

Content Warnings Reduce Aesthetic Appreciation of Visual Art

Payton J. Jones¹, Victoria M. E. Bridgland^{1, 2} and Benjamin W. Bellet¹

¹Department of Psychology, Harvard University

²College of Education, Psychology & Social Work, Flinders University

Content warnings are alerts about upcoming content that might be related to upsetting or traumatic experiences. Such warnings are increasingly used by artists and art curators around the world. Though the psychological literature on content warnings suggests they are typically functionally inert, warnings have not yet been studied in the context of art or aesthetics. In this preregistered, within-person, randomized controlled experiment, we showed diverse art pieces to 213 participants (six trials each). By random assignment, some art was prefaced with a content warning matching its specific content (e.g., “content warning: sexual assault” for Gérôme’s *Phryne before the Areopagus*). We found that content warnings decreased aesthetic appreciation (Cohen’s $d = -0.22$, Bayes factor = 54, $N = 1,278$). Content warnings also substantially increased negative emotional responses and decreased positive emotional responses (Cohen’s $d = 0.44$, Bayes factor = 9.6×10^9 , $N = 1,278$). Though we planned to test the effect of warnings on opting out of viewing art, we were surprised to find that none of the participants avoided viewing any of the art pieces regardless of whether they were prefaced with a warning.

Keywords: content warning, trigger warning, art, aesthetic appreciation

Content warnings (also called trigger warnings or content notes) are alerts about upcoming content that might be related to upsetting or traumatic experiences. Most frequently, the intent of these warnings is to help people emotionally brace themselves or completely avoid distressing content (e.g., Gust, 2016). In other cases, the intent relates to improving learning outcomes in an academic setting (e.g., Bentley, 2017). Accordingly, existing research has focused on the emotional, clinical, and pedagogical effects of warnings. This literature consistently indicates that warnings increase anticipatory anxiety, but have little impact on avoidance behavior, emotional reactions to content, or learning (see Bridgland et al., 2022 for a meta-analysis; see also Bridgland et al., 2019; Sanson et al., 2019). Although content warnings first became common in online forums and academic settings, their use has expanded to many other domains, including the world of art (e.g., Engle, 2018; Tobin, 2021). Content warnings are frequently used wherever art is found, from museums and galleries to art forums and social media. We sought to investigate how content warnings affect the aesthetic appreciation (AA) of visual artwork.

AA refers to the psychological processes associated with admiring beauty. AA encompasses not only visual processing, but also beauty found in representational or semantic characteristics and across non-visual modalities such as music and literature. Psychologists who study AA argue for a meaningful role of active cognitive processing and interpretation (e.g., Leder et al., 2004). Because AA relies at

least partially on semantic interpretation, humans frequently find beauty in the “how” or “why” of art (i.e., the process required to produce the art or its significance in the larger cultural context), even if the “what” (i.e., content of the art) is unremarkable or even distasteful (e.g., Marcel Duchamp’s 1964 “Fountain”; Camfield, 1990).

As is clear to anyone acquainted with the arts, there is a great deal of individual variation in AA. Despite this, researchers have demonstrated that AA can be fruitfully measured with high psychometric reliability and studied in a scientific context (Blijlevens et al., 2017; Faerber et al., 2010). Given variation across individuals, within-person experiments may be particularly appropriate. Though most conceptualizations and measurement scales include multiple separate aspects of AA (e.g., arousal, interestingness, innovativeness), these aspects tend to share moderate-to-strong intercorrelations (Blijlevens et al., 2017; Faerber et al., 2010), and at least some research has suggested a primacy of the beautiful–ugly dimension (Jacobsen et al., 2004).

Do content warnings affect the AA of art? Past research does not suggest a clear prediction. Considering the trivial effects of warnings on affective reactions to other types of content, we might predict that warnings would have no effect on AA (Bridgland et al., 2022).

Yet AA research suggests that cognitive interpretation of the content importantly interacts with affective reactions (Leder & Nadal, 2014). As such, it’s possible that warnings might result in semantic priming effects. Generally, when a semantic concept is presented, similar concepts presented soon thereafter are recognized more readily than unrelated concepts (e.g., when the word “nurse” is on a screen, the word “doctor” can be identified more quickly than the word “bread”; Henik & Carr, 2002). Past studies suggest this can extend to AA: Faerber et al. (2010) found that semantic primes sometimes influenced aesthetic judgments of visual designs. Content warnings could presumably cause a

Payton J. Jones  <https://orcid.org/0000-0001-6513-8498>

Correspondence concerning this article should be addressed to Payton J. Jones, Harvard University, 1240 William James Hall, 33 Kirkland Street, Cambridge, MA 02138, United States Email: paytonjones@gmail.com

similar effect. For example, prefacing Jean-Léon Gérôme's *Phyrne before the Areopagus* with "content warning: sexual assault" might cause participants to more readily notice the subject's visible shame rather than notice the piece's use of color or style. This could presumably result in lower perceptions of the painting's aesthetics—for instance, its valence or attractiveness.

Alternatively, warnings might serve as a value signal, allowing participants to more deeply trust the institution displaying the art (Bruce & Roberts, 2020). That said, it's unclear how institutional trust would impact appraisals of the artwork itself. Warnings might also enhance physiological arousal (Bruce et al., 2021), which may be related to at least some facets of AA (i.e., arousal, boringness). Without an empirical investigation, it is difficult to predict precisely what effect warnings might have on AA, if any.

In addition to testing for an overall effect on AA, we were interested in the effect of content warnings on each individual facet of AA. For example, warnings might decrease perceived attractiveness but increase perceived interestingness. We were also interested in the effects studied by past research on warnings, namely, effects on emotional response and stimulus avoidance (Bridgland et al., 2022). These are the most common goals of content warnings. Though past research has typically revealed null effects for emotional response and avoidance, they are relevant to include in this study as a replication of past studies and an extension to visual art as a stimulus. Given the politicization of content warnings in the media, we were also curious whether participants' reported political orientation would moderate effects.

To summarize, our primary preregistered research question was as follows:

1. Does the presence of a content warning impact overall AA of visual art?
2. Does the presence of a content warning impact each individual factor of AA (attractiveness, arousal, interestingness, valence, boredom, and innovativeness)?
3. Does the presence of a content warning impact overall negative emotion in response to visual art?
4. Does the presence of a content warning impact each individual facet of emotion (sad, happy, afraid, anxious, angry, content, excited, and disgusted)?
5. Does the presence of a content warning impact how often participants opt out of (i.e., avoid) viewing visual art?
6. Does the presence of a content warning interact with political orientation with respect to effects on AA or negative emotion?

To add context and insight to the primary research question, we preregistered the following secondary research questions:

Method

Study methods and analyses were preregistered on the Open Science Framework (OSF; <https://osf.io/myjfa>). Open data and analysis code are available via the corresponding OSF project (<https://osf.io/wcfzt>). Any departures from the preregistration are noted as such. We have reported all measures, conditions, and data exclusions. All procedures were approved by the Harvard Institutional Review Board.

Experimental Design

To test whether content warnings impacted AA of visual art, we designed a within-person randomized controlled trial. Participants rated stimuli that were randomly assigned to be prefaced simply with the artist's name and year (e.g., "Gerome, 1861"), or to be prefaced with the same information plus a content warning ("Gerome, 1861 | content warning: sexual assault"). All stimuli in the study contained potentially triggering content (e.g., sexual assault, war, violence), and thus had a relevant content warning that could potentially be assigned to it.

Participants were told they were taking part in a study on "Reactions to Classical and Modern Art." We concealed the true purpose of the study until a debriefing form at the end of the study. Participants were instructed that they would be rating their reactions to art pieces, each of which should be viewed in full screen for at least 10 s, though they were also told via the consent form that "all questions we ask are optional," and that they could proceed to the next screen after waiting 10 s. In the main phase of the study, participants were sequentially presented with the randomly assigned information about the art (e.g., "Gerome, 1861" or "Gerome, 1861 | content warning: sexual assault") and a gray prompt box with the text: "Click to view image." Participants were unable to advance until 10 s had passed but were not required to click the prompt box (i.e., they were allowed to opt out of viewing an image). After viewing each stimulus, they completed brief questionnaires regarding AA of the art and their emotional states.

Each participant was provided with six art pieces randomly chosen from a stimulus bank containing 13 total images. The randomization strategy ensured that for each participant, at least one stimulus was prefaced with a content warning and at least one stimulus was not. Allowing the number of total warnings to vary, as opposed to a more restricted randomization strategy, helps simplify causal inference and protects against some inferential biases (Lachin et al., 1988). After this main phase of the study, participants completed remaining questionnaires on demographic information, completed validity and attention checks, and reported which images they had viewed or avoided. For validity checks and avoidance questions, participants were informed that their responses would not affect study approval or payment.

Participants

We recruited adults in the United States from Amazon Mechanical Turk using the CloudResearch platform (Litman et al., 2017). Participants were recruited between October 27 and 28, 2022. In addition to our own set of validity checks described in detail below, we used CloudResearch's suite of features to limit bots and server farmers. We followed best practices for ensuring high-quality responses in a crowdsourcing environment, including several validity checks and collection of demographic data (see Screening Variables section; Aguinis et al., 2021; Moeck et al., 2022).

To determine the necessary sample size, we conducted a power analysis simulation in R (R Core Team, 2022).¹ In this simulation,

¹ This power analysis was mostly for convenience in selecting a sample size and number of trials; power is fundamentally a frequentist concept. Whereas our preregistered analysis used Bayes Factors, this power analysis used *p* values.

we created hundreds of sets of artificial data using various combinations of trial numbers, sample sizes, and effect sizes. We then automated a linear mixed-effects model analysis of the artificial data to identify how often our model would detect the effect (i.e., the true positive rate). This indicated that a sample of 200 participants rating six stimuli each was sufficient to detect a small effect size (Cohen's $d = 0.2$) with a power of 0.92. As preregistered, we collected participant responses in batches of 20 until we reached a sample of at least 200 valid participants after exclusions. This resulted in a total sample of 213 participants.

Measures

Aesthetic Appreciation

AA was the primary dependent variable. Following Faerber et al. (2010), we measured AA using six key variables identified from the literature (attractiveness, arousal, interestingness, valence, boredom, and innovativeness). Each variable was assessed on a 7-point Likert scale (e.g., "The painting is interesting," 1 = *strongly disagree*, 7 = *strongly agree*). To measure overall AA, items were summed (boredom was reverse coded).

Emotions

After each image, participants rated their emotional state by using slider bars ranging from 0 (*not at all*) to 100 (*very much*) on eight emotions: sad, happy, afraid, anxious, angry, content, excited, and disgusted. To measure negative emotional response in the aggregate, an emotion sum score was created by summing all items (with happy, content, and excited reverse coded).

Avoidance

Avoidance was measured by first asking participants if they avoided clicking any photos ("Did you click to view every painting shown in this study? [This question will not impact payment]"). If participants answered no, they were asked to check all reasons they did not view the photos ("I wanted to finish the study as quickly as possible," "I was anxious or afraid about the content of the art," "Other: ___"), and were asked to indicate specifically which paintings they did not view (descriptions and warnings from earlier in the study were repeated here to aid recall).

Demographic Variables

Participants were asked to report their gender, race, ethnicity, religiosity, political orientation, parents' educational attainment, student status, and age.

Screening Variables

A basic reCAPTCHA was used at the beginning of the study to exclude bots. We also used a five-question English proficiency test at the beginning of the survey (e.g., "Fill in the blank: You should not have a dog if you are not ___ to look after it" [prepared, adapted, arranged, decided]). Participants who answered fewer than four out of five questions correctly were immediately excluded from the study. An attention check item near the end of the survey was used to verify that participants were not randomly responding (a prompt that read "Please select each of the images you viewed

earlier [if any]," followed by a list of short descriptions of images). Participants failed this attention check and were excluded from the analysis if they selected an image not included in the available stimuli (i.e., "A large black dog"). A general validity screener was included at the end of the study ("Is there any reason you think that your data should not be used [this will not impact payment]?"). Participants whose responses were determined to be invalid based on this question (e.g., they admitted to randomly responding) were excluded from the analysis.

Analyses

For each research question, we computed a linear mixed-effects model with the presence of a content warning as the independent variable (fixed effect). As noted above, this was a within-person experimental design in which each participant viewed a subset of stimuli from a stimulus bank. A stimulus identifier and participant identifier were therefore included as random intercepts. The inclusion of these random intercepts functionally controls for any meaningful variation between the different stimuli (i.e., properties that are specific to a stimulus and not shared across all other stimuli, such as being highly recognizable) or across the different participants (such as being an art aficionado).

We then computed a Bayes factor (BF) comparing the probability of the null hypothesis ($b = 0$) to a nonspecific alternative hypothesis ($b! = 0$). More specifically, we computed the BF by using the *BayesFactor* R package (Morey & Rouder, 2018) to compare a full model ($y \sim \text{content_warning} + (1 \mid \text{participant_id}) + (1 \mid \text{stimulus_id})$) to a model lacking the content warning as a fixed effect ($y \sim (1 \mid \text{participant_id}) + (1 \mid \text{stimulus_id})$). We used the default prior distributions specified in the *BayesFactor* package (for specifics, see the *regressionBF* documentation, p. 48, Morey & Rouder, 2018). To perform a sensitivity check on the default prior distributions, we tested against an alternative ultrawide prior distribution, using an r -scale value of 1 rather than the original default of $\sqrt{2}/4$. This change did not meaningfully alter interpretation for any of the results. All values from this sensitivity analysis are available in the online supplemental materials (<https://osf.io/wcfzt/>).

We used thresholds of $1/3$ or 3 as a criterion for "substantial evidence" relative to the null or alternative hypothesis, respectively (Wetzels et al., 2011). Though we employed thresholds to minimize researcher ambiguity in the interpretation of the effects, we are at pains to emphasize that BFs are a continuous measure of relative evidence across two competing models. That is, a BF of 3.1 should not be interpreted with the same gravity as a BF of 3.1×10^9 .

Results

Screening

As preregistered, we initially recruited 200 participants and then collected an additional batch of 20 until at least 200 valid responses were identified. Participants who failed the English proficiency test were automatically excluded and replaced by another participant; there were 27 such responses found via Qualtrics. One participant was excluded due to submitting an incomplete survey. This left 219 recruited participants who completed the full survey and were eligible for payment. Four participants were excluded via the validity screener; three due to nonsense responses in the text field (e.g., "I

trust always”), and one due to indicating that their responses were invalid. Two participants were excluded due to reporting technical problems with the survey (e.g., images failed to load). This left 213 valid responses included in the final analysis.

Participant Demographics

Participants reported their sex as male (61%), female (39%), or other (0%) with a mean age of 39.4 ($SD = 11.7$). A majority of participants were nonreligious (54%). In terms of race, participants were Caucasian (75%), Asian/Pacific Islander (8%), Black/African American (7.5%), Hispanic/Latino (7%), or Multiracial (<1%).

Politically, participants leaned left-of-center, with an average score of 2.48 on a scale from 1 (*very liberal*) to 5 (*very conservative*). A small minority of participants reported being currently enrolled as undergraduate students (3%). Considering the highest education achieved by a parent, a bachelor’s degree was most common (40%), followed by some college, but no degree (22%), high school diploma or equivalent (GED; 15%), associate degree (14%), master’s degree (5%), graduate or professional degree (PhD, MD, JD; 4%), or less than a high school education (<1%).

Aesthetic Appreciation

Our primary research question was whether content warnings impacted AA of visual art. We found that content warnings decreased AA ($d = -0.22$, $BF = 54.70$). Full regression tables and sensitivity tests of BFs can be found in the online supplemental materials (<https://osf.io/wcfzt>).

Figure 1 displays side-by-side violin plots of the ratings (sum score of AA items) in each condition. The width of each violin reflects how many participants scored in that range (e.g., in the violins of Figure 1, more participants in the no warning condition had

AA scores above 35 than participants in the warning condition). A horizontal line in each violin represents the mean of the distribution. It should be noted that the violin plots simply display the aggregate distributions for convenience, whereas our conclusions and associated statistics are derived from linear mixed-effects models as specified above.

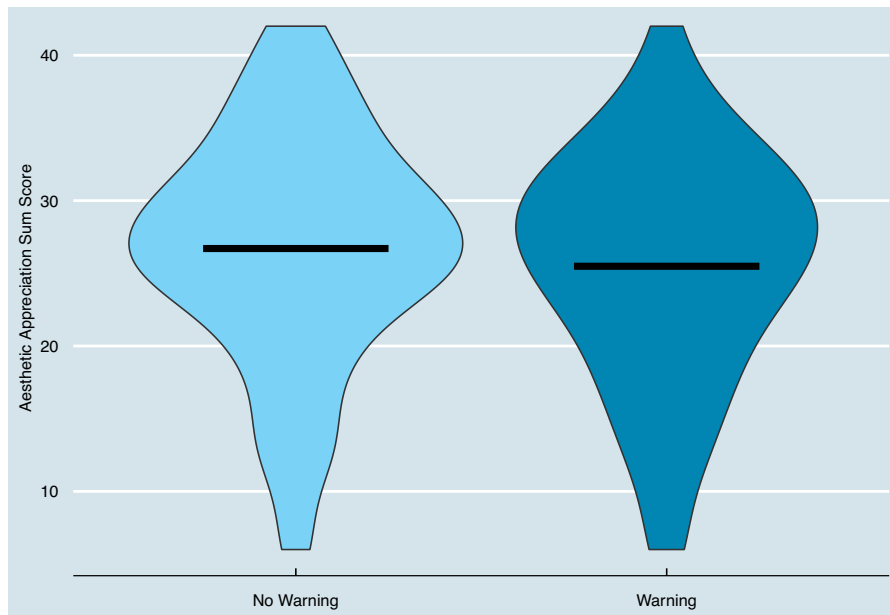
We also analyzed each facet of AA individually (Figure 2). Overall, results across the various aspects of AA were directionally consistent with the negative effect on the sum score. That is, content warnings did not enhance any aspect of AA. Adding a content warning substantially decreased the extent to which participants found a painting attractive ($d = -0.26$, $BF = 419.34$) or pleasant (valence; $d = -0.44$, $BF = 8.35 \times 10^9$). There was ambiguous evidence ($1/3 < BF < 3$) for decreases in arousal ($d = -0.14$, $BF = 0.76$) and interestingness ($d = -0.13$, $BF = 0.57$). BFs indicated relatively greater evidence for a null effect of content warnings on boringness ($d = 0.02$, $BF = 0.07$) and innovativeness ($d = -0.036$, $BF = 0.08$). In summary, it seems that decreases in AA were primarily driven by changes in the participants’ perception of aesthetic attractiveness and valence, whereas other aspects of AA saw relatively little movement.

Emotion

Next, we analyzed the impact of content warnings on emotional responses to visual art. Most prior research on content warnings has examined their effect on emotions, with most studies suggesting a null effect (see Bridgland et al., 2022 for a meta-analysis). Consequently, we expected to find similar results (i.e., a small or null effect). Contrary to this expectation, the presence of a content warning prior to an art piece substantially increased negative emotion ($d = 0.44$, $BF = 9.59 \times 10^9$; Figure 3).

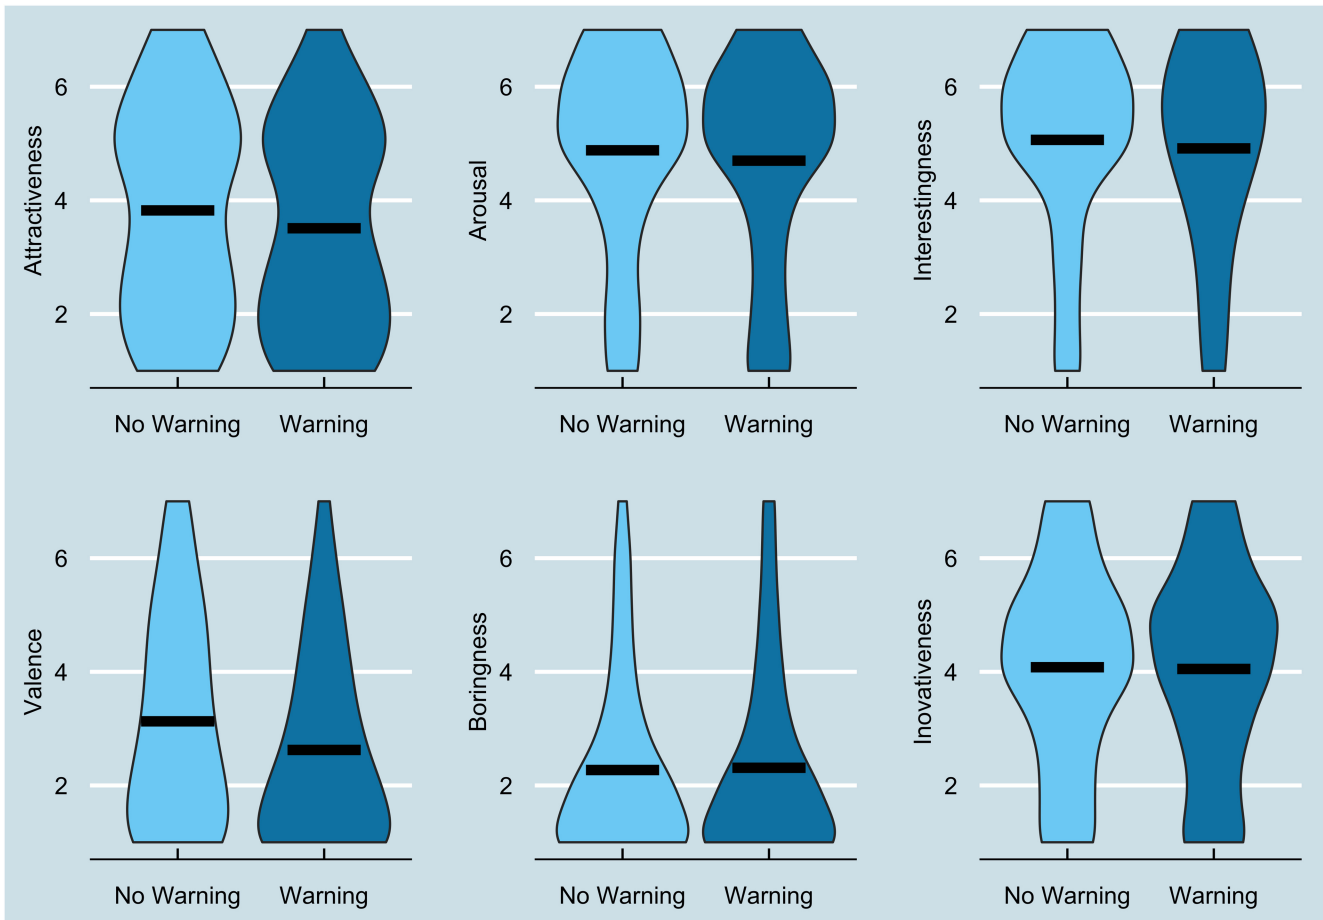
As with AA, the individual facets of emotion remained consistent with the overall trend (Figure 4). However, in the case of emotion, all

Figure 1
The Effect of Content Warnings on Aesthetic Appreciation



Note. See the online article for the color version of this figure.

Figure 2
The Effect of Content Warnings on Specific Facets of Aesthetic Appreciation



Note. See the online article for the color version of this figure.

individual facets showed substantial changes, with small-to-medium effect sizes. There were substantial increases in all negative emotions: sad ($d = 0.33$, $BF = 1.24 \times 10^5$), afraid ($d = 0.24$, $BF = 111.58$), anxious ($d = 0.22$, $BF = 29.55$), angry ($d = 0.35$, $BF = 7.03 \times 10^5$) and disgusted ($d = 0.40$, $BF = 1.25 \times 10^8$). Likewise, there were substantial decreases in all positive emotions: happy ($d = -0.39$, $BF = 4.42 \times 10^7$), content ($d = -0.29$, $BF = 5,114.50$), and excited ($d = -0.23$, $BF = 64.88$). Interestingly, this pattern suggests that content warnings did not uniformly impact emotional *arousal*, but rather impacted emotional *valence* specifically (e.g., anxiety increased but excitement and happiness decreased).

Avoidance

Our next preregistered research question was whether content warnings would impact how often participants would opt out of viewing visual art. In the study instructions, participants were asked to view each piece of art for 10 s. However, clicking the image was not required to advance through the study. For each stimulus, participants would only see the image after clicking a prompt (“Click to view image”). At the end of the study, we asked participants whether they had skipped any images, and if they had, to


identify each of the images they had skipped. Participants were informed that these questions would not affect payment.

We planned to fit a linear mixed-effects model with stimulus avoidance as the dependent variable. However, we were surprised to find that fitting this model was impossible, as no participants had avoided any stimuli. That is, out of the 213 participants and 1,278 total trials, there were zero cases in which participants opted out of viewing content. This is consistent with prior research showing that participants rarely opt out from viewing negative stimuli, including stimuli preceded by a trigger warning (e.g., Bridgland & Takarangi, 2022).

Interaction With Political Orientation

Our final preregistered research question was whether there was an interaction between content warnings and political orientation. Our results suggested a lack of evidence for an interaction. Content warnings were related to decreases in AA ($BF = 54.70$) and increases in negative emotion ($BF = 9.59 \times 10^9$), but political orientation did not moderate these effects. For AA, evidence favored a null hypothesis ($BF = 0.10$). For negative emotion, evidence for an interaction was ambiguous ($BF = 0.38$).

Figure 3
The Effect of Content Warnings on Negative Emotional Response to Art



Note. See the online article for the color version of this figure.

Discussion

We tested whether content warnings affect how individuals respond to visual art. We found that content warnings reduced the AA of art. No aspect of AA (attractiveness, arousal, interestingness, valence, boredom, innovativeness) saw improvements with a content warning. The drops in AA were primarily driven by decreases in perceptions of *attractiveness* and *valence*. Relatedly, we found that content warnings substantially increased all negative emotions (and decreased all positive emotions) in response to art. These effects were not moderated by the political orientation of participants. Whether or not warnings were given, participants did not opt out of viewing the art.

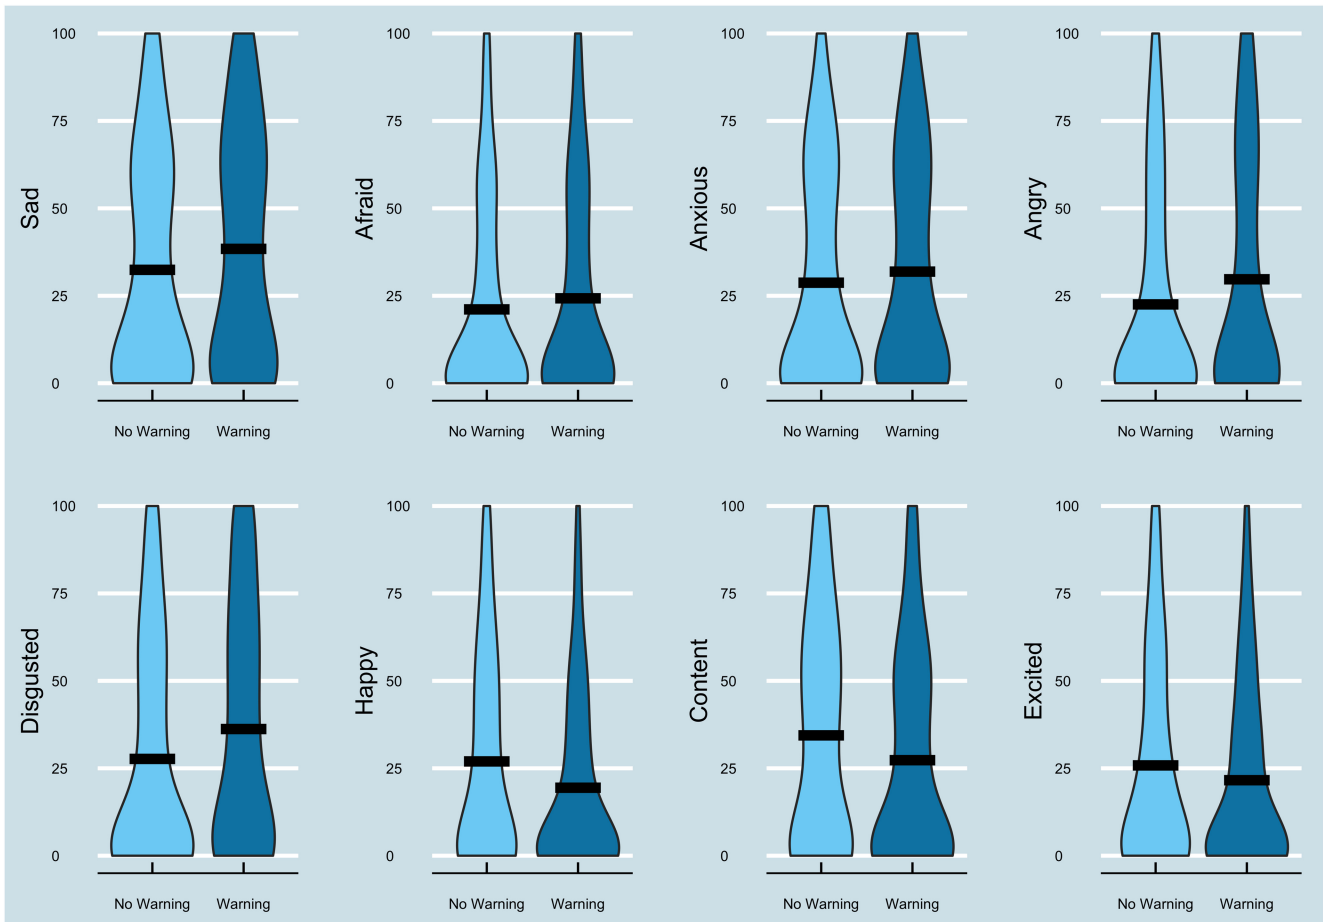
Why might content warnings reduce AA and increase negative emotion? Though this is the first study to examine AA, past research on content warnings has found that they generally have little if any impact on emotional response (Bridgland et al., 2022). The substantial impacts on emotional response in this study are therefore surprising: multiple studies have found that warnings have small (e.g., Bellet et al., 2020) or null (Bridgland et al., 2019; Sanson et al., 2019) effects on participant emotional responses (aside from a brief period of uncomfortable anticipation; Bridgland et al., 2022). What makes art unique? One possibility is that visual art is more ambiguous than stimuli used in past studies of content warnings. Whereas most literature passages, pictures, or videos might only have a few possible interpretations, visual art can take on a much larger variety of meanings. Because of this ambiguity, interpretation may be more easily influenced by various types of semantic priming, including content warnings (Faerber et al., 2010). Warnings may also narrow attention to the literal content of the art (i.e., a sexual assault) to the exclusion of factors such as the technical mastery of the artist or the artistic composition; this would be an interesting hypothesis to address in future research. Indeed, psychological

priming effects in general have been the subject of much controversy regarding replicability and reliability. This may be in part because their effect is limited to very specific scenarios (see Cesario, 2014; Molden, 2014). Content warnings may have trivially small effects in most cases, but nevertheless cause negative effects in the specific case of interpreting a highly ambiguous stimulus.

Another possibility relates to the subject of measurement. In past studies, measurement scales almost universally focused on the participant (“how do *you* feel right now”). In contrast, the present study focused more heavily on participants’ ratings of the content itself (“how attractive is the *painting*”). One counterpoint to this hypothesis is that the emotion scales used in the study did focus on the participants, and effects were still found. However, the context of the study may matter—participants were told that the purpose of the study was to understand ratings of art, and may have focused more on rating the content than accurately reporting their internal state.

Our study has several limitations that could be addressed in future research. Given that our results were somewhat surprising, a direct replication would be welcome. In addition, it would be useful for future studies to examine the role of content ambiguity, perhaps by first establishing the ambiguity of various stimuli in a pilot study. Our study tested the effect of adding a content warning to the artists’ last name and year; future studies could investigate warnings against no information or alternatively, against a more complete description of the art. Though past studies have indicated very little avoidance behavior following content warnings (e.g., Bridgland & Takarangi, 2022; Sanson et al., 2019), we were still surprised that not a single participant avoided any art in this study. Though our consent form mentioned that “all questions we ask are optional” and the study instructions stated that “after 10 s have passed, you will be able to proceed to the next screen,” it is possible that participants did not fully understand this option.

Figure 4
The Effect of Content Warnings on Specific Facets of Emotional Response to Art



Note. See the online article for the color version of this figure.

Conclusion

Many artists and art exhibits around the world currently furnish content warnings. The intention behind a content warning is noble: those providing warnings hope to provide a comfortable environment for those with a history of trauma, enable choice, and contextualize potentially provocative art.

A growing body of literature suggests that despite noble intentions, content warnings are ineffective or counterproductive when it comes to managing posttraumatic stress disorder (Bridgland et al., 2022; Jones et al., 2020). The present study adds to this literature by answering adjacent questions specific to art and aesthetics. The results suggest yet another reason for artists to forego content warnings: they reduce the degree to which individuals are able to process art in a fruitful and enriching way. They also decrease positive emotion and increase negative emotion in response to that art. The combination of these factors may be greater than the sum of their parts. Negative emotion is not inherently harmful, and can be a powerful tool in art. But if content warnings skew perceptions toward a lack of artistic beauty, intensifying of negative emotion, and sapping of positive emotion, it's hard to imagine that they enhance the aesthetic, pedagogical, or expressive goals of art. Finally, our results align with

past research to suggest that warnings do not encourage or meaningfully enable opting out of viewing content (which may be an ill-conceived goal in the first place; see Foa & Kozak, 1986; Jones et al., 2020).

If not by using content warnings, how might artists contextualize provocative, controversial, or potentially upsetting art pieces? One option is to do nothing and simply place greater trust and responsibility in the hands of the viewer to respond to art on their own terms, including potentially negative reactions. In the case that additional contextualization is needed or desired, artists and curators should act with caution. The results of this study suggest that attempting to frontload contextualization via content warnings will backfire. It is possible that useful contextualization could be added *after* the individual is given the chance to first process, interpret, and react to the art on their own, though it remains to future research to test this idea.

In summary, we conducted a randomized controlled trial to measure the causal impact of content warnings for visual art. We found that content warnings decreased AA, decreased positive emotion, and increased negative emotion. They made no difference in whether participants avoided viewing art. Considering these results in combination with a growing scientific literature, we suggest that artists and art curators suspend their use of content warnings.

References

- Aguinis, H., Villamor, I., & Ramani, R. S. (2021). MTurk research: Review and recommendations. *Journal of Management*, 47(4), 823–837. <https://doi.org/10.1177/0149206320969787>
- Bellet, B. W., Jones, P. J., Meyersburg, C. A., Brennehan, M. M., Morehead, K. E., & McNally, R. J. (2020). Trigger warnings and resilience in college students: A preregistered replication and extension. *Journal of Experimental Psychology: Applied*, 26(4), 717–723. <https://doi.org/10.1037/xap0000270>
- Bentley, M. (2017). Trigger warnings and the student experience. *Politics*, 37(4), 470–485. <https://doi.org/10.1177/0263395716684526>
- Blijlevens, J., Thurgood, C., Hekkert, P., Chen, L. L., Leder, H., & Whitfield, T. W. (2017). The Aesthetic Pleasure in Design Scale: The development of a scale to measure aesthetic pleasure for designed artifacts. *Psychology of Aesthetics, Creativity, and the Arts*, 11(1), 86–98. <https://doi.org/10.1037/aca0000098>
- Bridgland, V. M. E., Green, D. M., Oulton, J. M., & Takarangi, M. K. (2019). Expecting the worst: Investigating the effects of trigger warnings on reactions to ambiguously themed photos. *Journal of Experimental Psychology: Applied*, 25(4), 602–617. <https://doi.org/10.1037/xap0000215>
- Bridgland, V. M. E., Jones, P. J., & Bellet, B. W. (2022). A meta-analysis of the effects of trigger warnings, content warnings, and content notes. OSF Preprints. <https://doi.org/10.31219/osf.io/qav9m>
- Bridgland, V. M. E., & Takarangi, M. K. T. (2022). Something distressing this way comes: The effects of trigger warnings on avoidance behaviors in an analogue trauma task. *Behavior Therapy*, 53(3), 414–427. <https://doi.org/10.1016/j.beth.2021.10.005>
- Bruce, M. J., & Roberts, D. (2020). Trigger warnings for abuse impact reading comprehension in students with histories of abuse. *College Student Journal*, 54(2), 157–168. <https://www.ingentaconnect.com/content/prin/cs/2020/00000054/00000002/art00002>
- Bruce, M. J., Stasik-O'Brien, S. M., & Hoffmann, H. (2021). Students' psychophysiological reactivity to trigger warnings. *Current Psychology*. Advance online publication. <https://doi.org/10.1007/s12144-021-01895-1>
- Camfield, W. A. (1990). Marcel Duchamp's Fountain: Its history and aesthetics in the context of 1917. In R. E. Kuenzli & F. M. Naumann (Eds.), *Marcel Duchamp: Artist of the century* (64–94). MIT Press.
- Cesario, J. (2014). Priming, replication, and the hardest science. *Perspectives on Psychological Science*, 9(1), 40–48. <https://doi.org/10.1177/1745691613513470>
- Engle, J. (2018). Should art come with trigger warnings? *The New York Times*. Retrieved November 14, 2022, from <https://www.nytimes.com/2018/11/28/learning/should-art-come-with-trigger-warnings.html>
- Faerber, S. J., Leder, H., Gerger, G., & Carbon, C. C. (2010). Priming semantic concepts affects the dynamics of aesthetic appreciation. *Acta Psychologica*, 135(2), 191–200. <https://doi.org/10.1016/j.actpsy.2010.06.006>
- Foa, E. B., & Kozak, M. J. (1986). Emotional processing of fear: Exposure to corrective information. *Psychological Bulletin*, 99(1), 20–35. <https://doi.org/10.1037/0033-2909.99.1.20>
- Gust, O. (2016). I use trigger warnings—But I'm not mollycoddling my students. *The Guardian*. <https://www.theguardian.com/higher-education-network/2016/jun/14/i-use-trigger-warnings-but-im-not-mollycoddling-my-students>
- Henik, A., & Carr, T. H. (2002). Inhibition. In V. S. Ramachandran (Ed.), *Encyclopedia of the human brain* (pp. 571–585). Academic Press. <https://doi.org/10.1016/B0-12-227210-2/00177-1>
- Jacobsen, T., Buchta, K., Köhler, M., & Schröger, E. (2004). The primacy of beauty in judging the aesthetics of objects. *Psychological Reports*, 94, 1253–1260. <https://doi.org/https://doi.org/10.2466/pr0.94.3c.1253-1260>
- Jones, P. J., Bellet, B. W., & McNally, R. J. (2020). Helping or harming? The effect of trigger warnings on individuals with trauma histories. *Clinical Psychological Science*, 8(5), 905–917. <https://doi.org/10.1177/2167702620921341>
- Lachin, J. M., Matts, J. P., & Wei, L. J. (1988). Randomization in clinical trials: Conclusions and recommendations. *Controlled Clinical Trials*, 9(4), 365–374. [https://doi.org/10.1016/0197-2456\(88\)90049-9](https://doi.org/10.1016/0197-2456(88)90049-9)
- Leder, H., Belke, B., Oeberst, A., & Augustin, D. (2004). A model of aesthetic appreciation and aesthetic judgments. *British Journal of Psychology*, 95(4), 489–508. <https://doi.org/10.1348/0007126042369811>
- Leder, H., & Nadal, M. (2014). Ten years of a model of aesthetic appreciation and aesthetic judgments: The aesthetic episode—Developments and challenges in empirical aesthetics. *British Journal of Psychology*, 105(4), 443–464. <https://doi.org/10.1111/bjop.12084>
- Litman, L., Robinson, J., & Abberbock, T. (2017). Turkprime.com: A versatile crowdsourcing data acquisition platform for the behavioral sciences. *Behavior Research Methods*, 49(2), 433–442. <https://doi.org/10.3758/s13428-016-0727-z>
- Moeck, E. K., Bridgland, V. M. E., & Takarangi, M. K. T. (2022). Food for thought: Commentary on Burnette et al. (2021) “Concerns and recommendations for using Amazon MTurk for eating disorder research”. *International Journal of Eating Disorders*, 55(2), 282–284. <https://doi.org/10.1002/eat.23671>
- Molden, D. C. (2014). Understanding priming effects in social psychology: An overview and integration. *Social Cognition*, 32(Suppl), 243–249. <https://doi.org/10.1521/soco.2014.32.suppl.243>
- Morey, R. D., & Rouder, J. N. (2018). *BayesFactor: Computation of Bayes factors for common designs* (Version 0.9.12-4.4) [Computer software]. <https://CRAN.R-project.org/package=BayesFactor>
- R Core Team. (2022). *R: A language and environment for statistical computing* [Computer software]. R Foundation for Statistical Computing. <http://www.R-project.org/>
- Sanson, M., Strange, D., & Garry, M. (2019). Trigger warnings are trivially helpful at reducing negative affect, intrusive thoughts, and avoidance. *Clinical Psychological Science*, 7(4), 778–793. <https://doi.org/10.1177/2167702619827018>
- Tobin, J. S. (2021). Putting trigger warnings on great works misses the point of art. *The Federalist*. Retrieved November 14, 2022, from <https://thefederalist.com/2021/09/02/putting-trigger-warnings-on-great-masterpieces-misses-the-point-of-art/>
- Wetzels, R., Matzke, D., Lee, M. D., Rouder, J. N., Iverson, G. J., & Wagenmakers, E. J. (2011). Statistical evidence in experimental psychology: An empirical comparison using 855 *t* tests. *Perspectives on Psychological Science*, 6(3), 291–298. <https://doi.org/10.1177/1745691611406923>

Received November 18, 2022

Revision received February 25, 2023

Accepted March 7, 2023 ■