Still too good to be true: Reply to Bushman et al. (2022)

Joseph Hilgard

Author note

Correspondence concerning this article should be addressed to Joseph Hilgard at e-mail jhilgard@gmail.com.

Abstract

In this reply to Bushman et al. (2022), I explain why differences in method and sample demographics are not likely to explain why the obvious effect in Hilgard (2021) is smaller than the indirect effect in Hasan et al (2013). Additionally, Bushman and colleagues claim that the large effects can be explained by accumulation of effects across days. Comparison with longitudinal experimental studies, cross-sectional observational studies, and longitudinal observational studies indicate that this accumulation of effects across three days cannot explain the unparalleled size of effects from Hasan et al. (2013). Accumulation of effects also cannot explain the incredibly large effects observed in *single-session* studies published by Hasan and colleagues. Better evidence is needed for the claim that other credible studies have reported even larger media violence effects. The data reported in Hasan et al. (2013) are still too good to be true.

Keywords: Maximal positive controls, effect size, violent video games, hostile expectation bias, story completion task

Word count: 3,236

The two studies at hand, Hasan et al. (2013) and Hilgard (2021), each examine influences on the expected aggressiveness of a character in a story. Hasan et al. (2013) purports to show that one hour of violent video games causes pronounced changes in the expectations participants have about a generic person's behavior. Hilgard (2021) purports to estimate the largest plausible effect size on this measure by measuring participants' expectations about the explicitly violent or non-violent video game character. Because the effect sizes reported in Hasan et al. (2013) exceed that of this "maximal positive control", Hilgard (2021) suggests they are too large to plausibly represent the true effect.

In Hilgard (2021), the participants have some knowledge about the story character, having watched a two minute video in which the character draws a house, races a 4-wheeler, or shoots twenty people to death. Unsurprisingly, participants expect the story character to behave more aggressively when they saw the character commit mass murder than when they saw the character race a 4-wheeler ($d = 1.71$, 95% CI [1.24, 2.18]) or draw a house ($d = 2.51$, 95% CI [1.97, 3.04]). This is a *direct* influence on perceptions of the story character and represents the easily observed psychological phenomenon that humans form beliefs about other humans based on those other humans' actions.

In Hasan et al. (2013), the participants have no knowledge of the story character. Their expectations of the story character's thoughts, behaviors, and feelings are influenced by an unrelated task in which the participants played a violent or nonviolent video game for one cumulative hour over three days. This indicates an *incidental* influence on perceptions of the story character through some priming of aggressive schema or some change in worldview that is then projected onto the story character.

It would be remarkable if the effect of this incidental influence was at all similar in size to the direct influence of observed behavior. However, the effect of the incidental influence does not only match the direct influence in size — it exceeds it by nearly a factor of two ($d = 3.34$, 95% CI [2.61, 4.06])¹. According to these results, an observer's impressions of a target person are less sensitive to whether the observer has seen that target person shooting up a strip club and more sensitive to the video game the observer has played for one hour.

Reply to authors' noted seven differences

In their comment, Bushman et al. (2022) defend the plausibility of the result reported in Hasan et al. (2013). They list seven differences between the methods used in Hilgard (2021) and Hasan et al. (2013): 1) participants watched the stimulus, rather than played it; 2) the study lasted a single session, rather than three; 3) the task was directly related to the stimulus, rather than unrelated; 4) the participants were 20% male, rather than 50%; 5) analyses did not include covariates for ratings of the stimulus or participants' habitual exposure to violent video games; 6) participants were exposed to a single fixed stimulus rather than one of three stimuli; 7) participants did not complete a measure of aggressive behavior.

Criticisms out of scope

Two of the authors' points raise criticisms that are beyond the scope of the paper. Point 6 suggests that the result from Hilgard (2021) is not generalizable to some broader population of stimuli because only one stimulus per condition was used. In this case, the maximal positive

¹ The effect size was calculated from the Hasan et al. (2013) raw data by first calculating the mean and standard deviation of each group's hostile attribution bias scores, then entering the means, standard deviations, and cell sizes into the `escalc()` function of the `metafor` package (Viechtbauer, 2010) to calculate a standardized mean difference. I made an error in Hilgard (2021) when calculating the Day 3 hostile attribution bias effect, entering the wrong standard deviations per group. This caused a small overestimation of the effect size. What I misreported as $d = 3.46$, 95% CI [2.72, 4.21] should be, in fact, $d = 3.34$, 95% CI [2.61, 4.06]. Although I deeply regret this error, I feel that the substance of Hilgard (2021) holds, as 3.34 standard deviations is still an effect of implausible magnitude.

control was not intended to generalize to such a population of stimuli; the purpose was to compare against Hasan et al. (2013).

Point 7 suggests it was inappropriate for Hilgard (2021) to question the claimed correlation between hostile attribution bias and aggressive behavior reported in Hasan et al. (2013), as Hilgard (2021) did not itself test that correlation. However, if the hostile attribution bias data cannot be relied upon, it is inappropriate to draw inferences from it, whether as an outcome in itself or as a mediator of aggressive behavior. It is not necessary to collect new data on aggressive behavior to make this point.

Different psychological processes can still be relevant

The authors' point 3 argues that the design in Hilgard (2021), in which participants watch the stimulus and then complete a *related* task, is not relevant or comparable to the design in Hasan et al. (2013), in which participants play the stimulus game and then complete an *unrelated* task. In their view, the method used in Hilgard (2021) is so different that it cannot be used to evaluate the plausibility of the effect in Hasan et al. (2013).

The authors suggest that, because the methods in Hilgard (2021) are different, the results are not relevant to the results of Hasan et al. (2013). I would suggest that relevance depends not only on the methods, but what one hopes to learn. If we are conducting a meta-analysis of the effects of subliminal alcohol primes on driving ability, a study of the effects of six shots of tequila is not relevant. However, if we are identifying studies with implausibly large effects, it can be appropriate to point out when a 12ms presentation of a beer is reported to impair driving ability more than a blood-alcohol concentration twice the legal limit. For this purpose, studies

need not be comparable in the narrow “apples to apples” sense expected by meta-analysis; they need only be comparable such that one effect ought to be strictly larger than the other.

Could other factors have reduced the effect size?

The authors’ remaining four points indicate differences between Hasan et al. (2013) and Hilgard (2021) and suggest that these differences caused the effect size in Hilgard (2021) to be smaller. Although these sundry factors may increase the effect size observed in a media violence experiment, I feel they are not sufficient to explain why the indirect Hasan et al. (2013) effect is larger than the direct Hilgard (2021) effect. The combined influence of these factors is likely to be small compared to the main effect of seeing the story character kill twenty people. Impression formation can happen quickly (e.g., Funder & Colvin, 1988) and bad behavior is particularly informative (i.e., the “positive-negative asymmetry”; for a review, see Baumeister et al., 2001).

In point 1, the authors suggest that the effect size in Hilgard (2021) is smaller because participants watched the stimulus rather than played it. I suspect that, when observing a character kill twenty people, most humans are likely to infer that the character is an aggressive person. Furthermore, I do not think that this inference is necessarily dependent on controlling that character’s actions oneself. Watching vs. playing might be important in *unrelated* tasks like those in Hasan et al. (2013) or Polman et al. (2008), but it is not likely to matter much for a directly related task like that in Hilgard (2021). Additionally, asking participants to play the videogame themselves introduces variance into the manipulation, as different players will have different experiences according to how they played the game. Some participants might go on a shooting rampage, others might drive around and listen to the radio, and others might simply struggle with the controls. Such heterogeneous experiences might be expected to increase the variance and thereby decrease the effect size.

In point 4, the authors suggest that the effect size in Hilgard (2021) is smaller because the sample was 20% men, as compared to the 50% men sample in Hasan et al. (2013). Hasan et al. (2013) did not find any significant moderation by gender on day 3 hostile attribution bias, $t(66) = 0.150$, $p = .881$. Hilgard (2021) had too few men per cell to be informative with regard to gender. Nevertheless, we can restrict the analyses in both datasets to women only. In Hasan et al. (2013), the effect among women was $d = 2.92$, 95% CI [1.97, 3.88]; in Hilgard (2021), the effect among women was $d = 1.66$, 95% CI [1.14, 2.18]. The reported effect among women in Hasan et al. (2013) still exceeds that of women in Hilgard (2021), one-tailed Welch's $t(43.4) = 3.63$, $p < .001$.

In point 5, the authors suggest that, had Hilgard (2021) statistically adjusted the outcome for twelve covariates, it would have reduced the error variance and increased the effect size. However, neither the results reported in Hasan et al. (2013) nor the effect size I extracted from Hasan et al. (2013) were adjusted for covariates.² Further, if one adjusts the results of Hasan et al. (2013) for these twelve covariates, the effect size shrinks to $d = 1.10$, 95% CI [0.58, 1.61], in contrast to the authors' statement that "Of course, controlling for these variables reduces the error variance, which increases the magnitude of effect sizes." (p. 12).

Accumulation across time cannot account for such massive effects

In their point 2, the authors suggest that the effect size from Hilgard (2021) is smaller because the manipulation is administered once, rather than three times. They also argue that this accumulation makes comparison to the population of other studies using the story completion task inappropriate, as all of these involve a single session.

² "Because the same pattern of results was obtained with and without the covariates, we used the simpler analyses that excluded the covariates." Hasan et al. (2013), p. 226

Impression formation happens quickly, and it seems sensible to expect that observing a single mass murder is sufficient to form an impression of the murderer's character. Additionally, a review of the literature indicates converging evidence from experimental, cross-sectional, and longitudinal research that the accumulation of media violence effects is not likely to reach the incredible magnitude reported in Hasan et al. (2013).

To my knowledge, there are few experimental tests of the accumulating effects of violent video games outside of the presently contested Hasan et al. (2013). In one pre-post field experiment, participants were randomly assigned to receive (or not receive) a copy of a violent game and play it for one month. Self-reported hours of game play were modestly correlated with participants' post-test scores on the Normative Beliefs in Aggression General Scale, $r = .12$, 95% CI [-.02, .25] (Williams & Skoric, 2005, table 2). This study is sometimes criticized on the grounds that participants in the control group may have also played violent video games during the study period (Anderson et al., 2007, p. 35-36), but that criticism applies similarly to Hasan et al. (2013), which does not report measuring or prohibiting video game use outside the laboratory. In another pre-post experiment, participants were randomly assigned to play either *Grand Theft Auto V* or *The Sims 3* for thirty minutes a day for two months. On none of the 52 outcomes did the Game $\times$ Time interaction explain more than 8% of the variance; the mean variance explained was 0.8% (Kühn et al., 2019, table 2). A variance explained of 76%, as in Hasan et al. (2013), is far outside the range suggested by these studies.

We can also consider the evidence from observational studies. Meta-analysis of cross-sectional studies finds that violent video game use explains 3% of variance in aggressive cognition, on average (Anderson et al., 2010). We should not expect the effect in this experiment to exceed that estimate by much. On the one hand, the effect in an experiment might be larger

due to the increased experimental control of having all participants play the same games under the same conditions. On the other hand, the correlation in cross-sectional observation might be larger because one observes the effects of *many* hours of video games, rather than a *single* hour, on aggressive cognition. Additionally, the effect in cross-sectional studies represents a combination of both the causal effect of violent video games on aggressive thoughts (violent games make people more aggressive) and the reverse-causal selection effect of aggressive thoughts and other confounds on violent video game use (people who are more aggressive play more violent video games).

Similarly, meta-analyses of longitudinal observational studies find that time 1 violent video game use explains approximately 1% of the variance in time 2 aggressive outcomes when accounting for individual differences in time 1 aggressive outcomes (Anderson et al., 2010; Prescott et al., 2018). Again, these estimates represent the influence of many hours of violent video games over weeks and months; it is implausible that a single hour would explain an additional 75 percentage points of variance.

Finally, not all the implausibly large effects reported by Hasan and colleagues involve three-day manipulations. In Hasan et al. (2012), the authors report effects of $d = 0.66$ (medium) among women and $d = 1.92$ (very large) among men following a single twenty-minute game session. Similarly, in Hasan et al. (2015), the authors report an effect of $d = 2.98$ (very, very large) following a single twenty-minute game session (as reported in Bushman, 2016). The authors have offered no explanation for these large effects found in single-session experiments.

Claims of past studies with larger effects cannot be confirmed

The authors contend that their study's effect size of $d = 3.34$ is not implausible, as the histogram in one meta-analysis of violent television effects (Paik & Comstock, 1994, Figure 1) indicates one or more hypothesis tests with an effect size of Fisher's $Z = 1.5$, equivalent to a Cohen's d of 4.26, or a Pearson r of 0.91, or variance explained of 83%. This effect size is likely in error for the same reasons that Hasan et al. (2013) is likely in error; it is too big to be true. It would be helpful to check and confirm this effect size, its calculation, and its standard error, but the meta-analysis does not document which study provided this test statistic.

Although effects of this size are extremely rare in social psychology, Bushman and colleagues seem to suggest that the effects of violent video games would be larger still, as it is theorized that video game violence has stronger effects than television violence. If that is indeed what they mean, I am quite surprised at the idea. If video game violence could explain more than 83% of the variance in aggressive outcomes, violent video game effects would not be the subject of ongoing scientific inquiry — they would be among the largest effects in all psychology, obvious to casual observation, and society would have had to adapt to protect itself from the fearsome changes even a single hour of video games would cause in players.

Clarifying one statistical misunderstanding

The authors further suggest that, although their Day 3 effect size exceeds that of the maximal positive control, it does not exceed it by much, writing that “the upper limit of the Hilgard 95% CI (i.e., $d = 2.18$) is only slightly smaller than the lower limit of the Hasan et al. 95% CI (i.e., $d = 2.65$).” First, a difference of 0.5 units of Cohen's d is no slight difference; it is larger than the difference between “zero effect” and “the typical effect in social psychology” (Richard et al., 2003). Second, comparing the confidence intervals of these two effect sizes is only a heuristic; the proper test is provided by Welch's t test reported in Hilgard (2021), which

compares the two means and their standard errors. This test yields $t(111.29) = 5.40, p < .001$.

The APA style convention of reporting small p -values as “ $< .001$ ” understates the significance of this result; the exact p -value is four in ten million.³

It’s time to think seriously about effect size

The present conversation indicates the utility of maximal positive controls in providing informative empirical data and stimulating enlightening critical discussion. Application of positive controls are a useful but under-appreciated method for testing the sensitivity of research measures and the quality of research data (e.g., Calin-Jageman, 2018). Maximal positive controls may provide a comparatively quick and inexpensive method for researchers seeking to identify plausible ranges of effect sizes on various psychological measurements. I hope that this method increases the speed of scientific self-correction by making it easier for psychologists to identify, critique, correct, and retract implausible results.

³ This test was originally reported in Hilgard (2021) as $t(108.67) = 5.66, p < .001$, again a slight overstatement due to my regrettable miscalculation of the Hasan et al. (2013) day 3 effect size.

References

- Anderson, C. A., Gentile, D. A., & Buckley, K. E. (2007). *Violent video game effects on children and adolescents: Theory, research, and public policy*. Oxford University Press.
- Anderson, C. A., Shibuya, A., Ihori, N., Swing, E. L., Bushman, B. J., Sakamoto, A., Rothstein, H. R., & Saleem, M. (2010). Violent video game effects on aggression, empathy, and prosocial behavior in Eastern and Western countries: A meta-analytic review. *Psychological Bulletin*, 136(2), 151–173. <https://doi.org/10.1037/a0018251>
- Baumeister, R. F., Bratslavsky, E., Finkenauer, C., & Vohs, K. D. (2001). Bad is Stronger than Good. *Review of General Psychology*, 5(4), 323–370. <https://doi.org/10.1037/1089-2680.5.4.323>
- Bushman, B. J. (2016). Violent media and hostile appraisals: A meta-analytic review. *Aggressive Behavior*, 42(6), 605–613.
- Bushman, B. J., Hasan, Y., & Bègue, L. (2022). Applied to Video Game Violence, Maximal Positive Controls is Far From Even a Minimal Demonstration: Comment on Hilgard (2021). *Journal of Experimental Social Psychology*.
- Calin-Jageman, R. J. (2018). Direct replications of Ottati et al. (2015): The earned dogmatism effect occurs only with some manipulations of expertise. *Journal of Experimental Social Psychology*. <https://doi.org/10/gdsnhr>
- Funder, D. C., & Colvin, C. R. (1988). Friends and strangers: Acquaintanceship, agreement, and the accuracy of personality judgment. *Journal of Personality and Social Psychology*, 55(1), 149–158. <https://doi.org/10.1037/0022-3514.55.1.149>

- Hasan, Y., Bègue, L., & Bushman, B. J. (2012). Viewing the world through “blood-red tinted glasses”: The hostile expectation bias mediates the link between violent video game exposure and aggression. *Journal of Experimental Social Psychology*, 48(4), 953–956.
- Hasan, Y., Bègue, L., & Bushman, B. J. (2015). *To be cruel or kind? Hostile expectation bias mediates the link between violent video games and aggression and prosocial behavior*. International Society for Research on Aggression.
- Hasan, Y., Bègue, L., Scharkow, M., & Bushman, B. J. (2013). The more you play, the more aggressive you become: A long-term experimental study of cumulative violent video game effects on hostile expectations and aggressive behavior. *Journal of Experimental Social Psychology*, 49(2), 224–227.
- Hilgard, J. (2021). Maximal positive controls: A method for estimating the largest plausible effect size. *Journal of Experimental Social Psychology*, 93, 104082.
<https://doi.org/10.1016/j.jesp.2020.104082>
- Kühn, S., Kugler, D. T., Schmalen, K., Weichenberger, M., Witt, C., & Gallinat, J. (2019). Does playing violent video games cause aggression? A longitudinal intervention study. *Molecular Psychiatry*, 24, 1220–1234. <https://doi.org/10.1038/s41380-018-0031-7>
- Paik, H., & Comstock, G. (1994). The Effects of Television Violence on Antisocial Behavior: A Meta-Analysis. *Communication Research*, 21(4), 516–546.
<https://doi.org/10.1177/009365094021004004>
- Polman, H., de Castro, B. O., & van Aken, M. A. G. (2008). Experimental study of the differential effects of playing versus watching violent video games on children’s aggressive behavior. *Aggressive Behavior*, 34(3), 256–264.
<https://doi.org/10.1002/ab.20245>

Prescott, A. T., Sargent, J. D., & Hull, J. G. (2018). Metaanalysis of the relationship between violent video game play and physical aggression over time. *Proceedings of the National Academy of Sciences*, 115(40), 9882–9888.

Viechtbauer, W. (2010). Conducting meta-analyses in R with the metafor package. *Journal of Statistical Software*, 36(3), 1–48. <https://doi.org/10.18637/jss.v036.i03>

Williams, D., & Skoric, M. (2005). Internet Fantasy Violence: A Test of Aggression in an Online Game. *Communication Monographs*, 72(2), 217–233.
<https://doi.org/10.1080/03637750500111781>