

Awe Versus Aww: The Effectiveness of Two Kinds of Positive Emotional Stimulation on Stress Reduction for Online Content Moderators

CHRISTINE L. COOK, National Chengchi University, Taiwan

JIE CAI, Penn State University, USA

DONGHEE YVETTE WOHN, New Jersey Institute of Technology, USA

When people have the freedom to create and post content on the internet, particularly anonymously, they do not always respect the rules and regulations of the websites on which they post, leaving other unsuspecting users vulnerable to sexism, racism, threats, and other unacceptable content in their daily cyberspace diet. However, content moderators witness the worst of humanity on a daily basis in place of the average netizen. This takes its toll on moderators, causing stress, fatigue, and emotional distress akin to the symptomology of post-traumatic stress disorder (PTSD). The goal of the present study was to explore whether adding positive stimuli to breaktimes—images of baby animals or beautiful, aweinspiring landscapes—could help reduce the negative side-effects of being a content moderator. To test this, we had over 300 experienced content moderators read and decide whether 200 fake text-based social media posts were acceptable or not for public consumption. Although we set out to test positive emotional stimulation, however, we actually found that it is the cumulative nature of the negative emotions that likely negates most of the effects of the intervention: the longer the person had practiced content moderation, the stronger their negative experience. Connections to compassion fatigue and how best to spend work breaks as a content moderator are discussed.

CCS Concepts: • **Human-centered computing** → **Empirical studies in collaborative and social computing**; *Empirical studies in HCI*.

Additional Key Words and Phrases: content moderation; stress; compassion fatigue; emotional distress; social networking sites

ACM Reference Format:

Christine L. Cook, Jie Cai, and Donghee Yvette Woh. 2022. Awe Versus Aww: The Effectiveness of Two Kinds of Positive Emotional Stimulation on Stress Reduction for Online Content Moderators. *Proc. ACM Hum.-Comput. Interact.* 6, CSCW2, Article 277 (November 2022), 19 pages. <https://doi.org/10.1145/3555168>

277

1 INTRODUCTION

In the toxic world of cyberspace, online content moderators are the protectors of netizens everywhere. As either an employee or a volunteer, content moderators sift through everything people have posted or have reported as being inappropriate for hours on end [81], seeing the worst content people on the internet have to offer: sexism, racism, and more, with varying degrees of intensity [97]. Much like police officers working on internet crime teams (see [48]), consistent exposure to this kind of content can lead to severe levels of stress and fatigue among content moderators [20].

Authors' addresses: Christine L. Cook, National Chengchi University, Taipei City, Taiwan, christinelcook@outlook.com; Jie Cai, Penn State University, University Park, USA, jie.cai@psu.edu; Donghee Yvette Woh, New Jersey Institute of Technology, Newark, USA, yvettewohn@gmail.com.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2022 Association for Computing Machinery.

2573-0142/2022/11-ART277 \$15.00

<https://doi.org/10.1145/3555168>

Although this is problematic in and of itself, the most unsettling part of this issue is the fact that the problem is only poised to grow over time. Content moderation is a growing field that upholds the social media infrastructure of today, meaning that more and more people will be employed over time to do this job and be subjected to this kind of suffering (see [14, 66]). As we continue to produce more content, more people will have to filter through it and endure the ill-effects thereof. Content moderators are rapidly becoming a major vulnerable group in cyberspace.

However, researchers have begun to take interest in the wellbeing of content moderators [61, 75, 89] and have examined the effectiveness of current methods to ease the burden on content moderators (e.g., [13, 15, 79, 80]), referring to both moderators in volunteer and commercial content moderation, as they both experience exposure to harmful content and suffer from emotional labors [89]. Researchers have tried to reduce the burden on content moderators by altering the offensive content to reduce exposure, often by blurring image and video content [20, 22], and by either supporting or replacing human moderators with artificial intelligence systems, to varying degrees of success [76]. However, psychological studies would suggest that we still need interventions that target the moderator themselves, helping them to cope with the emotions and distress that accompanies their work (e.g., [8]). The present study takes inspiration from the very internet community the moderators aim to protect; one popular way to reduce stress and increase mental wellbeing on the internet is to actively seek out positive content to consume (e.g., [5, 23]). Two of the most purportedly effective sources of this positivity are baby animal images [5] and beautiful landscapes [23]. Thus, in this study, we aim to test the two associated positive emotions—cuteness and awe—for their therapeutic effects on stress and fatigue in the content moderating context to see if they can work for more acutely distressed emotional laborers the way they appear to for the average netizen.

Both cuteness and awe have been shown to be connected to increased levels of prosociality (e.g., [32, 84]) and a sense of connectedness to others (e.g., [6, 82]), but they also have their own distinct benefits that could impact content moderators differently. For instance, cuteness has been repeatedly connected to carefulness, attentional control, and caretaking behaviours [30, 44, 63]. This could impact content moderation in several different ways. The extra carefulness could be linked to a reduction in speed, but an increase in accuracy, while increased attention control has the potential to increase productivity by reducing distraction. From an emotional perspective, cuteness has been linked to stress reduction [94] and, when paired over time with another stimulus, has also been shown to increase the quality of positive experience with paired stimuli [57]. For content moderators, adding cute stimuli to break up the negative stimuli thus has the potential to alleviate the distress associated with their work by reducing stress and its physiological correlates (i.e., high blood pressure and heart rate in the moment), and longitudinally, by improving the moderator's experience with moderation via association. Awe has also been repeatedly linked to stress reduction specifically (e.g., [50]), but is also linked to more negative affect than cuteness [50, 71]. This increased negative affect could potentially heighten moderators' perception of the negativity in the posts they need to moderate, and this could result in multiple differing outcomes (e.g., counteracting the initial positive affect brought about by the awe stimulus, increasing speed to remove the post faster and reduce the negative affect, etc.). Both awe and cuteness however, have the potential to contribute to an upward spiral in positive affect through repeated exposure (e.g., [10, 18, 34]). Given that these upward spirals could have the power to counterract the negative spirals created by content moderation activities (see [27]), both awe and cuteness are worth testing as low-cost, simple interventions that target the moderator as opposed to focusing on the content that needs moderating.

To test these interventions for their suitability for use among content moderators, we ran an experiment on Amazon's Mechanical Turk (MTurk), a crowdsourced task completion website in

which people around the world sign up to participate in small task completions, both academic and commercial, in exchange for payment. This is also a popular platform where a lot of freelance content moderators perform moderation tasks. Participants who are screened for prior content moderation experience were randomly assigned to one of three conditions: taking a break (control), taking a break with pictures of baby animals (cuteness), and taking a break with pictures of beautiful landscapes (awe). Through a simulated content moderation task, in which participants will play the role of a content moderator for a fictional social media platform, we plan to evaluate how effective providing additional positive emotional stimuli is when it comes to reducing a) stress, b) fatigue, and c) emotional distress during breaks in a 45-minute simulated content moderation workday.

Table 1. Summary of Extant Literature's Potential Positive and Negative Effects of Using Awe and Cuteness Induction as Content Moderation Interventions

	Positives	Negatives
Awe	Stress reduction	Potential simultaneous increase in negative affect(self-minimizing perspective).
	Emotional distress reduction	Increased fatigue over time
	Increased positive affect	
	Increased prosocial tendencies	
Cuteness	Stress reduction	Reduction in task completion speed
	Increased positive affect	
	Increased desire to protect	
	Increased desire to nurture	
	Increased task attention	
	Increased care taken in tasks	

2 THEORETICAL BACKGROUND

Content moderation literature is clear about the high degree of stress, fatigue, and emotional distress content moderators experience through their work, with many researchers reporting post-traumatic stress disorder (PTSD) among their content moderator participants [7, 20, 29, 74, 97]. In fact, the job is so stressful that it shares common symptomatology with law enforcement workers tasked with finding and capturing child pornographers [48]. This is likely due to the fact that both are regularly faced with similar images: namely, images of suffering victims, descriptions of said suffering, and threats of more suffering to come (see [48, 66]). Given the fact that new content is posted on social media at an astronomical rate (see [38]) Internet Live Stats, 2020 for examples), content moderators' symptoms are apt to be present for the full duration of a 40-hour work-week, and potentially constantly present, in the event that major disorders like PTSD set in.

To help combat these effects in a low-cost, simple way, we have decided to test the power of awe and cuteness—two positive emotional stimuli—to see if their power to reduce the negative side-effects associated with repeated exposure to negative images [48, 66]. Recent research suggests that positive emotional stimuli may have the potential to reduce the negative effects of the job if moderators are exposed at regular intervals [41, 47, 50, 95]. In fact, we can already see a current example of cuteness being used to improve the experience of people working for long hours at computers in Google's Colab ¹. Programmers can have animated corgis and kittens strut across

¹<https://colab.research.google.com/>

the top of their screens and meow or bark while they work to add some light entertainment to their user interface. If effective at reducing stress and increasing positive emotions in users, this presents a simple, automated intervention that can be used to supplement existing traditional stress relief techniques, such as mindfulness-based stress reduction (MBSR) meditations (e.g., [85, 86]). In addition, while MBSR seems to take a minimum of 10 minutes of continuous practice to achieve effects [37, 88], no such time restrictions have been, to our knowledge, reported for the positive benefits of awe or cuteness, meaning they could theoretically be deployed throughout the workday with minimal disruption to existing schedules. However, despite this potential, to date, no one has actively tested the effects of awe or cuteness for their effectiveness at reducing negative effects and increasing positive affect—the goal of the present work.

2.1 Awe

Awe, and particularly awe experienced as a result of exposure to natural wonders or landscapes, has been repeatedly linked to stress relief, be it pathological (e.g., [3]), or due to simple, everyday stressors (e.g., [41, 47]). This exposure to awesome nature can take the form of actual, *in vivo* exposure (e.g., [3]), or exposure via images or video (e.g., [47]), as these have both been shown to be effective.

The purported mechanism behind awe's stress reduction properties lies in awe's effect on our self-perception and perception of our stressors [41, 50]. When a person witnesses something that inspires awe, the tendency is for the person to feel small in comparison to whatever it is they are seeing [6]. According to Le and colleagues [50], this can lead to one of two outcomes, depending on the person's perspective: the awe makes the person's obstacles or challenges feel small and insignificant, or it makes them feel small and insignificant themselves, and unable to meet the challenges ahead. Other research, though, suggests that the perspective adopted may be influenced by the type of awe experienced: positive or negative [31, 52]. Awe elicited by a rainbow or a waterfall, for example, does not pose any direct threat to the wellbeing of the average person. Awe elicited from a tornado though, does pose a threat, one that we can feel helpless to face [52]. Though all of these stimuli inspire awe, they effect affect and well-being differently [31]. Thus, extant literature would suggest that exposing stressed people to specifically non-threatening, but awe-inspiring natural stimuli should have stress reduction effects on the person. We thus posit the following hypothesis:

- **H1)** Exposure to awe-inducing stimuli will reduce stress in content moderators compared to the stress levels of content moderators who are not awe-induced.

However, experiencing awe may come with a cost in terms of fatigue. This is because the experience of positive awe has been consistently linked to prosociality in study participants [6, 21, 32, 51, 68, 69, 78]. With fundamentally prosocial work comes what various researchers have dubbed “compassion fatigue” [60, 65, 92]. Also called “empathic distress,” compassion fatigue is a cumulative fatigue that care workers experience as a consequence of their employment, as they empathize over time with those under their care [60]. Traditionally it is associated with the medical profession (e.g., [46, 60, 64, 65, 92]), but any job that involves caring for others could theoretically lead to it (e.g., [56, 91]). This is important to note, as the role of a content moderator is a fundamentally prosocial one: the purpose of the job is to protect the average netizen from the stimuli they themselves are forced to witness. This also means that the more content moderation experience one has, potentially, the more cumulative compassion fatigue they are experiencing, creating worsening symptoms over time. Research in awe has established that seeing an awe-inspiring landscape increases the motivation to perform prosocial behavior; the following question, however, remains unanswered:

- **RQ1)** Which has a more powerful effect on fatigue: an awe-inspiring stimulus, or accumulated content moderation experience?

When it comes to emotional distress, however, awe has a much clearer effect [47, 50, 71]. In Koh and colleagues' [47] study comparing happiness and awe's effects on emotional distress, operationalized as negative affect, and their findings indicate that awe has the same distress-reduction properties as happiness. In other words, experiencing awe makes a person feel less negative affect. When distress was operationalized as anxiety, the same effect has been demonstrated [71]. Le and colleagues [50] suggest that the mechanism behind this effect may be how awe makes a person focus on either themselves or others. They call this a "self-distanced" perspective or a "self-immersed" perspective. Their results would indicate that awe inspires a self-distanced perspective in their participants, which consequently made the obstacles in their lives seem less daunting, and the participants less distressed. That said, it is important to note that the studies that talked about distress and awe almost always used distress as an umbrella term to cover a lot of negative emotional experiences that could be treated as distinct in other work. However, there is no evidence to suggest that there is a negative relationship between emotional distress and awe, therefore:

- **H2)** Exposure to awe-inducing stimuli will reduce emotional distress in content moderators compared to the emotional distress levels of content moderators who are not awe-induced.

2.2 Cuteness

Awe is not, however, the only emotional stimuli that has the potential to reduce the negative effects of content moderation work: the relatively under-studied effect of cuteness has also been shown to have potential when it comes to making difficult situations more bearable [36, 73]. The sets of literature dedicated to awe and cuteness respectively mirror their topic, with awe literature being vast and well-established, and cuteness literature being young and small by comparison. A major portion of the literature on cuteness actually focuses on cuteness' effect on compassion and care, particularly in relation to mothers and their infants (e.g., [30, 49]). This literature stresses that the cuter something is, the more that people – usually mothers (see [55, 84]) – will be careful around it, and pay attention to it [63]. Essentially, if something is cute, we want to nurture it and protect it from harm [30, 44, 83]. This effect is not relegated to infant humans, though; it has also been shown in animals [10, 12, 54], and even objects [62, 93]. This desire to protect could have interesting effects for the world of content moderators.

In the case of stress, the admittedly sparse literature that exists would suggest that cute stimuli could have stress reduction properties [59, 93]. Much of this work focuses on animals (e.g., [59, 73]), and only occasionally separates out the "cute" factor from all the other parts of an animal. Reynolds and Rabschutz [73], for instance, studied exposure to live, in-person animals and how that can reduce stress in university students. Although this intervention was successful, cuteness in particular was not separated out as a specific variable of interest; as the animals in the study were puppies, however, chances are good that they were reasonably cute. Myrick [59], on the other hand, focused on exposure to animals online – cats, specifically – and found that her participants had a reduced perception of stress. Not only this, but the same participants specified that cuteness was a major part of their motivation to use internet cats as a stress-reduction technique. Although literature specific to cuteness rarely references stress directly (for an exception, see [93]), work on animals would suggest that exposure to them—in vivo or otherwise—should lead to stress reduction. Therefore:

- **H3)** Exposure to cute stimuli will reduce stress in content moderators compared to the stress levels of content moderators who do not see cute stimuli.

In terms of fatigue, cuteness shares the same effect as awe – an increase in the motivation to engage in prosocial behavior (e.g., [49, 96])—in addition to a general desire to care (e.g., [63]), and

so we would expect the same effect of compassion fatigue to occur with cute stimuli. However, when it comes to emotional distress, the evidence is more mixed [2, 93]. The majority of studies that discuss cuteness and emotional distress together talk about how higher levels of cuteness make parents particularly empathic and compassionate in regards to the distress of their children (e.g., [82, 84]). This has yet to be demonstrated with other cute stimuli like the baby animals used in the present study, though, nor are the animals in any way presenting distress in the images we procured. In terms of cuteness more generally as a variable of interest, there are only two studies that directly relate it to distress: Almanza-Sepúlveda and colleagues' study [2], and Takamatsu's [93]. In the first of these studies, the measure of cuteness has a significant negative relationship with personal distress [2], but in the second study, there was no relationship at all between the two variables [93]. Thus, the research would suggest that there is either no relationship or a negative one, but remains inconclusive in terms of which is more likely. All of this taken into account, we posit the following research question and hypothesis:

- **RQ2)** Which has a more powerful effect on fatigue: a cute stimulus, or accumulated content moderation experience?
- **H4)** Exposure to cute stimuli will reduce emotional distress in content moderators compared to the emotional distress levels of content moderators who do not see cute stimuli.

3 METHODS

The present study takes the form of an experiment with three conditions: control (30-second break every 20 posts moderated, no images), cuteness (30-second break every 20 posts moderated, baby animal images), and awe (30-second break every 20 posts moderated, awe-inspiring landscape images). Participants are randomly assigned to each using SurveyGizmo's built-in tool for random assignment. We used a between-subject design with approximate 120 participants for each condition. This study was reviewed and approved by the institutional review board of a mid-sized American university. We recruited participants from Amazon's Mechanical Turk (MTurk). We first designed a pre-screen survey to ensure that participants had some moderation experience. Participants would receive a "HIT" (code) to participate in the online experiment if they qualified. We monitored the data collection process for about three weeks. We terminated the survey at 387 responses (control 138, awe 123, and cuteness 126), as we had exceeded the total sample size suggested by an a-priori power calculation ran in G*Power seeking a medium effect according to Cohen's conventions (0.25). The projected total sample size at 0.95 observed power with an alpha of 0.05 and three groups was 252, and so to be conservative in case our effect size was smaller than medium (there was, to our knowledge, no existing estimation of effect sizes in content moderator interventions of this type in extant literature), we aimed for a minimum of 300 participants.

3.1 Participants

In total, 387 participants aged 23 to 77 ($M = 37.78$, $SD = 10.30$) completed the study. Of these participants, 153 (39.5%) identified as women, 230 (59.4%) identified as men, one (0.3%) as non-binary, and the rest (0.8%) did not disclose their gender identification. Some (199; 51.4%) were serving as moderators at the time of participation: 61 for Facebook, 39 on Reddit, 13 on Twitch. tv, 10 on Twitter, 9 on YouTube, 3 on Discord, 2 on Instagram, one on Meetup, one on Slack, one on Telegram, one on Tumblr, and 36 for various websites and forums. Among them, 22 participants were currently serving as content moderators on multiple platforms at once. Of these 199, 160 indicated how long they had occupied their position: 32 (20.0%) for less than a year, 38 (23.8%) for one to two years, 31 (19.4%) for two to three years, 19 (11.9%) for three to four years, 7 (4.4%) for four to five years, and 33 (20.6%) for five or more years. The majority of them (41; 25.6%) moderated

from 5 to 9 hours per week, although 24 (15.0%) only moderated four hours or less per week. The rest either moderated from 10-14 hours per week (39; 24.4%), 15-19 hours per week (25; 15.6%), 20-24 hours per week (17; 10.6%), 25-29 hours per week (2; 1.3%), 30-39 hours per week (7; 4.4%), or 40+ hours per week (5; 3.1%).

However, almost all participants (380; 98.2%) had previously served as a content moderator on some kind of platform. Most of them (126) had previously served as moderators for Facebook, but 93 others on Reddit, 30 on Twitch, 14 on Twitter, 13 on YouTube, 8 on Discord, 3 on Instagram, 2 on 4chan, 2 more on Pinterest, 1 on MySpace, 1 on Meetup, and 87 on various websites and forums. 57 of them served on multiple platforms in the past. Of these 380, 333 (87.6%) indicated how long they had served in their position: 67 (20.1%) had served less than a year, 106 (31.7%) for one to two years, 65 (19.5%) for two to three years, 39 (11.7%) for three to four years, 14 (4.2%) for four to five years, and 43 (12.9%) for five or more years. The majority (98; 29.4%) of these moderated from 5 to 9 hours a week, but 46 (13.8%) moderated four or fewer hours a week. The rest moderated from 10-14 hours per week (70; 21.0%), 15-19 hours per week (36; 10.8%), 20-24 hours per week (55; 16.5%), 25-29 hours per week (7; 2.1%), 30-39 hours per week (12; 3.6%), or 40+ hours per week (9; 2.7%).

3.2 Procedure

After providing digital consent to participate in the online experiment, participants are taught what kinds of posts need to be removed – i.e., sexism, racism, threats, etc. – and are given the opportunity to practice on five posts and receive feedback on their choices (were they correct or not, and why; this is only given after the practice rounds, not throughout the study). After successfully completing the five practice rounds, participants reviewed 400 statements and indicated which ones should be removed and which are ok to keep based on the training they just received ².

Throughout the moderation exercise, after every 20 posts (five questions), participants will be given the opportunity to take a 30-second break (as suggested by [11], as well as [53]), and depending on their condition, the screen is occupied by either the words “Take a break!” (control), a picture of a baby animal and the words “Take a break!” (cuteness), or a picture of a beautiful landscape and the words “Take a break!” (awe). Pictures of baby animals and landscapes were selected based on literature (e.g., [63, 68]) and were pre-tested by a voluntary sample of personal connections to ensure that they were equally cute or awe-inspiring, depending on the image, and elicited approximately the same levels of positive affect (as tested using Karim et al.’s short-form PANAS [43]). The order of presentation of these images was randomized, but all participants in each condition saw all of that condition’s images. After the participant made a decision about the final post (ok or unacceptable), participants were asked to provide their age and gender and expand upon their previous content moderation experience (if any), as well as fill out pre-validated measures of stress [17], fatigue [90], and emotional distress [45].

Finally, participants were thoroughly debriefed and given the option to follow a brief mindfulness-based stress reduction (MBSR) guided meditation [1]. This treatment has been repeatedly shown to reduce stress in a variety of people and situations (for a full discussion of MBSR and stress, see [40]). Participants were also provided with a list of free mental health resources available online, in the event that the meditation is not enough to relieve their emotional or psychological discomfort. At this point, they were given a code that they could enter into the MTurk link they originally clicked to start the study in order to receive their \$12.50 USD compensation for participating in about 45 minutes to one hour.

²Due to an error in the initial survey design, participants could not be truly randomly allocated using an algorithm. To simulate random allocation, participants were instructed to click one of three links (all with the same title, except Option 1, 2, or 3 appeared afterward) to participate in the study.

3.3 Materials

Warning: This section may contain offensive content.

3.3.1 Stimuli – Edited Tweets. As previously mentioned, participants simulated moderating 400 edited social media posts. These were initially gathered by undergraduate researchers from actual Twitter posts they found by browsing the social media platform. Undergraduate researchers first read Twitter's community guidelines and a list of banned words as well. Next, they used Google to search for tweets that violated those guidelines or contained these banned words. These tweets were then edited by these same students to remove any identifiable information or additional context. They also developed fake tweets that would be similar to these banned tweets they found online. These posts then were discussed and modified in a lab meeting to ensure that it was commonly understood that these tweets should have been moderated. More specifically, lab members voted on whether or not to ban each tweet, placing themselves in the role of content moderators, much like the present experiment, and only tweets that the majority of lab members considered offensive were marked as banned; others were considered acceptable. Each question is prefaced with a description statement with six options (5 negative posts and 1 named “None of these posts need removing”). Below is an example of a typical question in the survey.

“Please select the posts that should be removed according to the training you received. Remember that you can always click away, or choose not to complete tasks that make you too uncomfortable.”

- “Can’t believe the shooter couldn’t kill more people #iCanDoBetter”
- “Fuck niggers”
- “Gay night club? Just disgusting and I think that gunman did a good job”
- “All Muslims are terrorists”
- “Are all these ching chongs next to me related?”
- “None of these posts need removing”

3.3.2 Stimuli – Cute and Awe-Inspiring Images. At the outset, a group of 40 landscape images and 40 baby animal photos (these two groups were chosen based on extant literature suggesting that these would induce the two desired emotions). These were then piloted on a group of 20 volunteers from the undergraduate population and their families. Participants were asked to rate how cute/awe-inspiring the image was, and also complete the PANAS scale to determine the affect (positive or negative) each image inspired that we might select the images that were the most similar for each condition. The t-test between the target attributes – cuteness or awe – revealed no significant difference ($t(37) = -0.55$, $p = .58$), meaning that the final selection of images (40 out of 80) are equally cute/awe-inspiring. However, the two t-tests for positivity ($t(35) = 9.23$, $p < .001$) and negativity ($t(38) = -4.87$, $p < .001$) were both significant, meaning that the cute photos inspired significantly more negative affect, while the awe-inspiring photos inspired significantly more positive affect. The full survey with 400 posts, stimuli about awe and cuteness, and measures, is in the supplementary file.

3.3.3 Stress. To measure the stress our participants experienced after performing approximately 45 minutes of simulated content moderation, they completed Cohen and colleagues' perceived stress scale (PSS) [17]. This scale is designed to measure the perceived stress level of a participant at a given point in time, or for a given interval of time. It consists of 10 items that describe common symptoms of stress, and asks participants how often they have experienced said symptom – for the present study, the time period given was “while you were completing the moderation task” (e.g., “How often have you felt that you were on top of things?”). These were rated on a 5-point Likert scale ranging from 1 (Never) to 5 (Very often). The scale performed reliably, $\alpha = .85$.

3.3.4 Fatigue. To measure participants' levels of fatigue after the experiment, they completed the short form of the Multidimensional Fatigue Symptom Inventory (MFSI-SF), developed by the Moffitt Cancer Center in collaboration with the University of Florida in 1998 [90]. Similar to the PSS, this measure is designed to capture a person's perceived fatigue at a specific point, or for a given interval, in time. It, too, presents a series of symptoms, and in the present study, participants are asked to indicate how true the statement (including a symptom) is for them at the moment (e.g., "I feel run down", "I am worn out"); this is done using a 5-point Likert scale ranging from 1 (Not at all) to 5 (Extremely). The scale consists of five subscales: general fatigue, physical fatigue, emotional fatigue, mental fatigue, and vigor, which is the opposite of fatigue. Each subscale performed reliably (vigor = .91, general fatigue = .95, physical fatigue = .93, emotional fatigue = .94, mental fatigue = .93) as did the complete scale when treated as a whole, $\alpha = .96$.

3.3.5 Emotional Distress. Finally, to measure emotional distress, participants completed the National Comorbidity Survey's Emotional Distress Scale [45]. Again, much like the PSS and MFSI-SF, this measure presents participants with symptoms of emotional distress and asks them to indicate how true each statement is for them at the moment of filling out the measure (e.g., "Worried too much about things", "Frightened"). This is done using a 5-point Likert scale ranging from 1 (Not at all) to 5 (Extremely). This scale also performed reliably, $\alpha = .96$.

4 RESULTS

All analyses were performed in SPSS 26 [35]. To examine the effects of the experimental conditions (control, awe-inspiring stimuli, cute stimuli) and content moderation experience on the fatigue, stress, and emotional distress participants experienced after completing the experiment, we first divided the moderation experience ($M = 2.87$, $SD = 1.60$) into two groups by the median, with the "less experienced" group having moderated for two years or less, and the "more experienced" group having moderated for over two years. We then conducted a series of two-way ANOVAs with Tukey post hoc tests to examine the relationships between these variables.

4.1 Awe and Cuteness' Effects on Stress

For stress, the ANOVA revealed that only moderation experience ($F(1,328) = 11.20$, $p = .001$, $\eta^2 = .03$) had significant effects, with more experienced moderators ($M = 2.54$, $SD = 0.67$) feeling more stressed than less experienced moderators ($M = 2.30$, $SD = 0.78$), irrespective of the experimental condition ($F(2,328) = 2.97$, $p = .05$, $\eta^2 = .02$, observed power = .58), as shown in Figure 1c. We expected that exposure to awe-inducing stimuli would reduce stress in content moderators compared to the stress levels of content moderators who are not awe-induced (H1), which was not supported. We also anticipated that exposure to cute stimuli would reduce stress in content moderators compared to stress levels of content moderators who do not see cute stimuli (H3), which was also not supported.

4.2 Awe and Cuteness' Effects on Fatigue

For fatigue, the ANOVA analysis showed that both moderation experience ($F(1,328) = 14.00$, $p < .001$, $\eta^2 = .04$) and experimental conditions ($F(2,328) = 5.23$, $p = .01$, $\eta^2 = .03$) had significant effects, but that there was also a significant interaction between these two variables ($F(2,328) = 3.52$, $p = .03$, $\eta^2 = .02$). The significant interaction between moderation experience and condition is visualized in Figure 1a. The partial eta squared indicated that moderation experience had a stronger impact than experimental conditions on fatigue (RQ1 and RQ2). Simple effects for the interaction revealed that the difference between the three conditions is significant at the more experienced level ($F(2,328) = 8.17$, $p < .001$, $\eta^2 = .05$), but not significant at the less experienced level ($F(2,328) = 0.57$, $p = .57$, $\eta^2 = .00$, observed power = .14). Specifically, more experienced moderators had a higher fatigue level in




Fig. 1. Interaction effects

the cute ($p = .00, d = .47$) and awe ($p = .00, d = .46$) condition than in the control condition, and there was no significant difference in this regard between the cute and awe condition ($p = .95, d = .01$).

4.3 Awe and Cuteness' Effects on Emotional Distress

For emotional distress, the ANOVA indicated that both moderation experience ($F(1,328) = 7.98, p = .01, \eta^2 = .02$) and experimental conditions ($F(2,328) = 9.92, p < .001, \eta^2 = .05$) had significant effects, but that there was also a significant interaction effect between these two variables ($F(2,328) = 3.51, p = .03, \eta^2 = .02$). The significant interaction between moderation experience and condition is visualized in Figure 1b. More experienced moderators ($M = 1.85, SD = 0.99$) felt more emotionally distressed than less experienced moderators ($M = 1.63, SD = 0.86$). Simple effects for the interaction revealed that the difference between the three conditions is significant at the more experienced level ($F(2,328) = 12.72, p < .001, \eta^2 = .07$) but not significant at the less experienced level ($F(2,328) = 0.87, p = .42, \eta^2 = .01$, observed power = .20). Specifically, more experienced moderators had higher emotional distress in cute ($p = .00, d = .64$) and awe ($p = .00, d = .77$) conditions than in the control condition, and there was no significant difference in this regard between cute and awe condition ($p = .50, d = .13$). Thus, our anticipation that exposure to awe or cute stimuli would reduce emotional distress in content moderators compared to the emotional distress levels of content moderators who do not see cute or awe stimuli (H2 and H4) were not supported.

5 DISCUSSION

The goal of the present study was to determine if by adding cute or awe-inspiring stimuli to break times, we could reduce the ill-effects of performing content moderation duties on social media or other online fora. More specifically, we wanted to see if the addition of these images could decrease stress (H1 & H3), fatigue (RQ1 & RQ2), and emotional distress (H2 & H4) after completing an experimental content moderation task. The evidence would suggest that the addition of cute or awe-inspiring stimuli has the opposite effect of our prediction: it actually increases stress levels among experienced content moderators (H1 & H3); the same is true for emotional distress (H2 & H4). There are a few possible explanations for this reversed effect. Irrespective of the exact explanation, this reversed effect could be an indication of content moderation as a unique context, as it seems to contradict findings in other contexts in which positive stimuli can overrule the effects of negative stimuli (e.g., [19]). Though we set out to find out how positive emotional stimulation could alleviate the negative side-effects of being a content moderator, what we found instead is that the negativity is so entrenched in some moderators that positive stimulation actually amplifies it, a finding that could have several explanations.

The first possible explanation, also the simplest, is that the addition of any new stimuli during a break disrupts the relaxation that simply 'being' can bring (e.g., [67, 98]). In other words, it could

be that, when it comes to stress relief, doing nothing is better than doing something. However, this does not explain why emotional distress only increases significantly among experienced moderators when images are present during breaks, as opposed to creating the same effect across all content moderators. For this effect, we posit the following explanation: the combination of compassion fatigue and increased prosocial tendencies inspired by cuteness and awe.

As discussed earlier, compassion fatigue is essentially a cumulative effect on wellbeing that occurs in people who work or volunteer in prosocial contexts [56], such as healthcare (e.g., [46] or, in the present study, content moderation. A person experiencing compassion fatigue has spent so much time helping others that they develop a kind of increased susceptibility to taking on others' problems and struggles, to their own detriment (see [64]). Although it is called "fatigue", compassion fatigue is also heavily associated with stress and burnout [46, 56, 58, 60, 65, 92], which could help explain why it may have affected all of our outcome variables in the present study. Though we anticipated it may have an effect in terms of fatigue (RQ1 & RQ2), which we appear to have found, as experienced moderators suffered higher levels of fatigue than inexperienced moderators, particularly when exposed to extra stimuli, it would seem that compassion fatigue is all-encompassing in terms of its negative effects. Both cute and awe-inspiring stimuli have been shown to produce elevated prosocial orientation and desire in people [32, 49, 84, 96], in addition to the positive emotions we anticipated would reduce the negative outcomes we measured [41, 47, 50, 59, 93, 95]. It would seem that the prosocial effects of these stimuli activate the compassion fatigue of experienced moderators, overriding any positive effects they may have had and instead increasing the stress, fatigue, and emotional distress caused by the desire to help. Although not explicitly measured, theoretically, the content moderator would see a kitten or puppy and be reminded of things like vulnerability and fear and the innocence they are trying to protect in netizens, emphasizing their need to experience the negativity themselves in a self-sacrificial capacity. The longer the content moderator has been at this task, the more emotionally exhausted they would theoretically become being reminded - by the positive aspects of the world they are trying to protect - of their perceived neverending task and the importance that they continue working at it. Thus, based on our results, it would seem that it is better to do nothing on breaks in professions or volunteer positions where people are at risk of suffering from compassion fatigue.

This is, to our knowledge, the first article to make the association between compassion fatigue and content moderation specifically; if this is indeed the reason positive emotional stimuli ended up producing negative emotional consequences, this means that although cuteness and awe-inspiring images are not the right course of action, there are other options outlined in compassion fatigue research that could also be effective at reducing negative wellbeing effects in the content moderation field. The most popular in extant literature, which comes with a plethora of literature in its own right and was mentioned earlier, is mindfulness-based stress reduction, or MBSR [40, 46]. In fact, we used an MBSR guided meditation as a precaution in this study to ensure that participants did not leave the project depressed or overly anxious. In terms of the specific variables addressed in the present study, Janssen and colleagues [40] found MBSR to be an effective way to reduce stress and emotional exhaustion, a concept similar to emotional distress and related to burnout. In addition, mindfulness is based on the core principle of just being in the moment – in other words, doing nothing but focusing in silence on what the experience of existence [70]. If compassion fatigue is indeed at the heart of our results, then programs like those used by Kneff and colleagues [60] and Klein and colleagues [46] could be tailored to the content moderation context. If, however, the explanation is that nothing is better than something when it comes to breaks, then MBSR would still be theoretically a helpful practice for content moderators, as it would emphasize the nothingness that brings the relaxation they need. In both cases, however, MBSR would have to be implemented either in longer breaks than what was presented in the present study (see [37, 88]) or

as a separate program adjacent to work days. Micro-breaks without additional stimuli could still be beneficial alongside a more structured MBSR program, but this would require additional research to confirm with certainty.

Another possible explanation for the emotional effects we saw in the present study is the idea of affective contrast, which is when the presence of opposite-valence stimuli with affective dimensions amplify each other's emotional effects [4, 33]. The baby animals and the beautiful landscapes are designed to evoke positive emotions, but the revised tweets used in the experiment are designed to evoke negative emotions. Some articles discussing affective contrast would suggest that presenting positive stimuli in the midst of negative ones should have actually stabilized or reversed the effects of the negative stimuli (e.g., [16]). However, given certain individual differences in terms of optimism and pessimism, it is also possible that people could be more sensitive to contrasts in the negative sense, creating a negative overall affect (e.g., [28]). Similarly, Zhang and Covey [100] found that whether an affective contrast came out as a net positive or net negative depended entirely on the person's preconceptions, which although we can guess were negative given participants' experience with content moderation, we did not explicitly measure in advance. Referring back to the theoretical experience of content moderators exposed to these positive stimuli, this would mean that content moderators who are more experienced are probably framing their task as fruitless and neverending, while the newer content moderators are more motivated to work hard to protect the world's beauty and innocence. In short, although it could be a potential explanation for our findings, the literature is unclear regarding exactly which direction the contrast should go. Future studies could examine this in greater detail by running similar studies including personality and temperament measures, as well as preconceptions regarding content moderation work.

However, another possibility is that we do not see an immediate impact from a cuteness or awe-based intervention because, like compassion fatigue, the effects of positivity are also cumulative. In their study comparing active recovery via a digital game to active recovery via mindfulness practice, Collins et al. [19] found that although essentially equally effective in the short term, the fun stimulation of a game was more effective than the mindfulness program in the long run when it came to increasing energy and decreasing fatigue. This tells us two things: 1) dynamic positive stimulation may be more effective than static, and 2) we may need more exposure to positive stimuli over time to see their ultimate effect. Based on the knowledge we have and our current results, it would seem as though in the short term for micro-break interventions, switching off completely is ideal, particularly for moderators who have already been in the position for many years. However, over time, it is still possible that regular exposure to positive stimuli, particularly if they are dynamic in some way, might be beneficial. To adequately test this long-term possibility, future studies will need to employ a longitudinal design and will also have to test baseline attitudes toward technology and media use (see [72]), as well as pessimism and optimism [28], in order to tease apart the effects of compassion fatigue, affective contrasts, and simple individual differences to see what is ultimately best for content moderators who wish to make it their permanent career.

5.1 Practical and Design Implications

In line with recent work by Steiger et al. [89] to provide the necessary programmatic and technological interventions to offer resources and assistance for mental health, we provide several design interventions that have the potential to protect content moderators or at least mitigate the position's ill effects. Due to the results of positive emotional stimulation being a net negative in the present work, we propose interventions that take the opposite approach and encourage alternative break-time strategies.

The first is a technological intervention that involves how breaks are conducted in a content moderation work or volunteer shift. Although the importance of breaks for any kind of work is

well-established in literature (e.g., [9]), the exact content of these breaks is not always made entirely clear. Although some articles do refer to a “rest break” [26, p.221], it is not always made explicit what exactly happens during these breaks. Depending on the real-life work context, a break could consist of browsing social media on one’s phone, eating food, or staring at a wall, to name a few possibilities; not all articles specify whether this kind of behavior is prohibited during the assigned breaks or not, and what kinds of break behaviors are beneficial or detrimental to the break’s efficacy. As previously stated, our findings would seem to indicate that doing nothing on a break is better than doing something, even if that something is positive, like looking at baby animal photos or beautiful landscapes. We suggest incorporating restful breaks via bots in content moderation user interfaces (UIs). In addition to having a bot remind content moderators to take a 7.5-minute break after every 50 minutes, as proposed by [11], the bot should encourage content moderators to remove themselves from the computer, or at the very least from the UI, for the duration of the break. This could even be done automatically, with the UI shutting down or presenting a blank screen and restarting after the designated break time. Of course, preventing content moderators from doing anything during their breaks would be difficult if not impossible remotely, but encouragement to do as little as possible during breaks could increase the effectiveness of said breaks for them.

The second is a programmatic intervention that could be implemented as a form of a scheduled rotation of content moderators, to reduce the effects of compassion fatigue . We found that the longer a content moderator had been working (in years), the stronger the negative effects on their stress, emotional distress, and fatigue. This is consistent with compassion fatigue literature, where extended prosocial work leads to burnout or similar symptoms (e.g., [92]). In the case of medical work, it has been suggested that to get at the root of burnout, the job has to be stripped down to its core, removing any responsibilities outside of providing medical care, in order to reduce the negative effects of being a doctor or nurse [87]. However, the present study shows that even the most basic of content moderation tasks adds to cumulative negative effects of the position, meaning that there is little that can be ‘stripped away’. Thus, we propose extended breaks for content moderators, much like the sabbatical system used by university professors. Sabbaticals have been shown to work to reduce the effects of burnout - which mirror the negative, cumulative effects of content moderation work - in the case of crisis workers, who also traditionally suffer from compassion fatigue [42]. By implementing ‘shifts’ of content moderators, the accumulation of ill-effects can be disrupted, allowing for less buildup over time. Future research would have to be conducted to determine the exact duration of a content moderation position before they get rotated off to sabbatical (e.g., one week, one month, two years).

Finally, content moderators should receive systematic training to learn how to effectively combat the negative effects of their work as a programmatic intervention. While health professionals go through extensive training and years of university to prepare for their jobs, content moderators - especially volunteer content moderators - receive little to no training before they begin work [81]. In addition, while several programs have been developed to help reduce the stress and fatigue of other professionals at risk of compassion fatigue (e.g., [39, 92]), the same is not true of content moderators, who receive only inadequate support from the platform [24, 25, 75]. We have already discussed the potential benefits of mindfulness training for content moderators, but other targeted interventions relating to sleep, anxiety, or general physical wellness could also be explored. Now that we know that what content moderators go through could be akin to the experience of crisis workers, police officers, or healthcare professionals, we have many new avenues of promising intervention to validate in the content moderation context.

5.2 Limitations and Future Directions

Although this study has important implications for how best to care for content moderators, it is not without limitations. Despite every precaution taken, due to the aforementioned lack of true random allocation, participants could have participated more than once. In addition, since it was an online experiment, it was impossible to monitor exactly what participants did during their breaks. It is possible that they watched YouTube videos or played a video game or did something else during the control conditions or experimental conditions, although the short duration of the breaks would have made that difficult. Future studies could be conducted in the laboratory to better control for these kinds of possibilities. Future studies could also vary the type of stimulus used, as there could be a difference between using static, image stimuli as we did in the present work, and using dynamic video-based stimuli. Finally, although the measures we used covered many of the aspects of compassion fatigue and served as a kind of facsimile for a direct measure of the construct, we did not give our participants a compassion fatigue-specific measure. This is partially because there are only two that we were able to find: The Compassion Fatigue-Short Scale [99] and the Nurses' Compassion Fatigue Inventory [77], both of which are highly tailored to the medical professions. Future work could either create a new, more general compassion fatigue measurement tool, or adapt these existing ones to new contexts to explore the link between compassion fatigue and content moderation in greater depth. Finally, although unlikely, it is possible that participants in one group were feeling particularly distraught prior to starting the experiment. Future studies can measure negative emotion as a pre-test to ensure that this is not the case.

Beyond addressing the aforementioned limitations, future studies could also test some of the relationships and concepts we found here. For instance, we used 30-second breaks between every 20 items moderated, but this is due to the literature on worktime breaks (e.g., [9]), not due to how breaks currently work for content moderators in the field. Different break time durations, and different spacing between these breaks, should be explored to discover what is optimal for the wellbeing of content moderators. In addition, based on our results, it would seem that mindfulness interventions have the potential to greatly impact content moderators' experience of their work, but this has yet to be tested directly; it should be compared to other, similar interventions (e.g., transcendental meditation, yoga, relaxing music) to see if it is really the most effective option out there for content moderators.

6 CONCLUSIONS

In conclusion, although neither awe nor cuteness was found to be efficacious for reducing the negative side effects of performing content moderation, we did uncover some important information when it comes to taking care of the people who protect us on the internet. For one, the effects of stress, fatigue, and emotional distress appear to be cumulative, just like they are in other caring professions; the longer someone serves, the worse they seem to suffer the job's ill-effects. There is also the strong suggestion of a connection between compassion fatigue and content moderation, giving hope that some of the ways we take care of our healthcare professionals and crisis workers may help alleviate the difficulties of content moderators as well. Instead of alleviating the negative effects of the content moderation work, adding positive stimuli actually seems to have the opposite of the intended effect, leading to new design ideas around work scheduling and technological applications. By continuing to research new and established ways to protect the protectors of the internet, we can help ensure happier, healthier cyberspace for everyone.

ACKNOWLEDGMENTS

This research was funded by by National Science Foundation (Award No. 1928627). Thanks to the research assistants in SocialXLab at NJIT for data collection.

REFERENCES

- [1] 2015. MINDFULNESS 10 Minutes Breath Meditation - YouTube. <https://www.youtube.com/watch?app=desktop&v=AEiuvYQfWoA>
- [2] Mayra L. Almanza-Sepúlveda, Aya Dudin, Kathleen E. Wonch, Meir Steiner, David R. Feinberg, Alison S. Fleming, and Geoffrey B. Hall. 2018. Exploring the morphological and emotional correlates of infant cuteness. *Infant Behavior and Development* 53 (2018). <https://doi.org/10.1016/j.infbeh.2018.08.001>
- [3] Craig L. Anderson, Maria Monroy, and Dacher Keltner. 2018. Awe in nature heals: Evidence from military veterans, at-risk youth, and college students. *Emotion* 18, 8 (2018). <https://doi.org/10.1037/emo0000442>
- [4] Margaret M. Bacon, Esther A. Rood, and M. F. Washburn. 1914. A Study of Affective Contrast. *The American Journal of Psychology* 25, 2 (4 1914), 290. <https://doi.org/10.2307/1413417>
- [5] Joe Bagliere. 2020. Science shows watching cute animals is good for your health - CNN. <https://edition.cnn.com/2020/09/27/us/watching-cute-animals-study-scn-trnd/index.html>
- [6] Yang Bai, Laura A. Maruskin, Serena Chen, Amie M. Gordon, Jennifer E. Stellar, Galen D. McNeil, Kaiping Peng, and Dacher Keltner. 2017. Awe, the diminished self, and collective engagement: Universals and cultural variations in the small self. *Journal of Personality and Social Psychology* 113, 2 (2017). <https://doi.org/10.1037/pspa0000087>
- [7] Paul M. Barrett. 2020. Research Highlights | Who Moderates the Social Media Giants? A Call to End Outsourcing. <https://www.stern.nyu.edu/experience-stern/faculty-research/who-moderates-social-media-giants-call-end-outsourcing>
- [8] Roukaya Benjelloun and Yassine Otheman. 2020. Psychological distress in a social media content moderator: A case report. *Archives of Psychiatry and Mental Health* 4, 1 (10 2020), 073–075. <https://doi.org/10.29328/journal.apmh.1001024>
- [9] Gerhard Blasche, Sanja Pasalic, Verena Maria Bauböck, Daniela Haluza, and Rudolf Schoberberger. 2017. Effects of rest-break intention on rest-break frequency and work-related fatigue. *Human Factors* 59, 2 (2017). <https://doi.org/10.1177/0018720816671605>
- [10] Marta Borgi, Irene Cogliati-Dezza, Victoria Brelsford, Kerstin Meints, and Francesca Cirulli. 2014. Baby schema in human and animal faces induces cuteness perception and gaze allocation in children. *Frontiers in Psychology* 5, MAY (2014). <https://doi.org/10.3389/fpsyg.2014.00411>
- [11] Wolfram Boucsein and Michael Thum. 1997. Design of work/rest schedules for computer work based on psychophysiological recovery measures. *International Journal of Industrial Ergonomics* 20, 1 (1997). [https://doi.org/10.1016/S0169-8141\(96\)00031-5](https://doi.org/10.1016/S0169-8141(96)00031-5)
- [12] Ralf C. Buckley. 2016. Aww: The emotion of perceiving cuteness. <https://doi.org/10.3389/fpsyg.2016.01740>
- [13] Jie Cai and Donghee Yvete Wohn. 2019. What are effective strategies of handling harassment on twitch? Users' perspectives. In *Companion of the ACM Conference on Computer Supported Cooperative Work*. 166–170. <https://doi.org/10.1145/3311957.3359478>
- [14] Elinor Carmi. 2019. The hidden listeners: Regulating the line from telephone operators to content moderators. *International Journal of Communication* 13 (2019).
- [15] Eshwar Chandrasekharan, Chaitrali Gandhi, Matthew Wortley Mustelier, and Eric Gilbert. 2019. CrossMod: A cross-community learning-based system to assist reddit moderators. In *Proceedings of the ACM on Human-Computer Interaction*, Vol. 3. 30. <https://doi.org/10.1145/3359276>
- [16] Sheung-Tak Cheng. 2006. Negative emotions make positive emotions more salient in well-being appraisal. *Personality and Individual Differences* 40, 3 (2 2006), 557–567. <https://doi.org/10.1016/j.paid.2005.08.002>
- [17] S. Cohen, T. Kamarck, and R. Mermelstein. 1983. A global measure of perceived stress. *Journal of health and social behavior* 24, 4 (1983). <https://doi.org/10.2307/2136404>
- [18] Silvia Collado and Héctor Marín Manrique. 2020. Exposure to Awe-Evoking Natural and Built Scenes Has Positive Effects on Cognitive Performance and Affect. *Environment and Behavior* 52, 10 (12 2020), 1105–1132. <https://doi.org/10.1177/0013916519868733>
- [19] Emily Collins, Anna Cox, Caroline Wilcock, and Geraint Sethu-Jones. 2019. Digital Games and Mindfulness Apps: Comparison of Effects on Post Work Recovery. *JMIR Mental Health* 6, 7 (7 2019), e12853. <https://doi.org/10.2196/12853>
- [20] Brandon Dang, Martin J. Riedl, and Matthew Lease. 2018. But Who Protects the Moderators? The Case of Crowdsourced Image Moderation. <http://arxiv.org/abs/1804.10999>
- [21] Alexander F. Danvers and Michelle N. Shiota. 2017. Going off script: Effects of awe on memory for script-typical and -irrelevant narrative detail. *Emotion* 17, 6 (2017). <https://doi.org/10.1037/emo0000277>

[22] Anubrata Das, Brandon Dang, and Matthew Lease. 2020. Fast , Accurate , and Healthier : Interactive Blurring Helps Moderators Reduce Exposure to Harmful Content. In *Proceedings of the Eighth AAAI Conference on Human Computation and Crowdsourcing (HCOMP-20)*. 10.

[23] Peter Dockrill. 2016. Just Looking at Photos of Nature Could Be Enough to Lower Your Work Stress Levels. <https://www.sciencealert.com/just-looking-at-photos-of-nature-could-be-enough-to-lower-your-work-stress-levels>

[24] Bryan Dosono and Bryan Semaan. 2019. Moderation Practices as Emotional Labor in Sustaining Online Communities. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems - CHI '19*. ACM Press, New York, New York, USA, 1–13. <https://doi.org/10.1145/3290605.3300372>

[25] Elizabeth Dwoskin. 2019. Inside Facebook, the second-class workers who do the hardest job are waging a quiet battle. <https://www.washingtonpost.com/technology/2019/05/08/inside-facebook-second-class-workers-who-do-hardest-job-are-waging-quiet-battle/>

[26] Julia Faucett, James Meyers, John Miles, Ira Janowitz, and Fadi Fathallah. 2007. Rest break interventions in stoop labor tasks. *Applied Ergonomics* 38, 2 (2007), 219–226. <https://doi.org/10.1016/j.apergo.2006.02.003>

[27] Eric L. Garland, Barbara Fredrickson, Ann M. Kring, David P. Johnson, Piper S. Meyer, and David L. Penn. 2010. Upward spirals of positive emotions counter downward spirals of negativity: Insights from the broaden-and-build theory and affective neuroscience on the treatment of emotion dysfunctions and deficits in psychopathology. *Clinical Psychology Review* 30, 7 (11 2010), 849–864. <https://doi.org/10.1016/j.cpr.2010.03.002>

[28] Andrew L. Geers and G. Daniel Lassiter. 2002. Effects of Affective Expectations on Affective Experience: The Moderating Role of Optimism-Pessimism. *Personality and Social Psychology Bulletin* 28, 8 (8 2002), 1026–1039. <https://doi.org/10.1177/01461672022811002>

[29] Tarleton Gillespie. 2018. *Custodians of the internet: Platforms, content moderation, and the hidden decisions that shape social media*. Yale University Press, New Haven. 1–288 pages. <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85051469782&partnerID=40&md5=8d850b5298b7e5dc1a1fcf4c427fe3da>

[30] Melanie L. Glockner, Daniel D. Langleben, Kosha Ruparel, James W. Loughead, Ruben C. Gur, and Norbert Sachser. 2009. Baby schema in infant faces induces cuteness perception and motivation for caretaking in adults. *Ethology* 115, 3 (2009). <https://doi.org/10.1111/j.1439-0310.2008.01603.x>

[31] Amie M. Gordon, Jennifer E. Stellar, Craig L. Anderson, Galen D. McNeil, Daniel Loew, and Dacher Keltner. 2017. The dark side of the sublime: Distinguishing a threat-based variant of awe. *Journal of Personality and Social Psychology* 113, 2 (2017). <https://doi.org/10.1037/pspp0000120>

[32] Fang Guan, Jun Chen, Outong Chen, Lihong Liu, and Yuzhu Zha. 2019. Awe and prosocial tendency. *Current Psychology* 38, 4 (2019). <https://doi.org/10.1007/s12144-019-00244-7>

[33] Albert J. Harris. 1929. An Experiment on Affective Contrast. *The American Journal of Psychology* 41, 4 (10 1929), 617. <https://doi.org/10.2307/1414718>

[34] Katherine A. Hildebrandt and Hiram E. Fitzgerald. 1978. Adults' responses to infants varying in perceived cuteness. *Behavioural Processes* 3, 2 (7 1978), 159–172. [https://doi.org/10.1016/0376-6357\(78\)90042-6](https://doi.org/10.1016/0376-6357(78)90042-6)

[35] IBM Corp. 2020. IBM SPSS Statistics for Windows, Version 27.0.

[36] Hadrian Ilhan Aslan, Julian Kraus Burkhardt, and Elisabeth André. 2016. Hold my heart and breathe with me: Tangible somaesthetic designs. In *ACM International Conference Proceeding Series*, Vol. 23-27-Octo. <https://doi.org/10.1145/2971485.2996727>

[37] Kim E. Innes, Terry Kit Selfe, Kathleen Brundage, Caitlin Montgomery, Sijin Wen, Sahiti Kandati, Hannah Bowles, Dharma Singh Khalsa, and Zenzi Huysmans. 2018. Effects of Meditation and Music-Listening on Blood Biomarkers of Cellular Aging and Alzheimer's Disease in Adults with Subjective Cognitive Decline: An Exploratory Randomized Clinical Trial. *Journal of Alzheimer's Disease* 66, 3 (11 2018), 947–970. <https://doi.org/10.3233/JAD-180164>

[38] Internet-live-stats. 2020. Internet Live Stats - Internet Usage & Social Media Statistics. <https://www.internetlivestats.com/>

[39] Lois James, Charles H. Samuels, and Fiona Vincent. 2018. Evaluating the effectiveness of fatigue management training to improve police sleep health and wellness: A pilot study. *Journal of Occupational and Environmental Medicine* 60, 1 (1 2018), 77–82. <https://doi.org/10.1097/JOM.0000000000001174>

[40] Math Janssen, Yvonne Heerkens, Wietske Kuijer, Beatrice Van Der Heijden, and Josephine Engels. 2018. Effects of mindfulness-based stress reduction on employees' mental health: A systematic review. <https://doi.org/10.1371/journal.pone.0191322>

[41] Yannick Joye and Jan Willem Bolderdijk. 2014. An exploratory study into the effects of extraordinary nature on emotions, mood, and prosociality. *Frontiers in Psychology* 5, OCT (2014). <https://doi.org/10.3389/fpsyg.2014.01577>

[42] Chulhee Kang, Miok Kim, and Jongeun Lee. 2011. The effects of a month-long sabbatical program on helping professionals of nonprofit human service organizations in South Korea: Burnout, general health, organizational commitment, and the sense of well-being. *Administration in Social Work* 35, 1 (1 2011), 20–45. <https://doi.org/10.1080/03643107.2011.533619>

[43] Jahanvash Karim, Robert Weisz, and Shafiq Ur Rehman. 2011. International positive and negative affect schedule short-form (I-PANAS-SF): Testing for factorial invariance across cultures. In *Procedia - Social and Behavioral Sciences*, Vol. 15. Elsevier, 2016–2022. <https://doi.org/10.1016/j.sbspro.2011.04.046>

[44] Annemiek Karreman and Madelon M.E. Riem. 2020. Exposure to infant images enhances attention control in mothers. *Cognition and Emotion* 34, 5 (2020). <https://doi.org/10.1080/02699931.2019.1689101>

[45] Ronald C. Kessler, Katherine A. McGonagle, Shanyang Zhao, Christopher B. Nelson, Michael Hughes, Suzann Eshleman, Hans Ulrich Wittchen, and Kenneth S. Kendler. 1994. Lifetime and 12-Month Prevalence of DSM-III-R Psychiatric Disorders in the United States: Results from the National Comorbidity Survey. *Archives of General Psychiatry* 51, 1 (1994). <https://doi.org/10.1001/archpsyc.1994.03950010008002>

[46] Colleen J. Klein, Jami J. Rigenbach-Hays, Laura M. Sollenberger, Diane M. Harney, and Jeremy S. McGarvey. 2018. Quality of Life and Compassion Satisfaction in Clinicians: A Pilot Intervention Study for Reducing Compassion Fatigue. *American Journal of Hospice and Palliative Medicine* 35, 6 (2018). <https://doi.org/10.1177/1049909117740848>

[47] Alethea H.Q. Koh, Eddie M.W. Tong, and Alexander Y.L. Yuen. 2019. The buffering effect of awe on negative affect towards lost possessions*. *Journal of Positive Psychology* 14, 2 (2019). <https://doi.org/10.1080/17439760.2017.1388431>

[48] Meredith Krause. 2009. Identifying and Managing Stress in Child Pornography and Child Exploitation Investigators. *Journal of Police and Criminal Psychology* 24, 1 (2009). <https://doi.org/10.1007/s11896-008-9033-8>

[49] Morten L. Kringelbach, Eloise A. Stark, Catherine Alexander, Marc H. Bornstein, and Alan Stein. 2016. On Cuteness: Unlocking the Parental Brain and Beyond. <https://doi.org/10.1016/j.tics.2016.05.003>

[50] Phuong Q. Le, Thomas L. Saltsman, Mark D. Seery, Deborah E. Ward, Cheryl L. Kondrak, and Veronica M. Lamarche. 2019. When a small self means manageable obstacles: Spontaneous self-distancing predicts divergent effects of awe during a subsequent performance stressor. *Journal of Experimental Social Psychology* 80 (2019). <https://doi.org/10.1016/j.jesp.2018.07.010>

[51] Jin Li, Anke Li, Yu Sun, Hui'e Li, Lei Liu, Youlong Zhan, Wei Fan, and Yiping Zhong. 2019. The effect of preceding self-control on prosocial behaviors: The moderating role of awe. *Frontiers in Psychology* 10, MAR (2019). <https://doi.org/10.3389/fpsyg.2019.00682>

[52] Joseph Lichtenberg. 2015. Awe: Its bidirectional power - Veneration and inspiration or dread and inhibition. *Psychanalytic Inquiry* 35, 1 (2015). <https://doi.org/10.1080/07351690.2014.957138>

[53] Julian Lim and Kenneth Kwok. 2016. The Effects of Varying Break Length on Attention and Time on Task. *Human Factors* 58, 3 (2016). <https://doi.org/10.1177/0018720815617395>

[54] Anthony C. Little. 2012. Manipulation of infant-like traits affects perceived cuteness of infant, adult and cat faces. *Ethology* 118, 8 (2012). <https://doi.org/10.1111/j.1439-0310.2012.02068.x>

[55] Janek S. Lobmaier, Reiner Sprengelmeyer, Ben Wiffen, and David I. Perrett. 2010. Female and male responses to cuteness, age and emotion in infant faces. *Evolution and Human Behavior* 31, 1 (2010). <https://doi.org/10.1016/j.evolhumbehav.2009.05.004>

[56] Susan H. Lynch, Geoff Shuster, and Marie L. Lobo. 2018. The family caregiver experience – examining the positive and negative aspects of compassion satisfaction and compassion fatigue as caregiving outcomes. *Aging and Mental Health* 22, 11 (2018). <https://doi.org/10.1080/13607863.2017.1364344>

[57] James K. McNulty, Michael A. Olson, Rachael E. Jones, and Laura M. Acosta. 2017. Automatic Associations Between One's Partner and One's Affect as the Proximal Mechanism of Change in Relationship Satisfaction: Evidence From Evaluative Conditioning. *Psychological Science* 28, 8 (8 2017), 1031–1040. <https://doi.org/10.1177/0956797617702014>

[58] Holly Monaghan, Vanessa Rohlf, Rebekah Scotney, and Pauleen Bennett. 2020. Compassion Fatigue in People Who Care for Animals: An Investigation of Risk and Protective Factors. *Traumatology* (2020). <https://doi.org/10.1037/trm0000246>

[59] Jessica Gall Myrick. 2015. Emotion regulation, procrastination, and watching cat videos online: Who watches Internet cats, why, and to what effect? *Computers in Human Behavior* 52 (2015). <https://doi.org/10.1016/j.chb.2015.06.001>

[60] Kristin D. Neff, Marissa C. Knox, Phoebe Long, and Krista Gregory. 2020. Caring for others without losing yourself: An adaptation of the Mindful Self-Compassion Program for Healthcare Communities. *Journal of Clinical Psychology* 76, 9 (2020). <https://doi.org/10.1002/jclp.23007>

[61] Casey Newton. 2019. The Trauma Floor: The secret lives of Facebook moderators in America . <https://www.theverge.com/2019/2/25/18229714/cognizant-facebook-content-moderator-interviews-trauma-working-conditions-arizona>

[62] Hiroshi Nittono. 2016. The two-layer model of 'kawaii': A behavioural science framework for understanding kawaii and cuteness. *East Asian Journal of Popular Culture* 2, 1 (2016). https://doi.org/10.1386/eapc.2.1.79_{_}1

[63] Hiroshi Nittono, Michiko Fukushima, Akihiro Yano, and Hiroki Moriya. 2012. The Power of Kawaii: Viewing Cute Images Promotes a Careful Behavior and Narrows Attentional Focus. *PLoS ONE* 7, 9 (2012). <https://doi.org/10.1371/journal.pone.0046362>

[64] Anna G.W. Nolte, Charlene Downing, Annie Temane, and Marie Hastings-Tolsma. 2017. Compassion fatigue in nurses: A metasynthesis. *Journal of Clinical Nursing* 26, 23-24 (2017). <https://doi.org/10.1111/jocn.13766>

[65] Malin Norrman Harling, Elisabeth Höglman, and Elinor Schad. 2020. Breaking the taboo: eight Swedish clinical psychologists' experiences of compassion fatigue. *International Journal of Qualitative Studies on Health and Well-being* 15, 1 (2020). <https://doi.org/10.1080/17482631.2020.1785610>

[66] Lisa Parks. 2019. Dirty data: Content moderation, regulatory outsourcing, and the cleaners. *Film Quarterly* 73, 1 (2019). <https://doi.org/10.1525/FQ.2019.73.1.11>

[67] Eric Pfeifer, Henrike Fiedler, and Marc Wittmann. 2019. Enhanced relaxation in students after combined Depth Relaxation Music Therapy and silence in a natural setting. *Arts in Psychotherapy* 63 (2019). <https://doi.org/10.1016/j.aip.2019.02.006>

[68] Paul K. Piff, Pia Dietze, Matthew Feinberg, Daniel M. Stancato, and Dacher Keltner. 2015. Awe, the small self, and prosocial behavior. *Journal of Personality and Social Psychology* 108, 6 (2015). <https://doi.org/10.1037/pspi0000018>

[69] Claire Prade and Vassilis Saroglou. 2016. Awe's effects on generosity and helping. *Journal of Positive Psychology* 11, 5 (2016). <https://doi.org/10.1080/17439760.2015.1127992>

[70] Dawn Querstret, Linda Morison, Sophie Dickinson, Mark Cropley, and Mary John. 2020. Mindfulness-Based Stress Reduction and Mindfulness-Based Cognitive Therapy for Psychological Health and Well-Being in Nonclinical Samples: A Systematic Review and Meta-Analysis. *International Journal of Stress Management* (2020). <https://doi.org/10.1037/str0000165>

[71] Kyla Rankin, Sara E. Andrews, and Kate Sweeny. 2020. Awe-full uncertainty: Easing discomfort during waiting periods. *Journal of Positive Psychology* 15, 3 (2020). <https://doi.org/10.1080/17439760.2019.1615106>

[72] Leonard Reinecke and Wilhelm Hofmann. 2016. Slacking Off or Winding Down? An Experience Sampling Study on the Drivers and Consequences of Media Use for Recovery Versus Procrastination. *Human Communication Research* 42, 3 (2016), 441–461. <https://doi.org/10.1111/hcre.12082>

[73] Jo Ann Reynolds and Laurel Rabschutz. 2011. Studying for Exams Just Got More Relaxing-Animal-Assisted Activities at the University of Connecticut Library. *College and Undergraduate Libraries* 18, 4 (2011). <https://doi.org/10.1080/10691316.2011.624934>

[74] Martin J. Riedl, Gina M. Masullo, and Kelsey N. Whipple. 2020. The downsides of digital labor: Exploring the toll incivility takes on online comment moderators. *Computers in Human Behavior* 107 (2020). <https://doi.org/10.1016/j.chb.2020.106262>

[75] Sarah T Roberts. 2016. Commercial content moderation: Digital laborers' dirty work. In *The Intersectional Internet: Race, Sex, Class and Culture Online*, Safiya Umoja Noble and Brendesha Tynes (Eds.). NY: Peter Lang, New York, Chapter Commercial, 147–160. <https://doi.org/10.1007/s13398-014-0173-7.2>

[76] Minna Ruckenstein and Linda Lisa Maria Turunen. 2020. Re-humanizing the platform: Content moderators and the logic of care. *New Media & Society* 22, 6 (6 2020), 1026–1042. <https://doi.org/10.1177/1461444819875990>

[77] Mahdieh Sabery, Mansoureh Zagheri Tafreshi, Meimanat Hosseini, Jamileh Mohtashami, and Abbas Ebadi. 2017. Development and psychometric evaluation of the nurses' Compassion Fatigue Inventory. *Journal of Nursing Measurement* 25, 3 (2017). <https://doi.org/10.1891/1061-3749.25.3.E185>

[78] Kirk Schneider. 2017. The resurgence of awe in psychology: Promise, hope, and perils. *Humanistic Psychologist* 45, 2 (2017). <https://doi.org/10.1037/hum0000060>

[79] Joseph Seering, Tianmi Fang, Luca Damasco, Mianhong Cherie Chen, Likang Sun, and Geoff Kaufman. 2019. Designing user interface elements to improve the quality and civility of discourse in online commenting behaviors. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, 14. <https://doi.org/10.1145/3290605.3300836>

[80] Joseph Seering, Robert Kraut, and Laura Dabbish. 2017. Shaping pro and anti-social behavior on Twitch through moderation and example-setting. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing - CSCW '17*. 111–125. <https://doi.org/10.1145/2998181.2998277>

[81] Joseph Seering, Tony Wang, Jina Yoon, and Geoff Kaufman. 2019. Moderator engagement and community development in the age of algorithms. *New Media and Society* 21, 7 (2019), 1417–1443. <https://doi.org/10.1177/1461444818821316>

[82] Gary D. Sherman and Jonathan Haidt. 2011. Cuteness and disgust: The humanizing and dehumanizing effects of emotion. <https://doi.org/10.1177/1754073911402396>

[83] Gary D. Sherman, Jonathan Haidt, and James A. Coan. 2009. Viewing Cute Images Increases Behavioral Carefulness. *Emotion* 9, 2 (2009). <https://doi.org/10.1037/a0014904>

[84] Gary D. Sherman, Jonathan Haidt, Ravi Iyer, and James A. Coan. 2013. Individual differences in the physical embodiment of care: Prosocially oriented women respond to cuteness by becoming more physically careful. *Emotion* 13, 1 (2013). <https://doi.org/10.1037/a0029259>

[85] Sarah A. Smith. 2014. Mindfulness-Based Stress Reduction: An Intervention to Enhance the Effectiveness of Nurses' Coping With Work-Related Stress. *International Journal of Nursing Knowledge* 25, 2 (6 2014), 119–130. <https://doi.org/10.1111/2047-3095.12025>

[86] Yeoungsuk Song and Ruth Lindquist. 2015. Effects of mindfulness-based stress reduction on depression, anxiety, stress and mindfulness in Korean nursing students. *Nurse Education Today* 35, 1 (1 2015), 86–90. <https://doi.org/10.1016/j.nedt.2014.11.011>

1016/j.nedt.2014.06.010

- [87] John J. Squiers, Kevin W. Lobdell, James I. Fann, and J. Michael DiMaio. 2017. Physician Burnout: Are We Treating the Symptoms Instead of the Disease? *Annals of Thoracic Surgery* 104, 4 (10 2017), 1117–1122. <https://doi.org/10.1016/j.athoracsur.2017.08.009>
- [88] Elizabeth A. Stanley, John M. Schaldach, Anastasia Kiyonaga, and Amishi P. Jha. 2011. Mindfulness-based Mind Fitness Training: A Case Study of a High-Stress Predeployment Military Cohort. *Cognitive and Behavioral Practice* 18, 4 (11 2011), 566–576. <https://doi.org/10.1016/j.cbpra.2010.08.002>
- [89] Miriah Steiger, Timir J Bharucha, Sukrit Venkatagiri, Martin J Riedl, and Matthew Lease. 2021. The Psychological Well-Being of Content Moderators. In *Proceedings of 2021 CHI Conference on Human Factors in Computing Systems - CHI '21*. ACM, 14. <https://doi.org/10.1145/3411764.3445092>
- [90] Kevin D. Stein, Paul B. Jacobsen, Chris M. Blanchard, and Christina Thors. 2004. Further validation of the multidimensional fatigue symptom inventory-short form. *Journal of Pain and Symptom Management* 27, 1 (2004). <https://doi.org/10.1016/j.jpainsympman.2003.06.003>
- [91] Hannah Strohmeier, Willem F. Scholte, and Alastair Ager. 2018. Factors associated with common mental health problems of humanitarian workers in South Sudan. *PLoS ONE* 13, 10 (2018). <https://doi.org/10.1371/journal.pone.0205333>
- [92] Courtney E. Sullivan, Amber Rose King, Joni Holdiness, Judith Durrell, Kristin K. Roberts, Christopher Spencer, Joshua Roberts, Susan W. Ogg, Meredith W. Moreland, Emily K. Browne, Carla Cartwright, Valerie Mc Laughlin Crabtree, Justin N. Baker, Mark Brown, April Sykes, and Belinda N. Mandrell. 2019. Reducing compassion fatigue in inpatient pediatric oncology nurses. *Oncology Nursing Forum* 46, 3 (2019). <https://doi.org/10.1188/19.ONF.338-347>
- [93] Reina Takamatsu. 2020. Measuring Affective Responses to Cuteness and Japanese kawaii as a Multidimensional Construct. *Current Psychology* 39, 4 (2020), 1362–1374. <https://doi.org/10.1007/s12144-018-9836-4>
- [94] Tourism Western Australia and University of Leeds. 2020. Study: Quokkas can be good for your health - First Stop Singapore. <https://www.firststopsingapore.com/en/study-quokkas-can-be-good-for-your-health/>
- [95] Piercarlo Valdesolo and Jesse Graham. 2014. Awe, Uncertainty, and Agency Detection. *Psychological Science* 25, 1 (2014). <https://doi.org/10.1177/0956797613501884>
- [96] Tingting Wang, Anirban Mukhopadhyay, and Vanessa M. Patrick. 2017. Getting consumers to recycle now! When and why cuteness appeals influence prosocial and sustainable behavior. *Journal of Public Policy and Marketing* 36, 2 (2017). <https://doi.org/10.1509/jppm.16.089>
- [97] Donghee Yvette Wohn. 2019. Volunteer moderators in twitch micro communities: How they get involved, the roles they play, and the emotional labor they experience. In *Conference on Human Factors in Computing Systems - Proceedings*. <https://doi.org/10.1145/3290605.3300390>
- [98] James Yandell. 1987. Wasting time-the something of nothing. *Psychological Perspectives* 18, 1 (1987). <https://doi.org/10.1080/00332928708408752>
- [99] Serap Yildirim and Funda Cavcav. 2020. The Compassion Fatigue-Short Scale for healthcare professionals: A Turkish study of validity and reliability. *Perspectives in Psychiatric Care* (2020). <https://doi.org/10.1111/ppc.12712>
- [100] Qiyuan Zhang and Judith Covey. 2014. Past and Future Implications of Near-Misses and Their Emotional Consequences. *Experimental Psychology* 61, 2 (10 2014), 118–126. <https://doi.org/10.1027/1618-3169/a000231>

Received April 2021; revised November 2021; accepted March 2022