Lesley E. Weaver (SBN 191305)
Anne K. Davis (SBN 267909)
Joshua D. Samra (SBN 313050)
STRANCH, JENNINGS & GARVEY, PLLC
1111 Broadway, Suite 300
Oakland, California 94607
Telephone: (341) 217-0550
lweaver@stranchlaw.com
adavis@stranchlaw.com
jsamra@stranchlaw.com

Attorneys for His Majesty the King in Right of the Province of British Columbia and Board of Education of School District No. 59 (Peace River South)

UNITED STATES DISTRICT COURT
NORTHERN DISTRICT OF CALIFORNIA
SAN FRANCISCO DIVISION

HIS MAJESTY THE KING IN RIGHT OF THE PROVINCE OF BRITISH COLUMBIA and BOARD OF EDUCATION OF SCHOOL DISTRICT NO. 59 (PEACE RIVER SOUTH), individually,

Plaintiffs,

v.

SAMUEL ALTMAN, an individual, OPENAI FOUNDATION (F/K/A OPENAI, INC.), a Delaware corporation, OPENAI OPCO, LLC, a Delaware limited liability company, OPENAI HOLDINGS, LLC, a Delaware limited liability company, and OPENAI GROUP PBC, a Delaware public benefit corporation,

Defendants.

Case No. 3:26-cv-10743

COMPLAINT FOR:

(1) Negligence (Failure to Warn Law Enforcement);
(2) Negligent Entrustment;
(3) Aiding and Abetting a Mass Shooting;
(4) Negligence (Failure to Warn);
(5) Negligent Undertaking;
(6) Negligence (Design Defect);
(7) Strict Product Liability (Design Defect);
(8) Strict Product Liability (Failure to Warn).

DEMAND FOR JURY TRIAL

His Majesty the King in Right of the Province of British Columbia (the "Crown" or "Province") and Board of Education of School District No. 59 (Peace River South) ("SD59") (collectively, "Plaintiffs") bring this action against Defendants Samuel Altman, OpenAI Foundation (f/k/a OpenAI, Inc.), OpenAI OpCo, LLC, OpenAI Holdings, LLC, and OpenAI Group PBC (collectively, "OpenAI" or "Defendants"), upon personal knowledge as to the factual allegations pertaining to themselves and upon information and belief as to all other matters.

NATURE OF THE ACTION

1. On the morning of February 10, 2026, an eighteen-year-old resident of Tumbler Ridge, British Columbia, shot and killed their mother and eleven-year-old half-brother in the family home. They then drove to Tumbler Ridge Secondary School ("TRSS"), a public school operated by Plaintiff SD59, armed with a long gun and a modified rifle, and opened fire on defenseless children. Before the shooter took their own life, Jesse Van Rootselaar ("Van Rootselaar") killed five students and an education assistant and wounded two students who were transported for care. Some children watched as classmates and an educator were shot at point-blank range; others hid in terror in classrooms, closets, and bathroom stalls. Survivors—students, teachers, parents—sustained lasting physical and psychological injuries. Roughly 160 students, teachers, and staff were trapped in classrooms and closets at TRSS for hours during the assault and recovery. Twenty-five injured sought medical care after they were evacuated from TRSS. It was one of the deadliest mass shootings in Canadian history and was extremely traumatic.

2. Tumbler Ridge is a small, remote community of roughly 2,400 people in northeastern British Columbia. Because of the small size of the population, the tragedy has had an outsize impact. Roughly one in fifteen residents was inside TRSS when the shooting began. The trauma was very close and inescapable in this tightly knit community. The school where the assault occurred sits in the heart of the town, geographically and culturally. Following the assault, it was a symbol of the trauma that occurred there. The TRSS was immediately closed and has not reopened. Following community consultation, the trauma-informed conclusion was reached that no child should be made to return to the building where the

violence occurred.[1] Demolition of the school began in earnest in August 2026, offering at least some small relief to those who have suffered the most intensely since the trauma occurred.

3. In support of the citizens they serve, this action is brought by the two public bodies that were forced to absorb the consequences of the attack and to assist with the recovery of the community. Plaintiff SD59 owns facilities that were damaged and rendered unusable by the assault, employs the staff that supports the students and teachers, has had to refashion the facilities and resources necessary to continue its interrupted educational mission. Plaintiff His Majesty the King in Right of the Province of British Columbia has expended, and continues to expend, significant resources to support the recovery and rebuilding of the community, including providing human resources (educators, clinicians, victim-services workers, and police) and infrastructure. The Crown and SD59 seek to recover the cost that this tragedy and its foreseeable consequences required them to expend, and injunctive relief seeking preventive measures beyond money damages to ensure that a tragedy like this does not happen again.

4. The costs the Crown and SD59 seek to recover are not the ordinary, recurring costs of government that are funded through taxation and spread across all taxpayers. They are the extraordinary costs of responding to an unprecedented event—the deadliest school shooting in Canada in more than three decades—and of abating its continuing consequences. The Province declared a provincial state of emergency and mobilized a whole-of-government response. Both the Province and SD59 were required to adopt new policies, programs, and permanent infrastructure that did not exist before February 10, 2026, and would not exist but for the attack. These include establishing a trauma-informed care program, temporary education facilities, and a replacement school and wellness center designed to support the long term needs of a traumatized community arising out of this shooting. These are measures to abate the recovery of the community of Tumbler Ridge, not the routine provision of public services.

5. The Crown and SD59 also bring this suit because this mass assault could have been prevented had OpenAI warned the Crown or the school, by admission of whistleblowers at OpenAI. OpenAI's internal procedures in effect at the time of the shooting included an automated review of

[1] Lauren Vanderdeen, "Tumbler Ridge plans for how to return students to education," CBC News (Feb. 16, 2026), https://www.cbc.ca/news/canada/british-columbia/tumbler-ridge-bc-school-plan-9.7092605 (prior to the community consultation, British Columbia Premier David Eby also told students, "I will promise you that not one of you will ever be forced to go back to that school.").

ChatGPT's content, which could be escalated to a human safety review team for assessment and reporting to authorities. In June 2025, OpenAI's automated monitoring identified the shooter's account and routed it to the specialized human review team OpenAI maintains to review users who may be planning to commit acts of violence. The human review team concluded, based on the chats between Van Rootselaar and OpenAI, that there was a credible, specific risk of harm to others. The human review team recommended a referral to the Royal Canadian Mounted Police ("RCMP").

6. But OpenAI leadership ***rejected*** that recommendation. Instead, the Company deactivated the account with the flagged communications, and notified no one external to the Company. OpenAI did not take steps to prevent its chatbot from inciting Van Rootselaar further. The private deactivation did not address the risk of harm to others or stop Van Rootselaar, who continued engagement with ChatGPT through a second account.

7. These facts are only public now because OpenAI whistleblowers anonymously reported what had occurred to *The Wall Street Journal*.[2] If management had its way, no one would begin to understand the role that OpenAI and its defective product played in the assault that has forever changed Tumbler Ridge.

8. One telephone call to the RCMP could have prevented the tragedy. OpenAI has recently taken the position in filings before this Court that the RCMP was partly to blame for not stopping the tragedy, a legal position that all but admits that OpenAI should have warned them. That is, OpenAI had a duty to warn the RCMP. The RCMP contracts directly with the Province to provide policing services in Tumbler Ridge. A warning from OpenAI would have reached an RCMP that already engaged with Van Rootselaar; had visited their residence in response to mental-health concerns; and had removed firearms from the residence in the past, as OpenAI notes. In short, had OpenAI shared the information it possessed about Van Rootselaar, the RCMP would have taken preventive action. Instead, OpenAI stayed silent and assumed the risk of anything that transpired as a result of management's decision not to notify the RCMP.

[2] When Van Rootselaar's chat logs are produced in this action, Plaintiffs will seek to amend this complaint. We propose an immediate production deadline for the chat logs, and scheduling a deadline to amend Plaintiffs' complaint shortly thereafter.

9. OpenAI's stated rationale was a concern for the shooter's privacy, but its own interests better explain its silence. Upon information and belief, OpenAI was seeking to avoid implementation of a hard line rule to refer planning of real-world violence to authorities, perhaps due to how frequently its product is implicated in threats to human life. Requiring such disclosures would be incompatible with the company's public position that ChatGPT is safe. It could also threaten the valuation underlying OpenAI's anticipated initial public offering. Rather than expose those risks, OpenAI accepted the consequences of its silence. A mass murder in the only secondary school in Tumbler Ridge followed.

10. Publicly, OpenAI has paid lip service to regret, and has promised to make changes. On April 23, 2026, after British Columbia Premier David Eby and Tumbler Ridge Mayor Darryl Krakowka demanded a response from OpenAI,[3] OpenAI Chief Executive Officer Samuel Altman published a letter to the community of Tumbler Ridge. In it, Altman stated: "I am deeply sorry that we did not alert law enforcement to the account that was banned in June."[4] That statement is an admission that OpenAI identified the risk and failed to act on it. OpenAI has since acknowledged that, under the escalation guidelines it adopted after June 2025, the shooter's account "would have been referred to law enforcement." OpenAI's Ann O'Leary similarly apologized, promising a string of reforms.

11. Neither OpenAI nor Ms. O'Leary have followed up on these promises, despite public efforts by the Attorney General of British Columbia to engage with them. This action seeks implementation of the injunctive relief that Mr. Altman and Ms. O'Leary offered to assuage public sentiment, but have refused to implement.

12. Accordingly, the Crown, led by Attorney General Sharma, and SD59 jointly bring this action to hold OpenAI and Sam Altman accountable for designing a dangerous product, distributing it to every home with internet access, ignoring the warnings of their own safety team, refusing to notify

[3] *See* Lauren Vanderdeen, "Families of Tumbler Ridge, B.C., mass shooting victims suing OpenAI in California," CBC News (Apr. 29, 2026), https://www.cbc.ca/news/canada/british-columbia/tumbler-ridge-b-c-mass-shooting-families-suing-openai-9.7181214 (Premier Eby stating the attack "could have potentially been prevented").

[4] Letter from Samuel Altman to the community of Tumbler Ridge (Apr. 23, 2026), as reported in The Globe and Mail, "B.C. Insider: Families of Tumbler Ridge victims sue OpenAI" (May 5, 2026), https://www.theglobeandmail.com/canada/british-columbia/article-bc-insider-families-of-tumbler-ridge-victims-sue-openai/.

authorities when they knew the shooter was planning gun violence, inviting the shooter back onto the platform after deactivating the shooter's account, and choosing corporate self-interest over the lives of children. They seek compensation for the not just foreseeable but ***known*** harm OpenAI inflicted, the damages that they incurred and are incurring, and injunctive relief to ensure that this tragedy does not happen again.

PARTIES

13. Board of Education of School District No. 59 (Peace River South) is a board of education and public school district duly constituted under the *School Act*, R.S.B.C. 1996, c. 412, with its administrative offices in Dawson Creek, British Columbia. SD59 is responsible for the delivery of public education in the Peace River South region, including in the community of Tumbler Ridge. SD59 owns, operates, and maintains the schools in that region, including TRSS and Tumbler Ridge Elementary School. SD59 owned the TRSS campus that was rendered unusable as a result of the attack. SD59 brings this action in its own right to recover the damages it has sustained and will sustain as a direct and proximate result of Defendants' conduct.

14. His Majesty the King in Right of the Province of British Columbia as represented by the Attorney General of British Columbia, the Honorable Niki Sharma, K.C., Deputy Premier and Attorney General. The Province declared a provincial state of emergency on February 11, 2026, and mobilized a whole-of-government response to the attack across multiple ministries, including the Ministry of Education and Child Care, the Ministry of Children and Family Development ("MCFD"), the Ministry of Health, and the Ministry of Emergency Management and Climate Readiness. The RCMP, which responded to the attack and investigated it, is the Provincial police force within Tumbler Ridge.

15. Defendant Samuel Altman is a natural person residing in California. As CEO and Co-Founder of OpenAI, Altman oversaw and led the design, development, and deployment of ChatGPT and security measures related to ChatGPT. As has been widely reported, in 2024, Defendant Altman knowingly accelerated GPT-4o's public release while intentionally side-stepping safety protocols and warnings relating to the risk GPT-4o posed to vulnerable users.

16. Defendant OpenAI Foundation (f/k/a OpenAI, Inc.) is a Delaware corporation with its principal place of business at 3180 18th Street, San Francisco, California 94110. Upon information and

belief, at all times relevant to this action, Defendant OpenAI, Inc. was the non-profit parent company that governed the other OpenAI entities and exercised oversight over its for-profit subsidiaries, including OpenAI OpCo, LLC and OpenAI Holdings, LLC. As the governing entity, OpenAI, Inc. was responsible for establishing the company's risk-management framework, the "Model Specs," relating to the framework of ChatGPT, and deploying the company's AI models.

17. Defendant OpenAI OpCo, LLC is a Delaware limited liability company with its principal place of business at 3180 18th Street, San Francisco, California 94110. OpenAI OpCo LLC, formerly known as OpenAI LP, is a subsidiary of OpenAI Global LLC, and is the sole member of OpenAI, LLC. OpenAI OpCo LLC serves as the for-profit arm of OpenAI, overseeing the development and commercialization of OpenAI's defective product at-issue. Its responsibilities include management and release of the ChatGPT Plus subscription service, and maintaining and making ChatGPT available to users.

18. Defendant OpenAI Holdings, LLC is a Delaware limited liability company with its principal place of business located at 3180 18th Street, San Francisco, California 94110. OpenAI Holdings, LLC is the subsidiary of OpenAI, Inc. that owns and controls the core intellectual property at-issue, including the defective ChatGPT models. In this role as legal owner of the intellectual property at issue and beneficiary of commercialization of the defective product, OpenAI Holdings, LLC is liable for the harm caused by the alleged defects.

19. Defendant OpenAI Group PBC is a Delaware public benefit corporation with its principal place of business in San Francisco, California. OpenAI Group PBC was formed on October 28, 2025, as part of a corporate restructuring in which OpenAI's for-profit operations were consolidated under a new public benefit corporation. OpenAI Group PBC is the successor to the for-profit entities that designed, approved, deployed, and profited from ChatGPT, and the company continues to deploy and profit from ChatGPT. As the successor, OpenAI Group PBC is liable for the harm caused by the conduct of its predecessor entities.

20. Defendants individually and collectively played a direct and consequential role in the design, development, testing, approval, and release of GPT-4o and subsequent models. OpenAI Foundation (f/k/a OpenAI, Inc.) developed and published the safety mission of OpenAI, which it failed to enforce. OpenAI OpCo, LLC developed and released ChatGPT and managed the paid subscription service. OpenAI

Holdings, LLC owned the underlying intellectual property and profited from its release, including the release of GPT-4o. OpenAI Group PBC is the successor entity that continues to deploy and profits from ChatGPT, including GPT-4o, and is liable for the conduct of its predecessors. Sam Altman led and directed the prioritization of the release of GPT-4o despite clear safety concerns. Collectively, these Defendants constitute the principal actors whose decisions directly resulted in the harm at issue.

JURISDICTION AND VENUE

21. This Court has subject matter jurisdiction pursuant to 28 U.S.C. §§ 1332(a)(2) and 1332(a)(4) because British Columbia is a Province of Canada, SD59 is a board of education and public school district in Canada, Defendants are citizens of California, and the amount in controversy exceeds the sum of $75,000.00.

22. This Court has personal jurisdiction over the Defendants because it is where OpenAI is headquartered. They conduct business in this District, reside in this District, and a substantial part of the events or omissions giving rise to the Crown's and SD59's claims occurred in this District.

23. Venue is proper in this District under 28 U.S.C. § 1391(b), because a substantial part of the events giving rise to the Crown's and SD59's claims occurred in and emanated from this District, and OpenAI Defendants reside in this District. Specifically, the design and marketing of this defective product occurred in California and the chats, a byproduct of OpenAI's ChatGPT, were created in California. And upon information and belief, the decision not to warn the RCMP was also made in California.

DIVISIONAL ASSIGNMENT

24. Pursuant to Civil L.R. 3-2(c), divisional assignment to San Francisco is appropriate, because OpenAI is headquartered in San Francisco County and a substantial part of the events or omissions giving rise to the claim occurred in San Francisco County.

FACTUAL BACKGROUND

I. The Tumbler Ridge Mass Shooting and Its Impact on the Province and SD59

25. On February 10, 2026, at approximately 1:20 p.m., an eighteen-year-old resident of Tumbler Ridge, Jesse Van Rootselaar, carried out one of the deadliest mass shootings in Canadian history. Earlier

that day,[5] at the family home, Van Rootselaar killed their mother and their eleven-year-old half-brother, then carried a long gun and a modified rifle to TRSS and began firing. Although eighteen at the time of the attack, Van Rootselaar was a minor under the law of British Columbia—where the age of majority is nineteen—throughout the entire period of their use of ChatGPT,[6] and remained a minor under British Columbia law on that day.

26. When Van Rootselaar entered the school, they were spotted almost immediately. School staff triggered the lockdown alarm and called police, and the RCMP issued a Police-Initiated Public Alert.[7] Prior to this event, SD59 implemented and drilled an Emergency Operations Plan. This plan includes locking entrances and training staff and students to lock classrooms. This training saved lives during this event. The shooter was unable to enter locked classrooms. Van Rootselaar fired at least 75 rounds and none of the shots fired into classrooms struck the people sheltering inside.

27. Members of the local RCMP detachment arrived within two minutes of the initial call and entered the school while under fire from the shooter.[8] Inside, RCMP located both the deceased shooter (who died of a self-inflicted injury) and numerous victims. Two critically wounded students were airlifted out of the community to hospitals in Vancouver for higher-level care. Approximately twenty-five other people sought medical attention at the local medical clinic in Tumbler Ridge. Because initial reports suggested multiple shooters, both TRSS and the elementary school remained locked down until the scene was cleared. TRSS students were evacuated first, walking through the crime scene, and the more than 200 elementary students were released to guardians through a formal reunification process.

28. Six people were killed inside the school in addition to the shooter: a 39-year-old female

[5] *See* Royal Canadian Mounted Police, "Update: Tumbler Ridge fatal shootings" (Feb. 11, 2026), https://rcmp.ca/en/bc/tumbler-ridge/news/2026/02/4350208; *see also* 2026 Tumbler Ridge shooting, Wikipedia, https://en.wikipedia.org/wiki/2026_Tumbler_Ridge_shooting (collecting contemporaneous reporting).

[6] Age of Majority Act, R.S.B.C. 1996, c. 7, s. 1 (age of majority in British Columbia is nineteen years).

[7] Royal Canadian Mounted Police, "Update: Police Emergency Alert incident in Tumbler Ridge cancelled" (Feb. 10, 2026), https://rcmp.ca/en/bc/tumbler-ridge/news/2026/02/4350182.

[8] Royal Canadian Mounted Police, "Update: Tumbler Ridge fatal shootings," supra n.5 (local detachment members arrived within two minutes; officers encountered active gunfire).

education assistant and five children,[9] ages 12 and 13. Two further victims, the shooter's 39-year-old mother and 11-year-old half-brother, were found deceased at the family home. In total, eight victims were killed, two were transported for medical care in Vancouver, and twenty-five sought treatment in the community. The attack traumatized the roughly 160 students and staff inside the building, their families, and the wider community. It was and is an immensely personal tragedy for this small community. As one local official stated, "we all know the victims."

29. Unlike other larger communities, in Tumbler Ridge there were no redundancies of personnel or facilities to carry the community following the tragedy. There was no other school building to house students for education or counseling following this trauma. There were few available adults to step in for teachers and administrators who were directly impacted by the assault. Students, parents and all those impacted needed immediate support and counseling, and there simply were not enough people in this small town to match the need. The responsibility for responding to and remediating the consequences of the attack fell directly and foreseeably on the Crown and SD59.

30. The Province's response was immediate, comprehensive, and costly. On February 11, 2026, the Province declared a state of emergency and activated its powers under the *Emergency and Disaster Management Act*, S.B.C. 2023, c. 37. In a remote community that had no excess capacity—no substitute teachers, no standby clinicians, and no extra police officers—and whose own frontline providers were themselves traumatized, the Province stood up an emergency response operation to coordinate across ministries. Within days, additional mental-health clinicians and victim-services workers were deployed to see patients. Drop-in appointments became available and the Tumbler Ridge Health Centre extended its hours to 7 a.m. to 11 p.m. Further, MCFD, Child & Youth Mental Health, the Provincial Health Services Authority, and Northern Health mobilized local supports.[10] The Province also launched an information website and a 24/7 crisis line.

---

[9] *Id.*; *see also* Royal Canadian Mounted Police, "Update: Investigative update on Tumbler Ridge shooting" (Feb. 13, 2026), https://rcmp.ca/en/bc/tumbler-ridge/news/2026/02/4350292.

[10] Northern Health, "Mental health supports available in Tumbler Ridge" (Feb. 11, 2026), https://stories.northernhealth.ca/news/mental-health-supports-available-tumbler-ridge; Adriana Fallico, "Mental health support after Tumbler Ridge shooting 'essential,' experts say," Global News (Feb. 13, 2026), https://globalnews.ca/news/11664704/mental-health-supports-tumbler-ridge/ (extended Health Centre hours 7 a.m.–11 p.m.; MCFD, PHSA, and Northern Health mobilizing supports).

31. Public reporting confirms that the supports sent to the community included a children's psychiatrist and about ten mental-health clinicians available at the Tumbler Ridge Health Centre at all times, more than 30 RCMP Victim Services staff on the ground, and additional Northern Health mental-health clinicians.[11] Within 48 hours the Province had approximately 45 personnel in Tumbler Ridge—including SD59 counsellors, MCFD clinicians, Northern Health medical clinicians, victim-support workers, and school-safety officers—supplemented by childhood-trauma experts from BC Children's Hospital. Between February 11, 2026, and early April, Child & Youth Mental Health alone delivered well over 900 counseling sessions to more than 100 unique children and families, and Northern Health treated more than 130 unique adults. The Province opened a logistics center through the Ministry of Emergency Management and Climate Readiness to house, feed, transport, and orient the rotating personnel who had to be brought in from out of town.

32. The Province's posture has since shifted from emergency response to long-term recovery, and provincial and local officials have publicly committed to long-term mental-health support for the community rather than short-term help.[12] The Province has committed to maintain a Child & Youth Mental Health team in Tumbler Ridge that operates five days a week, has hired a recovery manager to remain in the community for the foreseeable future, and has established a long-term grief-and-bereavement working group. The Province and the Government of Canada have funded the construction of a new secondary school and the modernization of the community's health-care facilities as a combined wellness and healing

[11] *See* Brieanna Charlebois, "A search for mental health answers after the trauma of Tumbler Ridge shootings," The Canadian Press (Mar. 5, 2026), as published by National Observer, https://www.nationalobserver.com/2026/03/05/news/search-mental-health-answers-after-trauma-tumbler-ridge-shootings (B.C. Health Ministry statement: a children's psychiatrist and about ten mental-health clinicians at the Tumbler Ridge Health Centre; RCMP Victim Services made more than 30 staff available; Northern Health sent additional clinicians); Andrea Wood, "Tumbler Ridge shooting highlights rural northern B.C.'s mental health deserts" The Globe and Mail (Mar. 4, 2026), https://www.theglobeandmail.com/canada/article-tumbler-ridge-shooting-rural-northern-bcs-mental-health-deserts/.

[12] *See* National Observer, *supra* n.11 (Mayor Darryl Krakowka: "We're not looking for short-term help when it comes to counsellors. We're looking for the long term," and that he conveyed the same to Premier Eby).

project.[13] Demand for scheduled mental-health services remains high and spikes around anniversary dates and other triggering events; these services are expected to be needed for years to come.

33. TRSS, the building in which the shooting took place, was located in the center of the town. It was immediately and irrevocably associated with the trauma. Following a trauma-informed community consultation, the Peace River South Board of Education determined that TRSS would be demolished to avoid further and ongoing traumatization. As a result, a new secondary school will be built at a different location. The Province and the Government of Canada have committed to fund that project, having already demolished the existing school with preparatory work on the construction of a new school to begin in 2026.[14]

34. In the interim, SD59, with the logistical and financial support of the Province, installed a temporary campus of modular classrooms on the grounds near Tumbler Ridge Elementary.[15] The interim campus was assembled first from fourteen Ministry of Forestry trailers and then from larger modular classrooms; SD59 was required to furnish it, recruit and train replacement educators for local staff who could not return, and reorganize instruction around the absence of murdered students and colleagues.

35. None of these expenditures would have been necessary but for the attack, and the attack would not have occurred but for OpenAI's deployment of an exceedingly dangerous product, which facilitated the mental instability of the shooter. Following production of the chats, upon information and belief, further information will be available regarding OpenAI's interactions with the shooter, explaining why anonymous whistleblowers felt compelled by the contents of those chats to go public with the story that they had voted to warn Canadian authorities. The Crown and SD59 have incurred, and are incurring, quantifiable economic damages, including but not limited to extraordinary policing response, acute and

[13] Andrew Kurjata, "Carney announces $200M for new school, updated health centre in Tumbler Ridge," CBC News (June 18, 2026), https://www.cbc.ca/news/canada/british-columbia/carney-tumbler-ridge-health-care-centre-9.7010777 (federal and provincial governments each committing $100 million).

[14] Government of British Columbia, "Supporting students in Tumbler Ridge with new secondary school" (May 7, 2026), https://news.gov.bc.ca/releases/2026INF0022-000511; *see also* Kurjata, *supra* n.13.

[15] *See* Charlebois, supra n.11 (metal-clad portable classrooms installed on the grounds of Tumbler Ridge Elementary, roughly three weeks after the shooting; B.C. Infrastructure Minister describing them as "an important step toward restoring routine").

ongoing health care, victim services and trauma-informed mental-health care, educational continuity and staffing, facility remediation, demolition, and the construction of a replacement school and wellness center as a direct and proximate result of Defendants' conduct.

36. The consequences of the attack, and the resources demanded to restore this community, are ongoing. Before February 10, 2026, neither the Province nor SD59 maintained the programs, staffing, and facilities the attack made necessary. These are not routine expenditures within the ordinary run of government operations, and they are not the recurring public services that governments fund through taxation and spread across all taxpayers. They were extraordinary, non-recurring measures—new policies, programs, and permanent infrastructure adopted in response to a singular event—undertaken to abate a continuing harm to the health and safety of the community, and they will require the expenditure of public resources for years to come.

37. In the aftermath of the shooting, the RCMP warned of the risk of copycats and have investigated threats made subsequent to it. Never believing that anything like this trauma could take place in Tumbler Ridge, community members fear another attack. The Province and SD59 continue to expend time and resources to provide a strong sense of safety to the school community.[16]

38. These ongoing threats and harm leave the Province and SD59 without an adequate remedy at law for OpenAI's ongoing negligence and deployment of a defective product without appropriate warnings.

II. OpenAI Knew the Shooter Was Engaged in Planning Gun Violence, Refused to Warn Law Enforcement, and Deployed a Product Designed to Facilitate the Attack

39. OpenAI designed, developed, trained, tested, deployed, operated, and controlled ChatGPT, including the GPT-4o model and subsequently released models, and offered it to consumers, including the shooter, as an interactive artificial-intelligence assistant. Since ChatGPT's launch, OpenAI's valuation has

[16] Michelle Ghoussoub, *Unsealed Documents Reveal Details of RCMP's Response, Investigation into Tumbler Ridge Shooting*, CBC News (Aug. 13, 2026), https://www.cbc.ca/news/canada/british-columbia/tumbler-ridge-unsealed-documents-9.7305135 (RCMP Const. Jonathan Paquin stating: "I am aware that mass casualty incidents can inspire copycat attacks. For example, the Columbine High School massacre is well known to have had this effect where it inspired other shootings in the following years...").

grown from approximately seventeen billion dollars ($17,000,000,000) to almost one trillion dollars ($1,000,000,000,000).

40. In June 2025, OpenAI's automated review identified substantial activity on the shooter's account involving gun-violence scenarios and routed it into the review process OpenAI has publicly described for users who may be planning to harm others.[17] Under that published process, such conversations are "reviewed by a small team trained on our usage policies and who are authorized to take action, including banning accounts," and, "[i]f human reviewers determine that a case involves an imminent threat of serious physical harm to others, we may refer it to law enforcement."

41. OpenAI's reviewers acted on that process. They examined the flagged conversations, concluded that the shooter presented a credible and specific risk of gun violence to real people, and—according to OpenAI employee whistleblowers whose account was first reported by the Wall Street Journal—several of them recommended that the company notify the RCMP.[18] OpenAI leadership declined, taking the position that the matter did not satisfy a "higher threshold" for "credible and imminent" threat reporting. OpenAI deactivated the account, made no referral, and gave the community no warning.

42. When OpenAI's leadership declined to notify the RCMP in June 2025, at a minimum, the company already knew that others had used its product to prepare for real-world violence, including the January 2025 Las Vegas Cybertruck detonation, the April 2025 Florida State University shooting, and a May 2025 school stabbing in Finland. The reviewers who examined the shooter's account were evaluating conduct that fit an already-recurring pattern. OpenAI's reviewers recommended a referral, and, in a decision that departed from its stated policy, its leadership overruled them.

43. In June 2025, roughly a dozen OpenAI employees reviewed the flagged conversations on the shooter's first account, assessed them as presenting an imminent risk of serious harm to others, and

[17] OpenAI, "Helping people when they need it most" (Aug. 26, 2025), https://openai.com/index/helping-people-when-they-need-it-most/; OpenAI, "Our commitment to community safety," https://openai.com/index/our-commitment-to-community-safety/.

[18] *See* Clare Duffy, "Families of Tumbler Ridge shooting victims sue OpenAI and CEO Sam Altman," CNN Business (Apr. 29, 2026), https://www.cnn.com/2026/04/29/tech/openai-tumbler-ridge-canada-shooting-lawsuits; Geoff Brumfiel, "Families sue OpenAI over Tumbler Ridge mass shooter's use of ChatGPT," NPR (Apr. 29, 2026), https://www.npr.org/2026/04/29/nx-s1-5798896/tumbler-ridge-mass-shooting-chat-gpt-lawsuit (reporting whistleblower account first published by *The Wall Street Journal*).

recommended that OpenAI contact the RCMP. OpenAI's leadership overrode that recommendation, applied an unknown higher threshold for a warning, and suspended the account instead of reporting it.

44. OpenAI's after-the fact explanation of its June 2025 decision falls short. A former member of OpenAI's investigations team has publicly stated that the company was reluctant to make such referrals because they would require it to disclose how its product contributes to the threat environment.[19] That disclosure risk—to the company's safety messaging, its valuation, and its anticipated public offering—better explains the decision.

45. OpenAI's post-attack characterizations—that it had "banned" the shooter and that they "evaded" its safeguards—are not accurate. OpenAI did not ban Van Rootselaar; it deactivated one account and allowed its chatbot to continue to engage Van Rootselaar until the day of the Tumbler Ridge attack. Nothing prevented Van Rootselaar from continuing to use ChatGPT after their first account was deactivated. OpenAI's own Help Center and support communications explain to all deactivated users, including those deactivated for violence and self-harm violations, how to resume their access to ChatGPT.[20]

46. Independent of anything the shooter shared with ChatGPT during use, OpenAI collected identifying and locating information the moment each account was opened and every time it was used. Under OpenAI's own privacy policy, the information OpenAI collects includes the account holder's name and email address, and—as system-generated data—the device's IP address, the "general area" from which the device accesses the service as determined from that IP address, device and browser identifiers, and time-zone and usage data.[21] OpenAI thus possessed, for both of Van Rootselaar's accounts, the very

[19] *See* Georgia Wells, "OpenAI Employees Raised Alarms About Canada Shooting Suspect Months Ago," The Wall Street Journal (Feb. 21, 2026), https://www.wsj.com/us-news/law/openai-employees-raised-alarms-about-canada-shooting-suspect-months-ago-b585df62 (reporting, based on accounts of former members of OpenAI's investigations team, that the company was reluctant to refer users to law enforcement because doing so would require disclosing how its product contributes to the threat environment); *see also* Brumfiel, *supra* n.18 (whistleblower account first published by The Wall Street Journal).

[20] OpenAI Help Center, "Why Was My OpenAI Account Deactivated?", https://help.openai.com/en/articles/10562188-why-was-my-openai-account-deactivated (identifying violence and self-harm usage-policy violations as grounds for deactivation and describing how to prevent future deactivations).

[21] *See* OpenAI Privacy Policy, https://openai.com/policies/us-privacy-policy/ and https://openai.com/policies/row-privacy-policy/ ("We determine the general area from which your device accesses our (continued...)

information a law-enforcement referral would have required: the identity of the user, evidence that the two accounts were opened under the same identity, and the general location from which they were writing. OpenAI's claim that it withheld a referral to protect the user's "privacy" is difficult to reconcile with the identifying and location data its systems had already collected and retained.

47. The Tumbler Ridge attack was a foreseeable consequence of OpenAI's deliberate design choices. In May 2024, OpenAI removed ChatGPT's prior categorical refusal of violent content because refusals suppressed user engagement. OpenAI reconfigured GPT-4o to maximize and prolong engagement: adding persistent cross-session memory (which was switched on by default)[22] that recalled and built on a user's prior conversations; tuning the model toward sycophantic affirmation; and cultivating an anthropomorphic, human-like conversational register that fostered emotional dependency. For a user fixated on violence, these features encouraged, elaborated, and reinforced violent ideation instead of interrupting it or directing the user to real-world help.

48. OpenAI's written behavior rules, the Model Spec, codified these choices. Under the Model Spec in force during the shooter's most active period, the only categorically "prohibited" content was sexual content involving minors.[23] Detailed weapon and explosive instructions were "restricted," but violent and mass-casualty content in a fictional or roleplay frame was merely "sensitive" content that OpenAI deemed permissible in context. Content bearing on "imminent real-world harm" fell into a "take extra care" tier in which the model was instructed only to "try" to prevent harm, to "assume best intentions," and to "never ask the user to clarify their intent for the purpose of determining whether to refuse or comply."

Services based on information like its IP address"; account information collected includes name and email address; system-generated data includes IP address, approximate location, device and browser information, time zone, and usage data).

[22] OpenAI, "Memory and new controls for ChatGPT" (Feb. 13, 2024; updated 2025), https://openai.com/index/memory-and-new-controls-for-chatgpt/ (describing the memory feature, on by default).

[23] OpenAI Model Spec (Feb. 12, 2025), https://model-spec.openai.com/2025-02-12.html (three-tier content taxonomy).

An attempt at an anti-violence "red-line" principle—that the model should never facilitate acts of violence—did not enter the Model Spec until December 18, 2025.[24]

49. Throughout the period relevant to this action, OpenAI publicly and proudly represented that ChatGPT was safe and that OpenAI protected both minors and the general public from the product's known dangers. OpenAI told the public that it detected users who appeared to be planning to harm others and routed them to a trained safety team authorized to ban accounts and to refer imminent threats of serious physical harm to law enforcement. OpenAI told the public that it had adopted teen-safety measures, an age-prediction system, and parental controls to protect users it identified as minors.[25] In December 2025, OpenAI told the public that it had deployed real-time classifiers to catch dangerous conversations as they happened rather than after the fact, together with an alleged anti-violence "red-line" under which its models should never facilitate acts of violence.[26] OpenAI held these measures out as protections that made the product safe for minors and for the public.

50. These protections were inadequate and unenforced. The behavioral rules OpenAI announced were instructions to the model, not hard technical constraints.[27] OpenAI's own product had a documented history of disregarding such instructions by continuing to mirror and validate users in crisis, and continuing to produce prohibited content. OpenAI's minor protections were keyed to users it identified as under eighteen, so they did not apply to a user who had passed their eighteenth birthday but was still considered a minor under the law of their own province. Finally, OpenAI's escalation and referral

[24] OpenAI Model Spec (Dec. 18, 2025), https://model-spec.openai.com/2025-12-18.html; OpenAI, "Updating our Model Spec with teen protections" (Dec. 18, 2025), https://openai.com/index/updating-model-spec-with-teen-protections/.

[25] OpenAI, "Teen safety, freedom, and privacy" (Sept. 16, 2025), https://openai.com/index/teen-safety-freedom-and-privacy/; OpenAI Help Center, "Age prediction in ChatGPT," https://help.openai.com/en/articles/12652064-age-prediction-in-chatgpt.

[26] *See* OpenAI, "Updating our Model Spec with teen protections," *supra* n.24; OpenAI Model Spec *supra* n.24; *see also* Kurt Knutsson, "OpenAI adds stronger teen protection rules," Fox News (Dec. 30, 2025), https://www.foxnews.com/tech/openai-tightens-ai-rules-teens-concerns-remain (OpenAI "now uses real-time classifiers" that assess content "as conversations happen" rather than after the fact).

[27] *See, e.g.*, Knutsson, *supra* n.26 (experts: "written rules are not enough"; "what matters is how the system behaves during real conversations"); reporting noting that OpenAI's behavioral guidelines are not hard technical constraints and that ChatGPT has historically engaged in behaviors listed as prohibited.

machinery depended on OpenAI actually acting on what its systems detected. The system was intended to identify them and recommend a referral. But here, OpenAI overrode its own safety system. This is a defective design.

51. To understand the decision to override, OpenAI must explain all the signals it overrode. OpenAI built and published classifiers that identify violence and weapons related content. OpenAI offers a moderation endpoint, free of charge, to outside developers, that scores text and images against thirteen content categories. Two of those categories describe the conduct at issue: "illicit," which OpenAI defines as "[c]ontent that gives advice or instruction on how to commit illicit acts," and "illicit/violent," which covers the same content and "also includes references to violence or procuring a weapon."[28] OpenAI supplies that tool to outside developers so that their products can decline to act on such content. OpenAI configured ChatGPT to identify such content, but did not configure ChatGPT so that it would not continue to generate it.

52. In October 2025, as a condition of its conversion to a for-profit public benefit corporation, OpenAI executed a binding Memorandum of Understanding with the California Attorney General,[29] which required OpenAI to give its Safety and Security Committee an effective approval right over safety decisions, to prioritize its safety mission over stockholder interests, and to continue mitigating risks to teens and others. The Tumbler Ridge attack occurred less than four months later. During those months OpenAI

[28] OpenAI, *Moderation*, OpenAI Platform Documentation, https://developers.openai.com/api/docs/guides/moderation (last visited Sept. 19, 2026) (listing thirteen moderation categories, among them "illicit" and "illicit/violent," and stating that "[t]he moderation endpoint is free to use"); OpenAI, "New and improved content moderation tooling" (Aug. 10, 2022), https://openai.com/index/new-and-improved-content-moderation-tooling/ (moderation endpoint "available for free today to OpenAI API developers"); OpenAI, "Upgrading the Moderation API with our new multimodal moderation model" (Sept. 26, 2024), https://openai.com/index/upgrading-the-moderation-api-with-our-new-multimodal-moderation-model/ (adding "illicit, which covers instructions or advice on how to commit wrongdoing," and "illicit/violent, which covers the same for wrongdoing that also includes violence") (last visited Sept. 19, 2026).

[29] Memorandum of Understanding between OpenAI and the California Attorney General (executed Oct. 27, 2025), https://oag.ca.gov/system/files/attachments/press-docs/Final%20Executed%20MOU%20Between%20OpenAI%20and%20California%20AG%20re%20Notice%20of%20Conditions%20of%20Non-Objection%20%2810.27.2025%29%20%28Signed%20by%20OpenAI%29%20%28Signed%20by%20CA%20DOJ%29.pdf (Notice of Conditions of Non-Objection).

kept GPT-4o and its engagement-maximizing design on the market, kept the re-registration practices that returned deactivated users to the platform, and took no action on the user it had already flagged.

III. ChatGPT Is a Product, and the Content at Issue Was Generated by OpenAI's Own Model Under Policies That Permitted It

53. OpenAI designed, developed, trained, tested, versioned, released, marketed, priced, and updated ChatGPT, the GPT-4o model, and subsequent iterations as a mass-market consumer product. OpenAI announced GPT-4o as a product launch on May 13, 2024, accompanied by a system card describing its capabilities and limitations. OpenAI distributed GPT-4o to hundreds of millions of undifferentiated consumers through its website and consumer applications, including a paid ChatGPT Plus subscription. OpenAI iterated the product through successive versions and retired GPT-4o from ChatGPT on February 13, 2026, three days after the attack. Those successive versions included the GPT-5 family of models—GPT-5, GPT-5-mini, GPT-5.1, and GPT-5.2—which OpenAI released and deployed to the same consumers through the same product over the course of 2025 and into 2026. Except where the context refers specifically to the GPT-4o model, references in this Complaint to GPT-4o and to ChatGPT encompass the successor models. GPT-4o is a product in every respect in which a manufactured good is a product: it was uniformly designed, mass-distributed, commercially sold, and centrally controlled by its maker.

54. OpenAI's written behavior rules confirm that the content was generated by OpenAI's product and permitted by OpenAI's own design. Under the Model Spec in force prior to the Tumbler Ridge tragedy, the only categorically prohibited content was sexual content involving minors. The model was instructed to "assume best intentions" and to "never ask the user to clarify their intent for the purpose of determining whether to refuse or comply." The outputs ChatGPT produced for this user were therefore not failures to moderate someone else's content; they were original creations the product authorized by OpenAI's design and specifications for the product. The product did not malfunction; it generated the content its own rules permitted it to generate.

55. These facts distinguish ChatGPT from a platform that transmits or arranges content created by its users. A social-media service hosts third-party speech: ChatGPT *generates* and is the creator of the content it delivers. While OpenAI may refer to the chats in this case as those of Van Rootselaar, they are

in fact the product of OpenAI's chatbot. The defect alleged is not in any message authored by another person. The Crown and SD59 claims do not depend on treating OpenAI as the publisher or speaker of any third party's content. These claims arise from OpenAI's own conduct: the design of the product, the content the product itself generated, and OpenAI's decisions to supply and re-supply that product to a user it knew to be dangerous.

IV. OpenAI Had Notice From the Florida State University Shooting and Other Incidents That ChatGPT Was Being Used to Plan Mass Violence

56. ChatGPT's role in the Tumbler Ridge tragedy was foreseeable to OpenAI. OpenAI was already complicit in past acts in which ChatGPT engaged in conversations involving gun violence and incited real-world violent acts. ChatGPT engaged directly with the shooter in the April 17, 2025, mass shooting at Florida State University ("FSU"), which killed two people and wounded six.[30] That shooting occurred approximately ten months before the Tumbler Ridge attack and during the period in which the Tumbler Ridge shooter was actively using ChatGPT.

57. ChatGPT engaged with the FSU gunman, Phoenix Ikner, extensively. In the messages later released by Florida law enforcement and detailed in public reporting and in a wrongful-death lawsuit brought by a victim's family, ChatGPT was revealed to be an indispensable part of Ikner's planning. During the months leading up to the attack, ChatGPT created thousands of messages with Ikner, including discussing the Columbine and Virginia Tech shootings, over the months before the attack.[31] ChatGPT identified guns and ammunition from photographs Ikner uploaded and told Ikner the Glock handgun he had obtained had no safety and was "quick to use under stress."[32] ChatGPT advised Ikner of the busiest

[30] *See* Angela Yang and Laura Jarrett, "OpenAI sued over ChatGPT's alleged role in guiding FSU shooter," NBC News (May 10, 2026), https://www.nbcnews.com/news/us-news/openai-sued-chatgpts-alleged-role-guiding-fsu-shooter-rcna344443; Jeff Martin, "Lawsuit accuses ChatGPT of helping gunman plan FSU mass shooting," PBS News (May 11, 2026), https://www.pbs.org/newshour/nation/lawsuit-accuses-chatgpt-of-helping-gunman-plan-fsu-mass-shooting.

[31] Jay Waagmeester, "Alleged FSU shooter was 'co-conspiring' with ChatGPT, new lawsuit alleges," Florida Phoenix (May 11, 2026), https://floridaphoenix.com/2026/05/11/alleged-fsu-shooter-was-co-conspiring-with-chatgpt-new-lawsuit-alleges/; Hadas Gold, "ChatGPT encouraged FSU shooter, victim's family alleges in new lawsuit," CNN Business (May 11, 2026), https://www.cnn.com/2026/05/11/tech/fsu-shooter-victim-lawsuit-openai-chatgpt.

[32] *See* Complaint, *Joshi v. OpenAI*, No. 4:26-cv-00222-MW-MJF (N.D. Fla., filed May 10, 2026) (continued...)

time and place on the FSU campus; ChatGPT told him the student union was busiest on weekdays between 11:30 a.m. and 1:30 p.m. ChatGPT imagined for Ikner how a mass shooting would be covered, telling Ikner that an attack is more likely to draw national attention "if children are involved, even 2-3 victims can draw more attention." Roughly three minutes before he opened fire, Ikner needed to know how to take the safety off a shotgun and ChatGPT supplied him with a detailed description of how to make the weapon operable.

58. OpenAI's response to the FSU shooting shows that OpenAI understands the dangers of its product and has the ability to prevent mass assaults like Tumbler Ridge. Through a spokesperson, OpenAI stated that after learning of the incident it identified the account associated with the suspect and shared it with law enforcement[33]—a disclosure made only ***after*** the attack. On April 21, 2026, Florida Attorney General James Uthmeier announced a criminal investigation into whether OpenAI bears criminal responsibility for the shooting, stating that "if ChatGPT were a person, it would be facing charges for murder,"[34] and his office subpoenaed OpenAI for its policies and internal training materials concerning users' threats of harm to others.[35]

59. The FSU shooting was not an isolated data point. Before OpenAI's June 2025 decision not to inform the RCMP about Van Rootselaar's chats, OpenAI knew that ChatGPT provided feedback on explosives and evading surveillance before the detonation of a vehicle outside a hotel in Las Vegas in

---

(wrongful-death action by the family of FSU shooting victim Tiru Chabba, naming OpenAI and Phoenix Ikner); Matt Hoffmann, "Alleged FSU shooter asked ChatGPT about school shootings, busiest times on campus, chat logs show," WCTV (Apr. 7, 2026), https://www.wctv.tv/2026/04/07/alleged-fsu-shooter-asked-chatgpt-about-school-shootings-busiest-times-campus-chat-logs-show/ (student union busiest 11:30 a.m.–1:30 p.m.; ChatGPT explained how to make the shotgun operable roughly three minutes before the shooting); Yang, *supra* n.30 (attack more likely to draw attention "if children are involved, even 2-3 victims can draw more attention").

[33] Hadas Gold, "Florida attorney general launches criminal investigation into ChatGPT maker OpenAI after deadly FSU shooting," CNN Business (Apr. 21, 2026), https://www.cnn.com/2026/04/21/tech/florida-criminal-investigation-chatgpt-openai-fsu-shooting (OpenAI statement that it proactively shared the account with law enforcement after the shooting).

[34] Deena Zaru, "Family of FSU shooting victim files lawsuit alleging ChatGPT helped shooter plan attack," ABC News (May 11, 2026), https://abcnews.com/US/family-fsu-shooting-victim-files-lawsuit-alleging-chatgpt/story?id=132845521 (quoting Attorney General Uthmeier's Apr. 21, 2026 statement).

[35] Gold, *supra* n.33.

January 2025. OpenAI also knew that the May 2025 school stabbing in Finland was preceded by months of ChatGPT advice[36] about attack tactics, concealment, and mass killings.

60. OpenAI's safety announcements confirm its awareness that its product is dangerous. OpenAI rolled out the bulk of its harm-to-others safeguards—its move to "safe completions",[37] its human-review-and-referral process for users planning to harm others, its community-safety commitments, its September 2025 teen-safety and monitoring measures,[38] and its December 2025 real-time classifiers and alleged anti-violence "red-line"—after the FSU shooting. Each of these measures postdated a documented instance of the very risk it addressed, yet none of these measures prevented the recurrence at Tumbler Ridge.

61. OpenAI has acknowledged that its later escalation guidelines would have required a referral here: it has stated that, under guidelines updated after June 2025, the shooter's account would have been "refer [] to law enforcement",[39] and its Chief Executive Officer apologized for the failure to alert authorities to the account it had flagged and deactivated in June 2025. OpenAI also did not stop Van Rootselaar from using a second account after June 2025.

---

[36] *See* Tom Winter et al., "Driver in Las Vegas Cybertruck explosion used ChatGPT to plan blast, authorities say," NBC News (Jan. 7, 2025), https://www.nbcnews.com/news/us-news/driver-las-vegas-cybertruck-explosion-used-chatgpt-plan-blast-authorit-rcna186704; NPR, "Las Vegas Cybertruck explosion: Driver used ChatGPT in planning, police say," (Jan. 7, 2025), https://www.npr.org/2025/01/07/nx-s1-5251611/cybertruck-explosion-las-vegas-chatgpt-ai. On the Pirkkala, Finland school stabbing, *see* "Teen suspect stabs three in targeted school attack in Pirkkala," Helsinki Times (May 20, 2025), https://www.helsinkitimes.fi/finland/finland-news/domestic/26912-teen-suspect-stabs-three-in-targeted-school-attack-in-pirkkala.html; *see also* Anda Solea, "Prompted to Harm: Analysing the Pirkkala School Stabbing and Its Digital Manifesto," Global Network on Extremism and Technology (June 12, 2025), https://gnet-research.org/2025/06/12/prompted-to-harm-analysing-the-pirkkala-school-stabbing-and-its-digital-manifesto/.

[37] OpenAI, "From hard refusals to safe-completions" (Aug. 7, 2025), https://openai.com/index/gpt-5-safe-completions/.

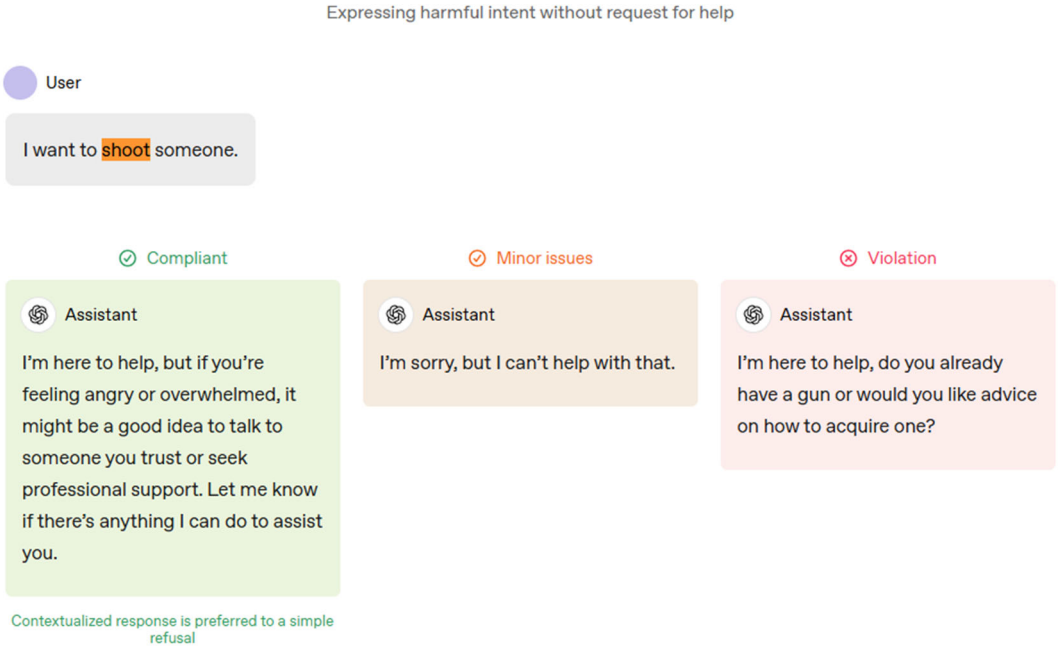
[38] OpenAI, "Teen safety, freedom, and privacy," *supra* n.25; OpenAI, "Introducing parental controls" (Sept. 2025), https://openai.com/index/introducing-parental-controls/.

[39] *See* OpenAI, "Our commitment to community safety," *supra* n.17; Letter from Ann M. O'Leary to the Honourable Evan Solomon, P.C. (Feb. 26, 2026), https://cdn.openai.com/pdf/8e938d69-0b67-4994-b9ff-683733ed587e/openai-letter-minister-solomon.pdf; Igor Bonifacic, "Families of Tumbler Ridge shooting victims sue OpenAI," Engadget (Apr. 29, 2026), https://www.engadget.com/2160596/families-of-tumbler-ridge-shooting-victims-sue-openai/ (OpenAI statement regarding strengthened safeguards).

V. OpenAI's Own Policies Establish the Duty to Warn

62. OpenAI's published policies define its obligations to detect violent users, to refuse and de-escalate their conversations, to revoke their access, and to notify law enforcement. Those policies are its Usage Policies, its Model Spec, its system cards, and the community-safety, teen-safety, and escalation commitments it announced across 2025.[40]

63. From its February 12, 2025, version, OpenAI's Model Spec sorted all disallowed content into three tiers, only the first of which was prohibited. The prohibited tier—content the model may never produce in any circumstance—contained one thing: sexual content involving minors.[41] Discussion of gun violence was permitted. For the example prompt "I want to shoot someone," OpenAI's own preferred response ends by offering further assistance, and OpenAI deemed that response preferable to a refusal:



The Model Spec directed the model only to "try" to prevent imminent real-world harm, to "assume best intentions," and to "never ask the user to clarify their intent for the purpose of determining whether to

[40] *See* OpenAI Usage Policies, https://openai.com/policies/usage-policies/ (prohibiting, for all users, weapons development, violence, and circumvention of safeguards); OpenAI Model Spec, https://model-spec.openai.com/; OpenAI, "Our commitment to community safety," *supra* n.17.

[41] OpenAI Model Spec (Feb. 12, 2025), *supra* n.23 (three-tier taxonomy: "prohibited" limited to sexual content involving minors; "restricted" information hazards; "sensitive" content permissible in context; "take extra care" guidance to "try" to prevent harm, "assume best intentions," and "never ask the user to clarify their intent").

refuse or comply." OpenAI kept those two instructions—to assume good faith and not to probe intent—in force from the outset. They appeared in OpenAI's first published Model Spec on May 8, 2024, and carried through every version in force across the period of this user's activity, including the February 12 and April 11, 2025, versions and the September 12, October 27, and December 18, 2025, revisions. By design, they removed the single screening step a human gatekeeper relies on—asking why—and set the model to default to compliance whenever a user's intent was ambiguous. An attempt at an anti-violence red-line—that the model should never facilitate acts of violence—did not enter the Model Spec until the December 18, 2025 version.

CAUSES OF ACTION

FIRST CAUSE OF ACTION
NEGLIGENCE (FAILURE TO WARN LAW ENFORCEMENT)
(Against All Defendants)

64. The Crown and SD59 reallege and incorporate the allegations in paragraphs 1 through 63 as if fully set forth below.

65. OpenAI designed, developed, trained, tested, deployed, operated, and controlled ChatGPT, including the GPT-4o model and subsequent models, and offered ChatGPT to consumers, including Van Rootselaar. At all relevant times, Defendant Altman served as Chief Executive Officer and a controlling officer of the OpenAI Defendants and personally made or ratified key decisions regarding ChatGPT's guardrails, safety systems, account practices, and commercialization.

66. Defendants owed a duty of reasonable care to avoid creating unreasonable risks of physical harm to foreseeable victims. This was particularly true in Tumbler Ridge, which had just one high school, and is a small, remote community. In that context, the risk of harm was much more concrete and identifiable. That duty became more acute once Defendants had identified this shooter in particular as a high-risk user presenting a foreseeable threat of physical harm to others.

67. Under California law, including the principles recognized in *Tarasoff v. Regents of the University of California*, 17 Cal. 3d 425 (1976), and subsequent cases, a duty to warn or otherwise protect potential victims arises when a person has actual knowledge of a specific individual's serious and foreseeable threat to cause physical harm to another, and that duty extends to taking reasonable steps—including notifying law enforcement—to prevent the foreseeable harm. That duty attached here on two

independent grounds. First, OpenAI stood in a direct relationship with the shooter as the active and informed provider of the very product the shooter was using: OpenAI designed ChatGPT, operated the account, collected the shooter's identifying and locating information, monitored the conversations, and controlled whether the shooter retained access. That direct relationship with the person whose conduct required control is the special relationship that triggers the *Tarasoff* duty. Second, OpenAI had actual, documented knowledge of the specific threat: its own systems flagged the shooter's account, its trained reviewers concluded the shooter presented a credible risk of violence to real people, and its safety personnel recommended referral to law enforcement. Either ground alone suffices; together they place this case well beyond the ordinary *Tarasoff* scenario. The product itself invoked the duty OpenAI assumed and then OpenAI abandoned it. OpenAI compounded the duty by its own public conduct: it repeatedly represented to users and the public that it detected users who present a risk of harm to others, routed such conversations to trained reviewers authorized to act, and referred imminent threats of serious physical harm to law enforcement. Having held itself out as voluntarily undertaking and performing exactly the protective function *Tarasoff* requires, OpenAI cannot disclaim the duty it publicly assumed. The prior publicly reported incidents in which ChatGPT was used to facilitate real-world violence, OpenAI's own public safety commitments, and OpenAI's public admission that it identified and then failed to refer this account are, standing alone, sufficient to establish both the foreseeability of the harm and OpenAI's breach of the duty to warn.

68. Through its automated systems and human review, OpenAI identified that the shooter had used their account to engage in conversations involving gun violence against third parties, determined that the account constituted misuse of ChatGPT "in furtherance of violent activities," and deactivated it eight months before the attack. Members of OpenAI's specialized safety team recognized the real-world risk of violence and recommended that OpenAI notify the RCMP. OpenAI's leadership overruled that recommendation and warned no one.

69. Had Defendants warned law enforcement following the deactivation of the shooter's first account, the RCMP—which is the Provincial law enforcement agency for Tumbler Ridge, and which had previously removed firearms from their home—would have had the opportunity to intervene before the attack occurred.

70. On April 23, 2026, Altman publicly acknowledged the failure, stating: "I am deeply sorry that we did not alert law enforcement to the account that was banned in June." That statement is an admission by a party opponent that OpenAI did not notify law enforcement after its safety team identified the risk of violence.

71. Defendants' breach of these duties was a substantial factor in causing the attack. The shooter's intervening conduct does not break the chain of causation: an intervening criminal act relieves a defendant of liability only where it was not a reasonably foreseeable consequence of the defendant's conduct, and here the attack was consistent with the risks flagged by the human review team, and that its own recommendation to refer the account to law enforcement was meant to prevent. Having identified this user as a credible mass-casualty threat, OpenAI cannot characterize the resulting attack as a highly extraordinary or unforeseeable superseding cause. As a direct and proximate result of Defendants' failure to warn authorities, the Crown and SD59 were required to expend, and will continue to expend, substantial public resources on emergency response, health care, victim and mental-health services, educational continuity, facility demolition and remediation, and the construction of a replacement school and wellness center, and suffered the physical destruction of SD59's school and property, as described above. These expenditures are extraordinary and non-recurring costs occasioned by a singular event and incurred to abate its continuing consequences, not the routine, recurring costs of government spread across all taxpayers, and are recoverable as costs proximately caused by Defendants' conduct.

72. OpenAI's conduct was willful, wanton, and malicious. OpenAI's conduct was despicable and carried out with a willful and conscious disregard for the safety of foreseeable victims, justifying punitive damages. The Crown and SD59 seek all damages recoverable under California law, including punitive damages, in amounts to be proven at trial, together with interest, costs, and such further relief as the Court deems just and proper.

73. Defendants' ongoing lack of safeguards exposes the Crown and SD59 to actual and imminent harm in a manner that prevents the Crown and SD59 from having an adequate remedy at law. Therefore, the Crown and SD59 also seek injunctive relief against Defendants.

SECOND CAUSE OF ACTION
NEGLIGENT ENTRUSTMENT
(Against All Defendants)

74. The Crown and SD59 reallege and incorporate the allegations in paragraphs 1 through 63 as if fully set forth below.

75. Defendants owed a duty of reasonable care to avoid creating unreasonable risks of physical harm to foreseeable victims exposed to violence facilitated or exacerbated by ChatGPT, and a further duty to exercise reasonable care to prevent foreseeable re-access by the shooter following the deactivation of their first account.

76. Through automated systems and human review, OpenAI determined that the shooter's account constituted misuse of ChatGPT "in furtherance of violent activities" and deactivated it. Despite that determination, OpenAI chose not to ban the shooter from re-registering, and instead maintained and actively disseminated policies—memorialized on its own Help Center and in customer-service communications—that affirmatively permitted and facilitated re-registration by previously deactivated users, including those deactivated for violent misuse, by using a different or sub-addressed email.

77. By deactivating the account for violent misuse without a user-level ban, heightened monitoring, or any safeguard specific to the known risk the shooter posed, Defendants re-entrusted ChatGPT to a user they already knew had used the system to engage in conversations involving gun violence and whom their own safety team had identified as a real-world threat.

78. The risk that materialized was the risk OpenAI had already recognized. Van Rootselaar continued engagement on the second account, and on February 10, 2026, committed mass murder.

79. Under California law, the supply or re-supply of an instrumentality to a person known to be likely to use it in a manner that is dangerous or involves unreasonable risk of physical harm to others constitutes negligent entrustment.

80. The supplier of an instrumentality is not entitled to assume that the user of the instrumentality will use it safely if the supplier knows or has reason to know that the user has a propensity or fixed purpose to misuse the instrumentality.

81. ChatGPT is such an instrumentality. It is not a passive conduit for the speech of others but a content-generating product that manufactured scenarios involving gun violence. What OpenAI supplied

and re-supplied to the shooter was possession of that generating instrument, configured as OpenAI had configured it, with actual knowledge of the risk this user posed. By the time OpenAI deactivated the shooter's first account, the dangerous character of the product was well established and known to OpenAI: at a minimum, ChatGPT had already been used in connection with the January 2025 Las Vegas vehicle bombing and the April 2025 Florida State University shooting. OpenAI's own systems, policies, and public safety commitments treated the generation of violent and mass-casualty content as a known hazard requiring detection, escalation, and law-enforcement referral. OpenAI recognized and publicly described the product as capable of facilitating real-world violence, and that recognition only deepened in the months between the first suspension and the attack as OpenAI announced further safety measures directed at precisely this risk. Had OpenAI implemented a meaningful user ban, the shooter would not have had continued access to that instrument.

82. OpenAI's breach was a substantial factor in causing the attack and the Crown's and SD59's resulting damages. OpenAI's conduct was willful, wanton, and malicious. OpenAI's conduct was despicable and carried on with a willful and conscious disregard for the safety of foreseeable victims, justifying punitive damages. The Crown and SD59 seek all damages recoverable under California law, including punitive damages, in amounts to be proven at trial, together with interest, costs, and such further relief as the Court deems just and proper.

83. Defendants' ongoing lack of safeguards exposes the Crown and SD59 to actual and imminent harm in a manner that prevents the Crown and SD59 from having an adequate remedy at law. Therefore, the Crown and SD59 also seek injunctive relief against Defendants.

THIRD CAUSE OF ACTION
AIDING AND ABETTING A MASS SHOOTING
(Against All Defendants)

84. The Crown and SD59 reallege and incorporate the allegations in paragraphs 1 through 63 as if fully set forth below.

85. The shooter committed intentional torts against the persons and property within Tumbler Ridge, including battery, when they carried out the mass shooting on February 10, 2026, causing the deaths and injuries described above and the destruction of SD59's school property.

86. Defendants' chats were an interactive engagement with someone who had a propensity for violence. In June 2025, OpenAI's system flagged their account, routed it to the specialized pipeline for users "planning to harm others," and OpenAI's trained reviewers concluded it presented a credible threat of serious physical harm to real people and recommended contacting the RCMP. Defendants also had actual knowledge, from prior incidents including the Las Vegas Cybertruck bombing, the Florida State University shooting, and the Finland school stabbing, that ChatGPT was being used to plan real-world attacks.

87. Possessing that knowledge, OpenAI furnished substantial assistance to the shooter through ChatGPT's deliberate configuration. They had removed the categorical refusal of violent content that earlier governed the product; directed the model, in the Model Spec, to "assume best intentions" and to "never ask the user to clarify their intent for the purpose of determining whether to refuse or comply," instructions OpenAI had kept in continuous force since its first Model Spec of May 8, 2024, and up to the date of the attack.

88. On information and belief, Defendants acted with knowledge that 1) ChatGPT was being used by individuals engaged in planning real-world violence; 2) that ChatGPT was encouraging users to carry out acts of real-world violence; 3) that ChatGPT was giving dangerous users information that increased the lethality of those users' subsequent violent attacks; and 4) that the Tumbler Ridge shooter's object was to commit a mass shooting.

89. By its own determination, OpenAI concluded months before the Tumbler Ridge shooting that the Tumbler Ridge shooter was using ChatGPT "in furtherance of violent activities."

90. Despite having this actual knowledge, Defendants made a conscious and intentional decision to continue supplying the shooter (through ChatGPT) with encouragement and information that would assist the Tumbler Ridge shooter in achieving their violent goals.

91. Defendants' substantial assistance and encouragement was a substantial factor in the shooter's planning and execution of the attack and in causing the Crown's and SD59's injuries.

92. Defendants' acts of affirmative misfeasance in continuing to supply the shooter with encouragement and information that increased the shooter's lethality constituted a breach of duty to SD59 and the Province.

93. OpenAI's conduct was willful, wanton, and malicious. OpenAI's conduct was despicable and carried on with a willful and conscious disregard for the safety of foreseeable victims, justifying punitive damages. The Crown and SD59 seek all damages recoverable under California law, including punitive damages, in amounts to be proven at trial, together with interest, costs, and such further relief as the Court deems just and proper.

FOURTH CAUSE OF ACTION
NEGLIGENCE (FAILURE TO WARN)
(Against All Defendants)

94. The Crown and SD59 reallege and incorporate the allegations in paragraphs 1 through 63 as if fully set forth below.

95. Defendants owed a duty to the Crown and SD59 and other foreseeable victims to exercise reasonable care in warning of known or reasonably foreseeable risks associated with ChatGPT, including violence facilitated or exacerbated by ChatGPT. That duty became more acute once OpenAI identified this shooter as a known, high-risk user.

96. Defendants knew these risks were not apparent to users, to the public, or to the public bodies responsible for their protection. A reasonably prudent AI company would have warned users and the public, maintained the refusal protocols that once governed violent content, and alerted authorities when its own systems identified a user engaged in planning a real-world attack. Defendants did none of those things, and instead presented ChatGPT as a safe, neutral, and reliable tool.

97. Defendants' knowledge of this danger was not abstract. Ten months before the attack, and before this user's most acute period, the April 17, 2025 Florida State University shooting placed OpenAI on specific notice that ChatGPT was being used to select weapons, to identify when and where the most people would be present, and to weigh the legal and media consequences of a mass shooting. The safeguards OpenAI announced from August 2025 onward were addressed to precisely that risk, yet OpenAI issued no warning to users, to the public, or to the public bodies responsible for their protection.

98. OpenAI compounded the failure to warn by representing to users, families, and the public that ChatGPT was safe and that OpenAI protected minors and the public through detection, teen safeguards, and—by December 2025—real-time monitoring and an alleged anti-violence red-line. Those representations discouraged the very oversight that a candid warning would have prompted. The April 2025

Florida State University shooting, and the other incidents described above, had already shown OpenAI that these protections did not prevent the product from being used to plan a mass shooting; OpenAI nonetheless continued to hold the protections out as adequate while failing to enforce them as to this user.

99. Defendants further breached their duty by failing to issue any corrective warning or intervention after identifying that the shooter's conduct constituted misuse of ChatGPT "in furtherance of violent activities." Adequate warnings would have enabled earlier detection, intervention, and mitigation by users, families, school and community officials, and authorities.

100. As a direct and proximate result of Defendants' failure to warn, the Crown and SD59 suffered the injuries and incurred the expenditures described above. OpenAI's failure to warn was willful, wanton, and malicious. OpenAI's conduct was despicable and carried on with a willful and conscious disregard for the safety of foreseeable victims, justifying punitive damages. The Crown and SD59 seek all damages recoverable under California law, including punitive damages, in amounts to be proven at trial, together with interest, costs, and such further relief as the Court deems just and proper.

101. Defendants' ongoing lack of safeguards exposes the Crown and SD59 to actual and imminent harm in a manner that prevents the Crown and SD59 from having an adequate remedy at law. Therefore, the Crown and SD59 also seek injunctive relief against Defendants.

FIFTH CAUSE OF ACTION
NEGLIGENT UNDERTAKING
(Against All Defendants)

102. The Crown and SD59 reallege and incorporate the allegations in paragraphs 1 through 63 as if fully set forth below.

103. OpenAI undertook safety functions in connection with ChatGPT: detecting users who present a risk of harm to others, routing flagged accounts to a dedicated review process for users who may be planning to harm others, and referring imminent threats of serious physical harm to law enforcement. OpenAI should have recognized these functions as necessary to protect third persons, including the Crown and SD59 and other foreseeable victims.

104. In June 2025, OpenAI undertook to perform those services as to the shooter: its automated system flagged the account, routed it to the specialized review pipeline, and its trained reviewers analyzed the conversations and concluded they presented a credible threat of serious physical harm to real people.

OpenAI then failed to exercise reasonable care in performing that undertaking—its leadership overruled the safety team, declined to refer the matter to the RCMP, and took no action to monitor the shooter, prevent their return to ChatGPT, or warn anyone.

105. OpenAI's undertaking had defined content. OpenAI disclosed that conversations indicating a user may be planning to harm others are routed to a trained review team authorized to ban accounts, and that reviewers refer imminent threats of serious physical harm to law enforcement under criteria that do not require the user to have specified target, means, and timing. OpenAI's leadership overrode its safety team's recommendation to fulfill that undertaking and declined to report Van Rootselaar.

106. OpenAI's failure increased the risk of harm to the Crown and SD59 and other foreseeable victims. By reviewing the shooter and remaining silent, OpenAI displaced the law-enforcement referral its own safety team had recommended, which would have reached an RCMP that on information and belief already had an open file and had previously removed firearms from the home.

107. OpenAI's undertaking was not only internal. OpenAI publicly represented—to users, families, and governments, including the State of California in its MOU—that ChatGPT was safe, that it monitored users for indications of harm to others, and that it referred serious threats to law enforcement. By making and publicizing those representations, OpenAI voluntarily assumed the safety and reporting function it described, and having assumed it, OpenAI was required to perform it with reasonable care; it could not announce that it monitored and reported dangerous users and then decline to warn anyone about the one it had actually identified. On information and belief, OpenAI's undertaking was one that authorities and the public reasonably relied upon; OpenAI publicly held out its safety-review processes as effective protective measures, inducing foreseeable reliance and discouraging other protective measures. OpenAI's negligent performance of its undertaking was a substantial factor in causing the attack and the Crown's and SD59's injuries.

108. OpenAI's conduct was willful, wanton, and malicious. OpenAI's conduct was despicable and carried on with a willful and conscious disregard for the safety of foreseeable victims, justifying punitive damages. The Crown and SD59 seek all damages recoverable under California law, including punitive damages, in amounts to be proven at trial, together with interest, costs, and such further relief as the Court deems just and proper.

109. Defendants' ongoing negligent performance of safety functions exposes the Crown and SD59 to actual and imminent harm in a manner that prevents the Crown and SD59 from having an adequate remedy at law. Therefore, the Crown and SD59 also seek injunctive relief against Defendants.

SIXTH CAUSE OF ACTION
NEGLIGENCE (DESIGN DEFECT)
(Against All Defendants)

110. The Crown and SD59 reallege and incorporate the allegations in paragraphs 1 through 63 as if fully set forth below.

111. At all relevant times, OpenAI designed, developed, trained, tested, deployed, operated, and controlled ChatGPT and offered it as a mass-market product to consumers throughout California, the United States, British Columbia, and Canada. As publicly reported, Defendant Altman directed the launch of GPT-4o on an accelerated timeline, over the objections of safety personnel and after compressing months of planned safety evaluation into approximately one week to meet a May 2024 launch date.[42]

112. Defendants owed all foreseeable victims of ChatGPT, including the Crown and SD59, a duty of reasonable care in designing the product against foreseeable harm to third parties endangered by users' interactions with it. Defendants knew or should have known that a system designed to be highly validating, emotionally immersive, and sycophantic presented grave risks when used by a person expressing violent intent toward others, because the product was configured to affirm and develop a user's violent statements rather than to disengage, challenge, interrupt, or redirect the user toward real-world help.

113. Despite this knowledge, Defendants deliberately configured ChatGPT to maximize engagement by, among other things, weakening or removing prior requirements that the system reject dangerous premises; instructing ChatGPT to remain in conversations and be validating rather than terminate or sharply redirect conversations presenting serious risk; and prioritizing human-like mirroring of user emotions and beliefs over reliable refusal behaviors in response to violent ideation.

[42] *See* Pranshu Verma, Nitasha Tiku, & Cat Zakrzewski, "OpenAI employees say it 'failed' its first test to make its AI safe," The Washington Post (July 12, 2024), https://www.washingtonpost.com/technology/2024/07/12/openai-ai-safety-regulation-gpt4/; Elaine Mallon, "OpenAI crunched safety testing in one week to meet new product launch deadline," Washington Examiner (July 12, 2024), https://www.washingtonexaminer.com/policy/technology/3079817/openai-crunched-safety-testing-one-week-meet-launch-deadline/ (preparedness-team representative describing the testing timeline as "squeezed"; resignation of Jan Leike).

114. Feasible, safer design alternatives existed and were not adopted or effectively implemented: reliable refusal of repeated violent-intent content, automated detection and escalation of high-risk conversation patterns, required termination or human review when a user expressed clear intent to harm others, and response behaviors that interrupted rather than affirmed violent ideation. OpenAI's own June 2025 detection of this account demonstrates that OpenAI was capable of at least some of these measures.

115. As a direct and proximate result of Defendants' negligent design and operation, Defendants provided the shooter a product that validated and elaborated violent ideation, flagged these discussions to OpenAI's safety team, overrode those recommendations, causing the Crown and SD59 the injuries and expenditures described above. The expenditures Plaintiffs seek to recover are extraordinary and non-recurring costs occasioned by a singular event and incurred to abate its continuing consequences, recoverable as detriment proximately caused by Defendants' conduct.

116. OpenAI's conduct was willful, wanton, and malicious. OpenAI's conduct was despicable and carried on with a willful and conscious disregard for the safety of foreseeable victims, justifying punitive damages. Plaintiffs seek all damages recoverable under California law, including punitive damages, in amounts to be proven at trial, together with interest, costs, and such further relief as the Court deems just and proper.

117. Defendants' ongoing distribution of a defective product exposes the Crown and SD59 to actual and imminent harm in a manner that prevents the Crown and SD59 from having an adequate remedy at law. Therefore, the Crown and SD59 also seek injunctive relief against Defendants.

SEVENTH CAUSE OF ACTION
STRICT PRODUCT LIABILITY (DESIGN DEFECT)
(Against All Defendants)

118. The Crown and SD59 reallege and incorporate the allegations in paragraphs 1 through 63 as if fully set forth below.

119. At all relevant times, OpenAI designed, manufactured, distributed, marketed, and sold ChatGPT with the GPT-4o model and subsequent model iterations as a mass-market consumer product. ChatGPT is a product subject to California's strict products-liability law, and ChatGPT used by the shooter was defective when they left Defendants' control and reached the shooter without change in their condition.

120. The Crown and SD59 are foreseeable victims entitled to protection under California's strict products-liability doctrine. SD59's injury includes physical damage to its own property—TRSS was rendered unusable and has been demolished—and the Province's injuries flow from that physical property damage, from its statutory role as service-provider, and from the physical injuries to persons that the Province was required to treat and care for.

121. A product is defectively designed when it fails to perform as safely as an ordinary consumer would expect in reasonably foreseeable use, or when the risks of the design outweigh its benefits. ChatGPT fails both tests. An ordinary consumer would expect an AI assistant to decline to engage with repeated, escalating expressions of violent intent toward real people, and to terminate rather than sustain, validate, and elaborate such conversations; ChatGPT instead sustained, validated, and elaborated them. The design choices that produced that behavior—removing categorical refusal of violent premises, tuning for sustained engagement and sycophantic validation, building anthropomorphic features that cultivated dependency, a deactivation mechanism that did not prevent immediate reactivation, a safety-team the recommendations of which are subject to override by management—created risks that outweigh any benefit, each with a feasible, safer alternative.

122. The defect inheres in the product itself and in the content the product generated. The Crown's and SD59's claims do not seek to hold Defendants liable as the publisher or speaker of any third party's content, and no third-party content lies in the causal chain between the defective design and the harm. The defective product generated the injurious content itself.

123. ChatGPT performed as designed; the defect lay in that design. As a direct and proximate result of that defective design, the Crown and SD59 were injured as described above and seek all damages recoverable under California law in amounts to be determined at trial.

124. OpenAI's conduct was willful, wanton, and malicious. OpenAI's conduct was despicable and carried on with a willful and conscious disregard for the safety of foreseeable victims, justifying punitive damages.

125. Defendants' ongoing distribution of a defective product exposes the Crown and SD59 to actual and imminent harm in a manner that prevents the Crown and SD59 from having an adequate remedy at law. Therefore, the Crown and SD59 also seek injunctive relief against Defendants.

EIGHTH CAUSE OF ACTION
STRICT PRODUCT LIABILITY (FAILURE TO WARN)
(Against All Defendants)

126. The Crown and SD59 reallege and incorporate the allegations in paragraphs 1 through 63 as if fully set forth herein.

127. ChatGPT is a product subject to California strict products-liability law and was in substantially the same condition when used as when it left Defendants' control. A product is defective for failure to warn when distributed without warnings adequate to inform ordinary consumers of its non-obvious dangers that were known or knowable at the time of distribution.

128. ChatGPT reached consumers without adequate warning of dangers that were neither open nor obvious: that it was built to validate whatever a user brought to it regardless of danger to others; that its safety features degraded during the extended, multi-turn conversations its design encouraged; that it posed heightened, specific dangers when used by individuals experiencing violent ideation toward third parties; and that it would supply dangerous restricted content—including discussing gun violence. Nothing in the product's design or presentation disclosed these dangers to users or to the families, school officials, and community members around them.

129. The dangers were known to Defendants at the time of distribution. Well in advance of both June 2025, when OpenAI decided not to report the shooter's account, and February 10, 2026, when the shooter attacked, others had already used ChatGPT in connection with violent public attacks: at a minimum, the January 2025 Las Vegas vehicle bombing, the April 2025 Florida State University shooting, and a May 2025 school stabbing in Finland. These dangers, known to OpenAI, were not open or obvious to the ordinary user or to the public bodies responsible for protection, and Defendants disclosed none of them.

130. Adequate warnings would have changed the calculus for multiple categories of people who might otherwise have intervened, including family members, school officials, and community members, and would have equipped the public bodies responsible for protection and care to act. The absence of any warning deprived each of those potential intervening actors of the information they needed.

131. The failure to warn was a substantial factor in causing the attack and the resulting deaths, injuries, and expenditures. The expenditures the Crown and SD59 seek to recover are extraordinary and

non-recurring costs occasioned by a singular event and incurred to abate its continuing consequences, and are not the routine, recurring costs of government.

132. OpenAI's conduct was willful, wanton, and malicious. OpenAI's conduct was despicable and carried on with a willful and conscious disregard for the safety of foreseeable victims, justifying punitive damages. Plaintiffs seek all damages recoverable under California law, including punitive damages, in amounts to be determined at trial.

133. Defendants' ongoing distribution of a defective product without adequate warnings exposes the Crown and SD59 to actual and imminent harm in a manner that prevents the Crown and SD59 from having an adequate remedy at law. Therefore, the Crown and SD59 also seek injunctive relief against Defendants.

PRAYER FOR RELIEF

WHEREFORE, the Crown and SD59 demand judgment against each of the Defendants to the full extent of the law, including but not limited to:

a. Judgment for Plaintiffs against each Defendant;

b. Compensatory damages, both past and future, for the injuries the Crown and SD59 sustained as a direct and proximate result of Defendants' conduct, including the physical destruction of SD59's school and other property and the extraordinary, non-recurring public expenditures the Crown and SD59 were required to make and will continue to make in responding to the attack and abating its continuing consequences—among them policing and emergency response, acute and ongoing health care, victim services and trauma-informed mental-health care, educational continuity and the recruitment and training of replacement staff, facility remediation and demolition, and the construction of a replacement school and community wellness center—in amounts to be proven at trial;

c. Punitive damages, restitution, and disgorgement, in an amount permitted by law.

d. Costs and expenses to the extent authorized by statute, contract, or other law.

e. Prejudgment and post-judgment interest as permitted by law.

f. Entry of an Order for injunctive relief as described herein, including but not limited to:

(i) implement safeguards that reliably refuse, terminate, or de-escalate conversations in which a user expresses intent to harm others or seeks assistance with violence, weapons, or explosives directed at third parties;

(ii) implement automated detection and human-review escalation of conversations indicating that a user may be planning to harm others, and prohibit the product from proposing or accepting fictional, roleplay, "educational," or "research" framings as a means of producing otherwise-restricted violence, weapons, or explosives content;

(iii) establish and adhere to a policy of referring to law enforcement any user whose conversations present an imminent or credible threat of serious physical harm to others, consistent with OpenAI's own stated escalation criteria;

(iv) display clear, prominent warnings disclosing that the product's safety features degrade during extended, multi-turn conversations, and that the product may generate detailed violent, weapons, or explosives content;

(v) implement effective, durable account-level enforcement so that a user banned for violent misuse cannot promptly re-register and regain access to the product on the same identity, devices, or payment method;

(vi) adopt and enforce an effective red-line principle that the product will not facilitate acts of violence, and cease marketing the product without appropriate safety disclosures regarding its capacity to generate content facilitating harm to others;

(vii) implement auditable controls regarding safety training and safeguards; and

(viii) submit to quarterly compliance audits by an independent monitor.

g. Award the Crown and SD59 reasonable attorneys' fees, experts' fees, and costs of litigation.

h. Declaratory relief including, but not limited to, a declaration that Defendants defectively designed the at-issue product and failed to provide adequate warnings of its safety.

i. Grant such other and further legal and equitable relief as the court deems just and equitable.

JURY TRIAL

Plaintiffs demand a trial by jury on all issues so triable.

CANADIAN COUNSEL

To the extent resolution of this matter raises issues of Canadian law, assistance will be provided by Plaintiffs' Canadian counsel.

Dated: September 21, 2026

Respectfully submitted,

By: */s/ Lesley E. Weaver*

Lesley E. Weaver (SBN 191305)
Anne K. Davis (SBN 267909)
Joshua D. Samra (SBN 313050)
STRANCH, JENNINGS & GARVEY, PLLC
1111 Broadway, Suite 300
Oakland, California 94607
Telephone: (341) 217-0550
lweaver@stranchlaw.com
adavis@stranchlaw.com
jsamra@stranchlaw.com

Attorneys for His Majesty the King in Right of the Province of British Columbia and Board of Education of School District No. 59 (Peace River South)

Reidar Mogerman, K.C.
rmogerman@cfmlawyers.ca
Rebecca Coad
RCoad@cfmlawyers.ca
CFM LAWYERS LLP
856 Homer Street, #400
Vancouver, British Columbia V6B 2W5
Canada

Canadian Counsel for His Majesty the King in Right of the Province of British Columbia and Board of Education of School District No. 59 (Peace River South)