because people have lower expectations of virtual humans' mental abilities, they expect lower levels of empathy and are more forgiving of 'insensitive' behaviors. For example, people expect robots to be worse at 'emotional' tasks (Waytz & Norton, 2014), are more willing to accept unfair offers from computers than from humans (Sanfey, Rilling, Aronson, Nystrom, & Cohen, 2003), and hold robots less (morally) accountable for making mistakes than humans (e.g., erroneoulsy depriving participants of prize money they had won; Kahn et al., 2012). Furthermore, when believing they were interacting with a computer (compared to a human), people were much more forgiving of the other's low interactivity (i.e., degree of acknowledgment and awareness of participants' prior responses; Go & Sundar, 2019). Thus, given the (perceived) non-judgmental and non-emotional nature of virtual humans, people may be fairly forgiving of less sensitive and normative support, leading them to be relatively accepting of cognitive support.

Some indirect evidence for the potential for virtual humans to provide effective cognitive support and thereby facilitate reappraisals comes from recently developed chatbots and other text-based smartphone interventions employing a wide variety of treatments, including the use of CBT principles (Fitzpatrick, Darcy, & Vierhile, 2017; Hoermann, McCabe, Milne, & Calvo, 2017; Inkster, Sarda, & Subramanian, 2018; Ly, Ly, & Andersson, 2017). For example, several prototypes for treating depression and anxiety have shown moderate success (Fitzpatrick et al., 2017; Fulmer, Joerin, Gentile, Lakerink, & Rauws, 2018; Inkster et al., 2018). However, given that these apps did not focus solely on CBT, but employed mixed methods (e.g., dialectical behavior therapy, motivational interviewing, positive behavior support, behavioral reinforcement, mindfulness), these findings are not conclusive regarding people's receptiveness to reappraisal: It is possible that other aspects of chatbots' responses were driving the positive outcomes. Furthermore, most of these apps were used for treating clinical symptoms. We were interested in examining whether a virtual human might foster effective non-clinical social sharing of the kind one would engage in in daily life with friends, family, or a therapist. And finally, previous studies have employed text-based conversational agents. Our aim was to examine how people experience sharing with a virtual human with whom they engage in an actual face-to-face interaction, thereby being a somewhat closer approximation of what real-life conversations may look like.

The Present Study

In sum, the first two conditions of effective social sharing can be fulfilled when sharing with virtual humans: (1) They create a safe environment for sharers to disclose sensitive information by fostering anonymity while simultaneously establishing rapport, and (2) sharers feel comfortable talking about their feelings and emotions in response to questions about the emotion-eliciting event. The present study aimed to examine how (human) sharers experience virtual humans' attempts to fulfill the third boundary condition: the provision of cognitive support in order to facilitate adaptive reappraisals of the situation. We compared sharers' responses to cognitive support with their responses to socio-affective support. Previous research and theoretical arguments inform two competing hypotheses: On the one hand, we may expect that sharers prefer socio-affective over cognitive support from virtual humans, as has been found when sharing with other humans. On the other hand, there may be no differences in the preference for these two types of support provided by virtual humans, given people's different expectations from virtual humans.

To examine these competing hypotheses, 115 participants engaged in social sharing with a virtual human who provided either socio-affective or cognitive support. Participants had two conversations: One in which they shared a situation that made them feel angry and one in which they shared a situation that made them feel worried. The virtual human provided support through the Wizard-of-Oz paradigm, a frequently used method in computer-human interactions (Dahlbäck, Jönsson, & Ahrenberg, 1993), in which participants are implicitly led to believe they are interacting with a computer, whereas in reality the avatar is controlled by a human experimenter ("wizard"). Rather than employing an automated agent, we used this procedure to ensure that any differences between the conditions would not be due to possible limitations of the dialogue systems. Depending on their assigned experimental condition, participants received either socio-affective or cognitive support across both conversations. Afterwards, we assessed emotional intensity, perceived helpfulness and relational closeness with the avatar. As such, the present study allowed us to examine to what extent people are receptive to cognitive support from a virtual human, by testing whether these two different forms of support have differential emotional and relational consequences.

Method

Participants

A total of 115 participants (M_{age} = 41.0, SD = 12.6; 50% male) were recruited via Craigslist in Los Angeles, United States. 45.2% of the participants identified

themselves as White/Caucasian, 30.4% identified themselves as Black/African American, 13.9% identified themselves as Asian/Pacific Islander, and 11.3% identified themselves as Hispanic American. Five participants had partially missing data because they did not have anything to discuss that they were either worried (2 participants) or angry (3 participants) about. Two participants ended up talking about a different event than what they had initially recalled, such that their pre-sharing measures did not correspond to the same event as their post-sharing measures, thereby also resulting in some missing data for one out of their two conversations. Finally, one participant failed to fill out one item (post-sharing measure of (anger) emotional intensity). The data of these participants could still be analyzed for at least one out of two conversations and they were therefore only excluded from those analyses for which they had missing data.

Procedure

Upon arrival, participants were told that they would share two personal issues with our virtual agent Julie. Depending on their randomly assigned order condition, participants were first instructed to recall either a situation that still worried them or that still made them angry. These two emotion conditions were included in order to create variance in support needs. Previous research has found that different emotions are associated with different support preferences: Whereas anger is linked to a strong preference for socio-effective support, worry has been associated with a desire for both socio-affective and cognitive support (Pauw et al., 2018). Before sharing, participants started the survey on an iPad, and were asked to write down the most important details of the situation, after which they answered questions about the event, including their experienced emotions. Next, they were instructed to call Julie by calling her name. Julie, the avatar, then appeared on a large screen in front of them, looking like a middle-aged woman sitting in a chair similar to a therapy setting (see Figure 1). Julie was controlled in real time by the experimenter.

In the first conversation, a short rapport-building phase was included in which Julie asked several introductory questions (e.g., "Where are you from originally?"; "Do you consider yourself more shy or outgoing?") in order for participants to get acquainted with talking to Julie. This also allows for building rapport, which has been shown to positively affect self-disclosure (Lucas et al., 2014). Then Julie prompted the participant to tell her something about their anger (or worry) evoking experience. Depending on the experimental condition, participants would receive three instances of either socio-affective or cognitive support provision. In order to standardize the conversations across participants, each conversation included three questions related to the emotional experience, and three instances of support provision by Julie. If a participant spoke too little to

allow contingent support provision, they were additionally prompted. After three rounds of support provision, Julie redirected participants back to the survey on their iPad. Participants then answered questions about their experienced emotional intensity, perceived support, perceived benefits and experienced closeness to Julie.

Next, participants went through another round of recall (i.e., an anger or worry evoking situation), a second conversation with Julie, and the same pre- and post-conversational measures. This time, however, after welcoming the participant back to the conversation, Julie directly prompted participants for the second emotional situation, rather than first engaging in small talk as was done in the initial rapport-building phase. At the end of the survey, participants answered some control questions regarding their experience of talking to a virtual human, and several demographic questions. Finally, as a behavioral measure of participants' willingness to interact with Julie further, participants were asked whether, while waiting until the experimenter ensured payment, they would prefer to fill out another survey or to talk to Julie some more. After indicating their preference, they were told that they in fact did not have to wait and could call the experimenter. In total, participation on average took about 25 minutes.[34]

Figure 1. Screenshot of our Virtual Human Julie.

[34] Additionally, before each conversation, we obtained background information regarding the shared event (e.g., recency, previous sharing, appraisals). We also included post-conversational ratings of desired support (see Supplemental Materials), interpersonal alerting and participants' previous experience with talking to a virtual human. These data fall outside the scope of this chapter, but are available upon request.

Design and Materials

This study involved a 2 (Support Type: Socio-Affective vs. Cognitive Support) x 2 (Emotion: Anger vs. Worry) x 2 (Emotion Order) mixed design. Participants were asked to share both an anger and worry evoking situation (within subjects) in randomized order (between subjects) and received one of the two types of support (between subjects).

Emotion recall

Participants were asked to recall both an anger and worry-evoking situation. More specifically, for anger, participants were asked to think of an issue they had with a friend or family member that still made them feel angry, and that had not yet been resolved. For worry, participants were instructed to think of an issue involving their work, finances or health that still made them feel worried, and that had not yet been resolved. We included these instructions in order to facilitate the construction of a repertoire of supportive responses that would fit any stories within these themes. Furthermore, in order to create variance and differentiate anger and worry instances, anger targeted a social situation, whereas worry did not. Participants were asked to take a moment to write down the most important details of the situation in order to recall what it was that made them so angry or worried.

Support provision

Participants received three instances of either socio-affective or cognitive support. These instances were taken from separate repertoires of socio-affective and cognitive support that were established before data collection and tailored to the specific emotion. All responses were recordings of a human voice. Socio-affective support consisted of sentences expressing sympathy (e.g. "I'm sorry to hear that"), and validation of the participants' emotional experience (e.g. "You have every right to be angry with them"), their appraisals (e.g. "That's so unfair") or difficulty of the situation itself (e.g. "It's stressful not knowing how things will turn out"). Cognitive support was always directed at reappraising the situation by either trying to find a more positive way of looking at the situation, or by putting it in perspective. This was done by trying to minimize the situation (e.g. "Maybe you don't need to assume the worst"), pointing the participant to the changeability of their feelings (e.g. "Maybe you won't continue to feel this way") or of the situation (e.g. "Maybe with time they'll come around"), or by providing a silver lining (e.g. "It sounds like you're learning though"). Furthermore, an additional form of cognitive support for anger included the provision of another perspective by minimizing the wrongdoer's bad intentions (e.g. "They probably didn't mean any harm") or foreknowledge (e.g. "Maybe they didn't realize how much it affects you").

Manipulation check: Perceived support
In order to assess whether participants were sensitive to the kind of support they received, after each conversation, they rated their perceived support using a 100-point slider bar (0 = *not at all*, 100 = *very much*). Socio-affective support was measured using four items (e.g. "Julie was empathic"; anger α = .94, worry α = .93). Cognitive support similarly was measured using four items (e.g. "Julie tried to help me look at the situation from a different perspective"; anger α = .93, worry α = .95). All items can be found in Appendix A.

Emotional experience
Participants rated their experience of anger, worry, guilt, shame, fear and sadness on a slider ranging from 0 (*not at all*) to 100 (*very much*) when thinking back at the recalled event and before sharing the emotional situation. This allowed us to test whether participants indeed experienced the target emotion most strongly (i.e., manipulation check of emotion recall). In order to assess the potential effect of support provision on participants' emotional experience, participants rated their experience of the target emotion again after sharing. In addition, before and after each conversation, participants rated their general affect on a scale from -5 (very negative) to +5 (very positive).

Perceived benefits
Perceived benefits were measured using three questions asking the extent to which participants felt that the conversation had made them feel better, helped them deal with the situation better, and reduced their negative emotions regarding the specific event (anger α = .91, worry α = .92). All items were rated on a 7-point Likert scale (1 = *not at all*, 7 = *very much*).

Interpersonal closeness
To assess experienced closeness, after each conversation, participants rated how connected they felt to Julie on a 7-point Likert scale (1 = *not at all*, 7 = *very much*).

Desire to interact with Julie
At the end of the study, participants were told they would need to wait five minutes until the experimenter could finish up the study and arrange payment. They were asked whether they preferred to talk to Julie a bit more, or to fill out another survey. Choosing to talk to Julie was considered a behavioral indicator of a greater desire to interact with Julie further.

Control questions
Participants were asked several control questions regarding their overall experience talking to Julie. They were asked to what extent it bothered them that the conversation had been video recorded (M = 2.20, SD = 1.54); How similar this conversation was to a conversation they normally have in their daily life (M = 3.83,

SD = 1.61); Whether they found it difficult to talk about their personal situations with Julie (*M* = 2.48, *SD* = 1.61); How well they thought Julie did as an conversational partner (*M* = 4.58, *SD* = 1.57). All questions were answered on a 7-point Likert scale (1 = *not at all | not well at all*, 7 = *very much | very well*). Overall, these questions revealed that participants had a relatively positive experience sharing with Julie.[35]

Results

Manipulation Check: Emotion Recall

To verify whether participants indeed experienced more anger when recalling an anger experience, and more worry when recalling a worry-evoking situation, we conducted a 2 (Emotion Condition: Anger vs. Worry) x 6 (Experienced Emotion: Anger, Worry, Sadness, Fear, Guilt and Shame) Repeated Measures ANOVA with both factors varying within subjects. Mauchly's test indicated that the assumption of sphericity had been violated, χ^2 (14) = 60.13, $p < .001$. Therefore, degrees of freedom were corrected using Huynh-Feldt estimates of sphericity (ε = .85). Main effects were observed for Emotion Condition (F[1, 107] = 11.50, p = .001, η_p^2 =.10), as well as Experienced Emotion (F[4.23, 107] = 63.09, $p < .001$, η_p^2 =.37). More importantly, a significant interaction was observed between Emotion Condition and Experienced Emotion (F[4.23, 107] = 67.17, $p < .001$, η_p^2 =.39). All means and standard deviations are displayed in Table 1.

Table 1
Means (M) and Standard Deviations (SD) of Emotional Intensity Ratings of Experienced Emotions split by Emotion Condition (Conversation)

	Emotion Condition (Conversation)	
	Anger	Worry
	M (SD)	*M (SD)*
Experienced Emotions		
Anger	70.87 (28.48)	36.99 (34.10)
Worry	38.95 (32.59)	73.13 (26.05)
Guilt	21.37 (28.98)	30.02 (33.79)
Shame	22.50 (30.16)	27.08 (31.58)
Fear	23.50 (31.35)	58.35 (31.75)
Sadness	51.38 (33.69)	49.38 (33.92)

[35] Ratings on these control questions did not differ as a function of support condition (F[4, 108] = 0.41, p = .802, η_p^2 =.02), order condition (F[4, 108] = 0.34, p = .852, η_p^2 =.01), or their interaction (F[4, 108] = 0.95, p = .437, η_p^2 =.03).

Simple main effect analyses showed a significant effect of Emotion Condition on anger (F[1, 107] = 99.19, p < .001, η_p^2 =.48), indicating that those in the anger condition experienced more intense anger compared to those in the worry condition. Furthermore, Emotion Condition had a significant effect on worry (F[1, 107] = 77.40, p < .001, η_p^2 =.42), indicating that those in the worry condition experienced more intense worry compared to those in the anger condition. Moreover, pairwise comparisons showed that within the anger condition, anger was experienced more intensely compared to all other emotions (all p's < .001). Similarly, within the worry condition, worry was experienced more intensely compared to all other emotions (all p's < .001). In sum, these findings demonstrate that our manipulation of emotion succeeded.

Manipulation Check: Support Provision

In order to test whether participants in the socio-affective support condition indeed felt they had received more socio-affective support, and participants in the cognitive support condition felt they had received more cognitive support, a Repeated Measures ANOVA was conducted with Support Condition (Socio-Affective vs. Cognitive Support) as a between-subjects variable, and Perceived Support (Socio-Affective vs. Cognitive Support) and Emotion Condition (Anger vs. Worry) as within-subjects variables. There was a main effect of Perceived Support (F[1, 108] = 144.32, p < .001, η_p^2 = .57), and Support Condition (F[1, 108] = 8.90, p = .004, η_p^2 = .08). These effects were qualified by a significant interaction between Support Condition and Perceived Support (F[1, 108] = 33.15, p < .001, η_p^2 = .24).

Contrary to our expectations, simple main effect analyses showed that there was no significant effect of Support Condition on perceived socio-affective support (F[1, 108] = 0.11, p = .740, η_p^2 < .01). Thus, regardless of the type of support that was provided by Julie, participants perceived an equally high amount of socio-affective support (see Table 2 for all means and standard deviations). For perceived cognitive support, however, simple main effect analyses showed that there was a significant effect of Support Condition (F[1, 108] = 23.00, p < .001, η_p^2 = .18). In line with our expectations, those in the cognitive support condition perceived a significantly higher amount of cognitive support compared to those in the socio-affective support condition.[36] Finally, there was no significant three-way interaction between Support Type, Support Condition and Emotion Condition (F[1, 108] = 0.53, p = .470, η_p^2 = .01), indicating that our findings did not

[36] There was no main effect of Emotion Condition, F(1, 108) = 0.29, p = .589, η_p^2 < .01. Furthermore, there was no significant interaction between Support Condition and Emotion Condition, F(1, 108) = 0.38, p = .542, η_p^2 < .01.

vary depending on Emotion Condition.[37] In sum, our manipulation was thus partially successful: While participants experienced an equally high degree of socio-affective support across both support conditions, they did experience significantly more cognitive support when Julie provided cognitive support.

Table 2
Means (M) and Standard Deviations (SD) of Perceived Support, Perceived Benefits, Closeness, Emotional Intensity and General Affect split by Support Condition and Emotion Condition.

	Support Condition			
	Socio-Affective Support		Cognitive Support	
	Anger	Worry	Anger	Worry
	M (SD)	*M (SD)*	*M (SD)*	*M (SD)*
Perceived Support				
Socio-Affective Support	68.73 (26.87)	69.09 (26.63)	68.42 (25.34)	72.43 (24.30)
Cognitive Support	32.99 (28.58)	32.38 (29.17)	57.65 (30.13)	57.70 (30.50)
Perceived Benefits	3.57 (1.66)	3.74 (1.75)	4.08 (1.62)	4.19 (1.70)
Closeness to Julie	4.20 (1.75)	3.96 (1.73)	4.29 (1.63)	4.54 (1.68)
Emotional Intensity of Target Emotion				
Before Sharing	70.74 (25.87)	74.06 (24.93)	70.61 (31.18)	73.11 (26.75)
After Sharing	50.43 (26.84)	52.02 (26.17)	38.91 (27.86)	49.57 (30.04)
General Affect				
Before Sharing	0.74 (2.69)	0.41 (2.57)	0.22 (2.45)	0.63 (2.59)
After Sharing	1.67 (2.19)	1.56 (2.26)	1.54 (2.02)	2.17 (1.85)

Effect of Support Condition and Emotion on Perceived Benefits

In order to examine whether the provision of socio-affective versus cognitive support differentially impacted participants' perceived benefits, a 2 (Support Condition: Socio-Affective vs. Cognitive Support) x 2 (Emotion Condition: Anger vs. Worry) Repeated Measures ANOVA was conducted, with Support Condition as a between-subjects variable, Emotion Condition as a within-subjects variable, and perceived benefits as the dependent variable. No main effect of Support Condition was observed (F[1, 108] = 2.75, p = .100, η_p^2 =.03), indicating that the provision of socio-affective and cognitive support similarly impacted perceived benefits. Furthermore, there was no effect of Emotion Condition (F[1, 108] = 1.08, p = .300, η_p^2 =.01), nor an interaction effect between Support Type and Emotion Condition, (F[1, 108] = 0.06, p = .808, η_p^2 < .01). All means and standard deviations are displayed in Table 2.

[37] Including Emotion Order in our models revealed no statistically significant interactions with Emotion Order, indicating that all observed effects remain the same regardless of whether participants first shared an anger-evoking event and then a worry-evoking event, or vice versa. Therefore, we present all our analyses without Emotion Order in the model.

Effect of Support Condition and Emotion on Interpersonal Closeness

Similarly, in order to examine whether the provision of socio-affective versus cognitive support differentially impacted experienced closeness towards Julie, a 2 (Support Condition: Socio-Affective vs. Cognitive Support) x 2 (Emotion Condition: Anger vs. Worry) Repeated Measures ANOVA was conducted. Again, no main effect of Support Condition was observed (F[1, 108] = 1.26, p = .264, η_p^2 = .01). Furthermore, there was no effect of Emotion Condition (F[1, 108] < 0.01, p = .974, η_p^2 < .01), nor an interaction effect between Support Type and Emotion Condition predicting interpersonal closeness (F[1, 108] = 3.08, p = .082, η_p^2 = .03). Thus, these findings indicate that the provision of socio-affective and cognitive support similarly impacted how close participants felt to Julie after the interaction.

Effect of Support Condition and Emotion on Desire to Interact with Julie

To test whether support provision differentially impacted participants' desire to continue interacting with Julie, we conducted a Chi-Square test examining the relation between Support Type (Socio-Affective vs. Cognitive Support) and Willingness to Interact (Yes vs. No). Participants were equally willing to continue interacting with Julie regardless of whether she had provided socio-affective support (34.5% of participants chose to continue interacting) or cognitive support (43.1% of participants chose to continue interacting), $\chi^2(1) = 0.87$, p = .351, Phi = 0.09.

Effect of Support Condition on Emotional Intensity

To examine whether the provision of socio-affective and cognitive support was effective in reducing the experienced negative emotion, we tested its effects on the intensity of the target emotion as well as general affect. First, a 2 (Time: Before and After Sharing) x 2 (Support Condition: Socio-Affective vs. Cognitive Support) x 2 (Emotion Condition: Anger vs. Worry) Repeated Measures ANOVA was conducted, with Time and Emotion Condition as within-subjects variables, Support Condition as between-subjects variable, and intensity of the target emotion (worry or anger) as the dependent measure. There was a main effect of Time (F[1, 105] = 127.43, p < .001, η_p^2 = .55), indicating that overall, participants showed a reduction in the intensity of the target emotion after sharing with Julie (see Table 2 for all means and standard deviations). The absence of an interaction effect between Time and Support Condition (F[1, 105] = 2.23, p = .139, η_p^2 = .02) indicates that this reduction in intensity was equally strong across socio-affective and cognitive support provision. Furthermore, no three-way interaction was observed between Time, Support Type and Emotion condition (F[1, 105] = 1.80, p = .183, η_p^2 = .02). Finally, no other effects emerged (ps > .315), except for a small

effect of Emotion Condition (F[1, 105] = 4.88, p = .029, η_p^2 = .04), indicating that on average, participants experienced somewhat higher emotional intensity when discussing the worry-evoking event as compared to the anger-evoking event. In sum, regardless of the type of support that was provided, all participants indicated feeling less angry and worried about the events after sharing.

Second, to assess whether socio-affective and cognitive support made participants generally feel better after sharing, we conducted a 2 (Time: Before and After Sharing) x 2 (Support Condition: Socio-Affective vs. Cognitive Support) x 2 (Emotion Condition: Anger vs. Worry) Repeated Measures ANOVA, predicting general affect. Here too, there was a main effect of Time (F[1, 106] = 38.70, p < .001, η_p^2 = .27), indicating that overall, participants generally felt better after sharing with Julie (see Table 2 for all means and standard deviations). The absence of an interaction effect between Time and Support Condition (F[1, 106] = 0.97, p = .328, η_p^2 = .01) indicated that the observed increase in positive affect was equally strong across socio-affective and cognitive support provision. There was an interaction between Support Type and Emotion Condition (F[1, 106] = 4.46, p = .037, η_p^2 = .04). Simple slope analyses, however, showed that the provision of socio-affective versus cognitive support did not affect general affect differently across the two emotion conditions (all p's > .273). No three-way interaction was observed between Time, Support Type and Emotion condition (F[1, 106] < 0.01, p = 1.000, η_p^2 < .01). In sum, regardless of the type of support that was provided, all participants indicated that they felt better after sharing.

Exploratory Analyses: The Effects of Perceived Support

Finally, we examined whether support *perceived* by the participants predicted positive relational and emotional outcomes. To that end, we conducted multiple regression analyses of the anger and worry conversations separately, with perceived socio-affective and cognitive support as predictors of perceived benefits and closeness. As can be seen in Table 3, both for worry and anger, perceived socio-affective and cognitive support were significant positive predictors of perceived benefits, as well as experienced closeness. Thus, overall, the more support participants felt they had received, the more they perceived the conversation as having been helpful in coping with the event, and the closer they felt to Julie. However, a logistic regression analysis including perceived socio-affective and cognitive support for both the worry and anger conversation as predictors did not significantly predict the desire to talk to Julie (χ^2[4]= 3.98, p = .408), indicating that the degree of support that they felt they had received did not affect their willingness to have another conversation with Julie (see Table 4 for separate coefficients).

Table 3
Outcomes of Four Regression Analyses Predicting Perceived Benefits (or Closeness) by Perceived Socio-Affective and Cognitive Support, for Anger and Worry Separately.

	Perceived Benefits		Closeness	
	β	*p*	β	*p*
Anger Conversation				
Perceived Socio-Affective Support	.25	.003	.49	< .001
Perceived Cognitive Support	.54	< .001	.36	< .001
Worry Conversation				
Perceived Socio-Affective Support	.44	< .001	.40	< .001
Perceived Cognitive Support	.34	< .001	.43	< .001

Table 4
Outcomes of a Logistic Regression Analysis Predicting Desire to Talk to Julie (0 = No, 1 = Yes) by Perceived Socio-Affective and Cognitive Support during the Anger and Worry Conversations.

	Desire to Talk to Julie	
	B	*p*
Perceived Socio-Affective Support (Anger)	0.02	.259
Perceived Cognitive Support (Anger)	< -0.01	.593
Perceived Socio-Affective Support (Worry)	< -0.01	.918
Perceived Cognitive Support (Worry)	0.01	.427

Furthermore, we examined whether perceived socio-affective and cognitive support would also predict improved emotional outcomes. To this end, we controlled for emotional intensity reported prior to sharing when predicting post-sharing emotional intensity, and we controlled for general affect reported prior to sharing when predicting post-sharing general affect (see Table 5). These findings showed that perceived support did not predict a reduction in emotional intensity experienced in response to the anger or worry-evoking event. In fact, perceived cognitive support predicted an increase in worry experienced after sharing. Perceived socio-affective and cognitive support did, however, predict enhanced general affect in both the worry and anger conversation: The more support participants felt they had received, the better they reported feeling afterwards.

Table 5

Outcomes of Four Regression Analyses Predicting Post-Sharing Emotional Intensity (or Post-Sharing General Affect) by Perceived Socio-Affective, Perceived Cognitive Support and Pre-Sharing Emotional Intensity (or Pre-Sharing General Affect), for Anger and Worry Separately.

	Post-Sharing			
	Emotional Intensity		General Affect	
	β	*p*	β	*p*
Anger Conversation				
Perceived Socio-Affective Support	.15	.183	.30	.001
Perceived Cognitive Support	.03	.807	.29	.001
Pre-Sharing Emotional Outcome	.31	.001	.43	< .001
Worry Conversation				
Perceived Socio-Affective Support	-.04	.683	.40	< .001
Perceived Cognitive Support	.22	.036	.20	.024
Pre-Sharing Emotional Outcome	.51	< .001	.39	< .001

Note. 'Pre-sharing emotional outcome' denotes pre-sharing *emotional intensity* experienced in response to the specific event in the models predicting post-sharing emotional intensity, but denotes pre-sharing *general affect* in the models predicting post-sharing general affect.

Discussion

Main Findings and Theoretical Implications

In light of the many obstacles people may face when seeking support in daily life, the present study was set up to examine virtual humans' potential for providing effective social support. To this end, participants shared two personal emotional events with a virtual human that provided either cognitive support (facilitating reappraisal) or socio-affective support (providing comfort and validation). Our findings showed that the provision of cognitive support was experienced as equally helpful in coping with the emotional experience as compared to socio-affective support, and led to a similar level of experienced interpersonal closeness and desire to interact with the virtual human again. Furthermore, both socio-affective and cognitive support were (similarly) effective in reducing experienced emotional intensity relating to the upsetting event and led participants to feel better in general. Exploratory analyses revealed that the more socio-affective and cognitive support participants felt to have received, the more helpful they experienced the conversations to be, the better they felt afterwards, and the closer they felt to the virtual human. Perceived support did not reduce the emotional intensity experienced in response to the discussed upsetting event, however. Together, these findings do not support the hypothesis that sharers respond more favorably to socio-affective support than cognitive support provided by a virtual human, as they do when sharing with humans.

Instead, they speak to a less pronounced preference for socio-affective support when sharing with virtual humans, thereby lending indirect support for our alternative hypothesis.

The finding that cognitive support fostered similar levels of perceived benefits, relational closeness and emotional improvement as socio-affective support suggests a certain degree of receptivity to cognitive support. These observations are promising in light of prior research on social sharing between humans, which has revealed a strong preference for socio-affective support (Duprez et al., 2014; Pauw, Sauter, Van Kleef, & Fischer, in press; Pauw et al., 2018). This preference for socio-affective support may be due to sharers' need for perceived understanding (Reis, Lemay, & Finkenauer, 2017): Socio-affective support validates the sharer's thoughts and feelings, which signals understanding and communicates that it is accepted and appropriate to feel the way they feel (Shenk & Fruzzetti, 2011). Cognitive support, on the other hand, may be experienced as invalidating, as the listener tries to encourage *another* way of interpreting the event, which may lead the sharer to refuse the listener's support attempt (see Marigold, Holmes, Wood, & Cavallo, 2014). Interestingly, in the current study, the provision of cognitive support did not come at the expense of perceiving the virtual human as sensitive: Those receiving cognitive support experienced the virtual human as equally empathic as did those who received socio-affective support, and also felt equally understood. Furthermore, greater *perceived* cognitive support generally enhanced experienced benefits and closeness. It is thus possible that virtual humans may provide a way to circumvent the typical trade-off between empathetic and effective support provision.

One potential reason for greater receptivity to cognitive support from virtual humans concerns the different expectations that people may have about support from virtual humans as compared to real humans. The fulfillment of support expectations might be more important in shaping sharers' responses than the mere support itself (Bar-Kalifa & Rafaeli, 2015; Sarason, Pierce, & Sarason, 1990). Given that people generally do not ascribe actual mental experiences to agents, and assume that agents are worse at emotional tasks than humans (Gray et al., 2007; Waytz & Norton, 2014), they might be more forgiving when provided with relatively insensitive or more direct ways of communication as compared to if a human would violate such norms. Consequently, they may adopt a more tolerant mindset, leading them to perceive cognitive support as less judgmental (cf. Pickard et al., 2016), thereby potentially being more inclined to try to reappraise the situation. Furthermore, the provision of cognitive support – though not explicitly sympathetic in its verbal content – might have still communicated emotional involvement: By responding to the participant's story, these support instances may reflect that the virtual human was attending to and interpreting the

participant's emotional experience, and trying to help them feel better (Bodie & Jones, 2012; Burleson, 2008; Burleson et al., 2005; Jones, 2011). Finally, the inclusion of a rapport-building phase may have contributed to the overall high levels of perceived empathy, and thereby to sharers' receptivity to cognitive support. In this rapport-building phase, participants disclosed some personal information, to which the virtual human responded encouragingly. The mere act of talking about personal events with someone who listens and asks follow-up questions might have been experienced as considerate. Indeed, work on listening has shown that active listening, by for example asking questions, is a key component of supportive communication (Bodie & Jones, 2012; Jones, 2011).

Together, our findings thus speak to the potential of virtual humans in fostering adaptive coping. Prior research has shown that virtual humans are capable of fulfilling the first two conditions of effective support (Burleson & Goldsmith, 1998): By creating anonymity and rapport they establish a safe and non-threatening environment in which people feel comfortable disclosing their thoughts and feelings regarding personally upsetting events (Birnbaum et al., 2016b; Liu & Sundar, 2018; Lucas et al., 2014, 2017). The present study suggests that virtual humans can also fulfill the third condition: the facilitation of adaptive reappraisals of the situation (Burleson & Goldsmith, 1998; Rimé, 2009). Participants were equally receptive to cognitive support as to socio-affective support. This is of particular importance given that the facilitation of adaptive reappraisals is key to emotional recovery when sharing in daily life as well as in clinical settings (Butler et al., 2006; Lepore, Fernandez-Berrocal, Ragan, & Ramos, 2004; Nils & Rimé, 2012; Rimé, 2009). The fact that the provision of cognitive support did not come at the cost of perceiving the virtual human as sensitive is promising, given that people prefer to interact with more empathic virtual humans and are more motivated to have them as a companion in difficult times (Birnbaum et al., 2016b; Liu & Sundar, 2018). Taken together, the current findings suggest that people may be receptive of virtual humans providing cognitive support, thereby helping them to adopt a more positive view on the situation.

Limitations and Future Directions

While pointing to the promise of virtual humans as an effective additional source of cognitive support, the present study has several limitations. The finding that sharers showed similar responses to cognitive support provision as to socio-affective support provision – a typically highly appreciated form of support – speaks to sharers' general receptivity to cognitive support. Admittedly, the absence of a significant effect does not equate evidence for the absence of an effect. Yet the finding that greater perceived cognitive support predicted greater perceived benefits and experienced closeness to the virtual human does hint at

sharers' appreciation of cognitive support. Furthermore, both the provision of socio-affective and cognitive support were associated with a reduction in emotional intensity after sharing. However, we did not include a control condition in the present study, thereby making it impossible to attribute emotional improvement to the obtained support rather than venting or the mere passage of time. Future research is thus warranted to more conclusively map sharers' preferences for and responses to support from virtual humans. Ideally, these studies would include a comparison between human and virtual human sharing partners, for example by letting participants believe that the virtual human they are interacting with is operated by a human versus automation. Alternatively, a control condition could be included in which participants do not share or share in the absence of a sharing partner (e.g., by writing, or recording a voice or video message), in order to examine the added value of socio-affective and cognitive support for fostering (perceived) benefits of sharing.

Another shortcoming is that our findings only speak to self-reported short-term benefits, but not to actual (long-term) emotional recovery. In addition to sharers' experienced support, we assessed their *perceptions* of helpfulness, as well the immediate reduction in emotional intensity. It is thus possible that our findings merely reflect *perceived* benefits as typically reported after sharing (see Zech & Rimé, 2005). The finding that perceived support predicted perceived benefits, as well as improved wellbeing, but did not reduce emotional intensity experienced in response to the discussed event is in line with that possibility. Nevertheless, the finding that sharers experienced cognitive support as beneficial is promising, as it reflects receptivity to a form of support that enhances long-term emotional recovery. Yet it remains an open question whether they actually adopted the reappraisals that they were pointed towards. Future research is warranted to examine more long-term consequences of support provision by virtual humans, ideally assessing emotional and relational outcomes across a wider domain (e.g., appraisals, emotions, physical wellbeing, continued use of the virtual human).

Furthermore, an important question is the extent to which the current findings can be applied. The present study employed the Wizard-of-Oz paradigm (Dahlbäck et al., 1993), with a human selecting the most appropriate response from a support repertoire to allow us to match the support to the specific situation at hand. Currently, this technique outperforms agents in terms of conversational abilities, because of humans' emotional intelligence (Grosz, 2018; Mavridis, 2015). Our conclusions regarding the appreciation of support from a virtual human are thus in fact limited to support indirectly provided by a human. The next step would be to automatize support provision. While still in its infancy, we do believe there is potential for doing so: With technology rapidly improving, artificially intelligent virtual humans are being built that are capable of controlling computer-

generated bodies and interacting verbally and non-verbally with other users in virtual environments (see Rizzo et al., 2011). For example, SimSensei Kiosk is a virtual human interviewer that is designed to create engaging face-to-face interactions in which people feel comfortable sharing their psychological distress (Devault et al., 2014). She is used for semi-structured interviewing, in which she engages in follow-up questions and simple empathic responses to sharers' expressions, as well as provides non-verbal behavior that is contingent upon the sharer. An experimental study comparing the automated SimSensei Kiosk to a human-operated (Wizard-of-Oz) version and an actual human interviewer showed that participants were equally willing and comfortable sharing personal information to the automated version as to the human-operated version, and were generally satisfied and willing to recommend it others. While they did perceive the human-operated version as a better listener and experienced greater closeness to her, surprisingly, participants experienced similar levels of rapport to the (automated) agent as to a human interviewer. Thus, while not living up to human standards, the findings are promising. However, as SimSensei Kiosk currently is only capable of generic empathic responding (e.g. "I'm sorry to hear"), we opted for Wizard-of-Oz paradigm, in which a human selected supportive responses that were matched to what the sharer was saying.

Given that several chatbot agents are currently being relatively effectively used to treat a variety of depression and anxiety symptoms (Fitzpatrick et al., 2017; Hoermann et al., 2017; Inkster et al., 2018; Ly et al., 2017), a related question concerns whether virtual humans have added value compared to chatbots when it comes to support provision. One advantage may be that by talking, people emit clearer expressions of what they are feeling compared to when they are writing, given that their speech includes both verbal and non-verbal content. Certain virtual human agents (such as SimSensei Kiosk) are capable of inferring psychological distress (e.g., depression, anxiety and post-traumatic stress disorder symptoms) from people's non-verbal behaviors (e.g., speech prosody, downward eye gaze; Devault et al., 2014; Lucas, Gratch, Scherer, Boberg, & Stratou, 2015; Ring, Bickmore, & Pedrelli, 2016; Scherer et al., 2014; Scherer, Stratou, Gratch, & Morency, 2013), thereby not requiring sharers to explicitly indicate how they feel. Chatbots often demand users to indicate their emotional experience by Likert scales, emojis or other pre-formatted options (e.g., Fitzpatrick et al., 2017; Inkster, 2018), thereby impeding the natural flow of a conversation, and relying on self-report, rather than more implicit indicators of psychological distress. Furthermore, by their non-verbal active listening behaviors, virtual humans may feel more real and human-like than chatbots. This may foster greater rapport and thereby greater self-disclosure and motivation to continue using the virtual human (cf. Birnbaum et al., 2016b; Gratch et al., 2013; Lucas et al., 2014). Finally,

the question is whether writing versus talking about one's emotions is differentially effective in and of itself. People do appear to use different words when describing their experiences depending on whether they are writing or talking (Balon & Rimé, 2016), but we are unaware of any empirical studies directly comparing the effectiveness of chatbots and virtual humans in fostering emotional recovery. In sum, then, future research is warranted to examine whether virtual humans constitute better conversational partners than chatbots.

Conclusions

Taken together, the present study suggests that virtual humans may be capable of fulfilling all necessary steps for providing effective social support. Extending prior research showing that people are willing to disclose personal information and talk about their emotions with a virtual human (Devault et al., 2014; Gratch et al., 2014; Lucas et al., 2017), we show that people also are receptive to cognitive support – a form of support that is considered necessary for effective long-term recovery. The provision of cognitive and socio-affective support yielded similar levels of perceived helpfulness and experienced closeness to a virtual human, and both effectively reduced short-term emotional distress. Hence, while we do not conceive virtual humans as a replacement for humans, the current findings suggest that virtual humans may be a valuable addition for those who may, at times, lack appropriate support from close others or clinicians.

Appendix A: Perceived Support Measure

In order to assess the support participants perceived the avatar Julie to have provided, they were asked the following question: "Looking back at the conversation with Julie about the anger evoking situation, to what extent do you agree with the following statements? Julie..", followed by the eight different types of support:

Socio-affective support:

(1) Showed compassion
(2) Conveyed understanding
(3) Was supportive
(4) Was empathic

Cognitive support:

(5) Tried to help me look at the situation from a different perspective
(6) Tried to help me find meaning in what occurred
(7) Provided an outside perspective
(8) Put what occurred into perspective