Protecting the internet after Christchurch: Far easier said than done

Matthew Feeney

March 20, 2019

Last week a man in Christchurch, New Zealand filmed himself murdering 50 Muslims at their places of worship. Daoud Nabi, an Afghan who fled to New Zealand to escape the Soviet-Afghan war, greeted the shooter saying, "Hello, brother," before the shooter murdered him. The shooter's livestream cut out after he conducted his slaughter inside the Al Noor Mosque and before he continued his rampage at the Linwood Islamic Centre.

Since the shooting, there have been calls for internet giants to do something to address the spread of graphic content. Anyone with an ounce of sympathy understands these calls, but we should keep in mind the difficulties associated with content moderation and the risks of crackdowns on particular kinds of content.

The scale of last week's attack may be difficult for many Americans to appreciate at first glance. New Zealand is a small country, about the same size as Colorado with a population a little higher than Louisiana's. It's also a country with a strong outdoor culture. As the son of a New Zealander, I grew up with stories of my father and his brothers playing rugby, spear-fishing, and rabbit hunting. On recent trips to New Zealand, I've shot handguns and semi-automatic rifles at a gun club. There are about three citizens for every gun in New Zealand.

Despite what many people think about the relationship between the number of guns in a country and crime rates New Zealand is a peaceful country, with a homicide rate far lower than that of the United States. In such a small and peaceful place an attack like the one last week is rare, and its impact profound. Although New Zealand authorities have yet to finalize recent homicide data, it looks likely that the gunman killed more people than were murdered in the entire nation in 2017. The shooter deployed this barbarity in about half an hour.

It's unclear how many people have watched the Christchurch murders video. According to Facebook, users viewed the original video 4,000 times. The social media giant removed 1.5 million videos of the attack within 24 hours. Other websites, including YouTube, scrambled to remove it.

While Americans are no strangers to mass shootings, there is something alien and at the same time uniquely 21st century about the Christchurch murders. Video of murders posted on social media sites are relatively rare, though not unheard of. The Christchurch shooter, a Millennial digital native, seems to have been radicalized online. His "manifesto" is replete with memes, jokes and symbolism well-known to those unfortunate enough to find themselves in the corners of the internet where alt-right ghouls waste their time.

Sadly, the attacks in Christchurch won't be the last to be livestreamed. Lawmakers know this. In the UK, Labour Party leader Jeremy Corbyn has said that social media companies should better

"deal with" graphic content. British Home Secretary Sajid Javid wants Google, Facebook YouTube and Twitter to "do more" to tackle extremism. Sen. Richard Blumenthal (D-Conn.) wants to get involved in private company policy and hold hearings on what he perceives to be social media companies' "abject failure" to halt the spread of graphic content. His colleague Sen. Mark Warner (D-Va.) wants social media companies to do more about extremist content, going so far as to propose amendments to Section 230, the legislation that protects internet platforms from legal action associated with content users upload or post.

Corbyn and Javid seem to think that social media companies can throw their computer science talent at the problem. But it behooves us to consider the scale and complexity of the problem.

Content like the Christchurch shooter's livestream shares a lot in common with legitimate content. YouTube and Facebook users often upload first-person shooter video-game footage. Audio of gunfire is found all over the Internet in content that is legal and historically valuable. Gun users and hobbyists regularly feature firearms in videos. Using machine learning to address graphic content is useful, but it's not going to be able to flag every piece of offensive footage.

False positives are a serious worry, as is the fact that users can edit videos, and stills from graphic footage are used by news organizations. Social media companies have a clear interest in eliminating graphic content from their platforms, but we should accept that solutions won't be perfect.

Using legislation as the bludgeon to use against social media companies has its own issues. Attempts to amend Section 230 so that social media companies can be held legally accountable for the proliferation of graphic content will harm user experience. Lawsuit-averse companies will overcorrect in response. After all, Craigslist shut down its "personal" pages section in the wake of the passage of sex trafficking legislation. Were lawmakers to amend Section 230 to address shooting videos, we should expect a similar overreaction.

Social media sites and other internet giants have the unenviable task of moderating content in an environment where billion of pieces of content are created every day. In engaging with such moderation there will be failures, but it's not clear that such failures should prompt government intervention, which could end up stifling legitimate and valuable content.

Feeney is the director of Cato's Project on Emerging Technologies, where he works on issues concerning the intersection of new technologies and civil liberties.