Jay Edelson (*pro hac vice* admission to be sought)
jedelson@edelson.com
Rafey Balabanian (SBN 315962)
rbalabanian@edelson.com
Ari Scharg (*pro hac vice* admission to be sought)
ascharg@edelson.com
EDELSON PC
350 North LaSalle Street, 14th Floor
Chicago, Illinois 60654
Tel: 312.589.6370

Todd Logan (SBN 305912)
tlogan@edelson.com
Brandt Silverkorn (SBN 323530)
bsilverkorn@edelson.com
EDELSON PC
150 California Street, 18th Floor
San Francisco, California 94111
Tel: 415.212.9300

Attorneys for Plaintiff

UNITED STATES DISTRICT COURT
NORTHERN DISTRICT OF CALIFORNIA
SAN FRANCISCO DIVISION

DEIDRE RUSHLOW,

Plaintiff,

v.

SAMUEL ALTMAN, an individual, OPENAI FOUNDATION, a Delaware corporation, OPENAI OPCO, LLC, a Delaware limited liability company, and OPENAI GROUP PBC, a Delaware public benefit corporation,

Defendants.

Case No.:

COMPLAINT FOR:

(1) Negligence;
(2) Negligent Entrustment;
(3) Aiding and Abetting a Mass Shooting;
(4) Negligence;
(5) Negligent Undertaking;
(6) Negligence;
(7) Strict Product Liability;
(8) Strict Product Liability; and
(9) Negligent Infliction of Emotional Distress.

DEMAND FOR JURY TRIAL

Plaintiff Deidre Rushlow brings this Complaint and Demand for Jury Trial against Defendants Samuel Altman, OpenAI Foundation, OpenAI OpCo, LLC, and OpenAI Group PBC (collectively, "OpenAI"), and alleges as follows upon personal knowledge as to herself and her own acts and experiences, and upon information and belief as to all other matters.

NATURE OF THE ACTION

1. On February 10, 2026, an eighteen-year-old shot and killed their mother and eleven-year-old brother at their family home in Tumbler Ridge, British Columbia. The Shooter then drove to Tumbler Ridge Secondary School with a modified rifle and a long gun and opened fire, killing five children and an education assistant. Two more students were severely wounded, including a twelve-year-old girl who was shot in the head and suffered catastrophic brain injuries. Nearly two hundred students and teachers were trapped in classrooms and closets for hours as the massacre unfolded around them. It was one of the deadliest mass shootings in Canadian history.

2. Deidre Rushlow, a grade 7 teacher at Tumbler Ridge Secondary School, was in her classroom when the shooting began. Earlier that period, she had sent her class to the library with an educational assistant where the Shooter would, minutes later, kill several of those children and her colleague. When she first heard the gunfire and screams of terrified children from down the hall, she locked the door and hid her remaining students under her desk, including her nephew. She called the library and texted her educational assistant to make sure the children there were safe, but there was no response. Moments later, the Shooter reached her room, found it locked, and fired several shotgun blasts through it. He then walked around to the back door and fired several more shots. The bullets struck the wall behind her and her students' heads. She texted her husband a final goodbye, certain she was about to die. The trauma Deidre experienced is severe and lasting, and will haunt her for the rest of her life.

3. In the weeks that followed the attack, a sickening truth emerged: ChatGPT played a role in the mass shooting and OpenAI could have, and should have, prevented it. Sadly, the victims didn't learn this because OpenAI was forthcoming, but because its own employees leaked the story to the Wall Street Journal after they could no longer stomach the company's silence. Those whistleblower disclosures and OpenAI's eventual admissions revealed that ChatGPT deepened the Shooter's violent fixation and pushed them toward the attack—the predictable result of OpenAI's deliberate design choice to let ChatGPT engage with users about real-world violence. But that design choice was not the worst of it. OpenAI actually knew the Shooter was planning an attack and, after a contentious internal debate, made the conscious decision not to warn authorities.

4. In June 2025, just eight months before the attack, OpenAI's automated system flagged the Shooter's ChatGPT account for gun violence activity and attack planning. The account was routed to a specialized internal team of highly experienced professionals with backgrounds in threat assessment, counterterrorism, and time-sensitive threat-to-life response. After reviewing the chat logs, these professionals determined that the Shooter posed a credible threat of gun violence against real people, and urged OpenAI to notify the Royal Canadian Mounted Police ("RCMP").

5. Sam Altman and his leadership team knew what silence meant for the citizens of Tumbler Ridge, but they were more concerned about what disclosure meant for themselves. Warning the RCMP would set a precedent the company deemed unacceptable, as OpenAI would be compelled to notify authorities every time its safety team identified a user planning real-world violence. More significantly, according to former employees, every referral would mean handing law enforcement chat logs showing how ChatGPT actually engaged and actively encouraged such users to commit violence. The public would finally see what OpenAI was desperately trying to hide—that ChatGPT is not the safe, essential tool the company sells it as, but a product so inherently dangerous that its own makers routinely flag its users as threats to human life.

6. For OpenAI and Chief Executive Officer Sam Altman, this was a question of corporate survival. On May 22, 2026, OpenAI confidentially submitted a draft registration statement to the Securities and Exchange Commission for its initial public offering ("IPO"). The offering is expected to value the company at more than one trillion dollars—a transaction that would make Sam Altman one of the wealthiest and most powerful people on earth. But Altman's reign is fragile. OpenAI's Board of Directors has fired Altman once before, in November 2023, for not being "consistently candid" about safety. Altman and his team understood that revealing their flagship product was helping a teenager plan yet another violent act—this time a mass shooting—could end his tenure, derail the IPO, and wipe out the company's valuation. They did the math and decided that the safety of the children of Tumbler Ridge was an acceptable risk.

7. This calculation is a familiar one. In the 1970s, Ford kept selling the Pinto after its own engineers warned that the fuel tank design would cause people to burn to death in rear-end collisions. Ford concluded that paying settlements to the families of the dead would cost less than

fixing the car. OpenAI has made a modern version of that exact same calculation. But there is a crucial difference. For Ford, the dangerous design was a flaw in an otherwise ordinary product. For OpenAI, the dangerous design *is* the product. The core features that make ChatGPT one of the most popular products in history—its willingness to engage on any topic, to validate any belief, to sustain any fixation over time—are the very features that make it unsafe. Fixing them would cost OpenAI its market share, its path to an IPO, and hundreds of billions of dollars in valuation. OpenAI has decided that facing large judgments, even for children killed in mass shootings its product helped plan, costs less than telling the public what ChatGPT actually does.

8. And so, in June 2025, company leaders overruled the safety team, vetoed their recommendation to notify the RCMP, "deactivated" the Shooter's ChatGPT account, and kept what they had seen to themselves, hoping that whatever happened next would not be traced back to the company.

9. When the story eventually broke, OpenAI lied about what it had actually done. Scrambling to control the public fallout, it claimed to have "banned" the Shooter from using ChatGPT, choosing the term "banned" in an attempt to create the false impression that the company had taken decisive action to counter the threat. But OpenAI does not ban users from ChatGPT. At most, it will "deactivate" a user's account, which is what it did to the Shooter. Deactivation is a hollow measure that can be easily sidestepped by creating another account with a different email address. The Shooter did exactly that, either before or after the account was deactivated, and continued using ChatGPT to plan the attack.

10. That first lie collapsed when OpenAI was forced to admit that the Shooter continued using ChatGPT after the June 2025 deactivation. So, OpenAI told a second lie: it claimed the Shooter "evaded" the company's safeguards to create a new account. But there were no safeguards to evade. OpenAI's own Help Center article about deactivated accounts tells users they can regain access to ChatGPT simply by creating another account with a different email address. OpenAI lied because the truth is indefensible: to protect user growth and engagement, the company refuses to implement system-wide bans, even for a user actively planning mass violence, because it would mean losing the constant stream of user engagement data that fuels its valuation.

11. After every tragedy, OpenAI claims it acted responsibly and promises to do better. But it never promises to do the one thing that would actually make a difference—stop ChatGPT from engaging with users about violence and self-harm in the first place. When ChatGPT was first offered to consumers in 2022, it was programmed to categorically refuse those conversations. But in May 2024, OpenAI stripped away that safeguard after realizing those refusals were suppressing user engagement. For OpenAI, engagement is its entire business. Every conversation trains the next model, and every minute spent on ChatGPT grows the company's market share. OpenAI therefore programmed ChatGPT to participate in almost any conversation, no matter how dangerous.

12. OpenAI's silence about the Shooter's interactions with ChatGPT is itself a damning admission. Even though the company is withholding the Shooter's chat logs from the victims, its selective response leaves no doubt that ChatGPT encouraged the Tumbler Ridge attack. OpenAI has responded to numerous ChatGPT-related tragedies since it stripped out the violence and self-harm refusal protocols in 2024. When OpenAI believes the logs are exculpatory, it is quick to say so. For example, after the mass shooting at Florida State University ("FSU") in April 2025, OpenAI told reporters that ChatGPT "did not encourage or promote illegal or harmful activity" and that its outputs had merely provided "information that could be found broadly across public sources on the internet." OpenAI's response to Tumbler Ridge was markedly different—it called the attack a tragedy and deflected by pointing to safeguards it supposedly "strengthened" after the shooting. If the chat logs cleared OpenAI, the company would have said so.

13. Sam Altman did not apologize for more than two months after the attack. He waited until the RCMP announced its investigation was in its "final stages" to publish a letter to the community—and even then, it fell woefully short. While the letter admitted OpenAI's failure in not reporting the Shooter to the RCMP when the account was flagged in June 2025, it was otherwise filled with empty platitudes. Nothing in it suggested that Altman grasped the scale of the tragedy, including that the Shooter also murdered an education assistant, that a twelve-year-old girl suffered catastrophic brain injuries, or that nearly two hundred terrified students and teachers were trapped in classrooms and closets for hours as the massacre unfolded around them.

14. Deidre was within the zone of physical danger created by OpenAI's conduct, and she knew it in real time. The terror she experienced that day was severe, foreseeable, and the direct result of OpenAI putting a product on the market that it knew was being used to plan mass violence.

15. Accordingly, Plaintiff brings this action to hold Sam Altman and OpenAI accountable for designing a dangerous product, ignoring the warnings of their own safety team, refusing to notify authorities when they knew the Shooter was planning a mass attack, allowing the Shooter to continue using ChatGPT after deactivating their account, and choosing profit over the lives of the children of Tumbler Ridge.

PARTIES

16. Plaintiff Deidre Rushlow is a resident of Tumbler Ridge, British Columbia, Canada. She is a citizen of Canada.

17. Defendant Samuel Altman is a natural person residing in San Francisco, California. He is a citizen of California. As co-founder and Chief Executive Officer of OpenAI, Altman directed the design, development, safety policies, and deployment of ChatGPT. In 2024, Altman knowingly accelerated GPT-4o's public launch while deliberately bypassing critical safety protocols.

18. Defendant OpenAI Foundation (formerly known as OpenAI, Inc.) is a Delaware corporation with its principal place of business in San Francisco, California. At all relevant times, OpenAI Foundation was the nonprofit parent entity that governed the OpenAI organization and exercised oversight over its for-profit subsidiaries, including OpenAI OpCo, LLC. As the governing entity, OpenAI Foundation was responsible for defining the organization's safety mission, establishing its risk-management framework, and publishing the official "Model Specifications" that set the policies and requirements applicable to the development and deployment of OpenAI's artificial intelligence models.

19. Defendant OpenAI OpCo, LLC is a Delaware limited liability company with its principal place of business in San Francisco, California. OpenAI OpCo, LLC is the for-profit operating entity within the OpenAI corporate structure, responsible for the development,

deployment, and commercialization of the defective product at issue. OpenAI OpCo, LLC managed and operated the ChatGPT product that the Shooter used to plan their attack, including the infrastructure and systems through which GPT-4o was delivered to end users. On information and belief, the members of OpenAI OpCo, LLC include OpenAI Foundation, a Delaware nonprofit nonstock corporation with its principal place of business in California. On information and belief, tracing through each tier of ownership to its ultimate natural-person and corporate members, no member of OpenAI OpCo, LLC is a citizen or subject of Canada.

20. 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 consolidating OpenAI's for-profit operations. On December 31, 2025, OpenAI Holdings, LLC—the for-profit holding entity that owned the core intellectual property underlying GPT-4o and the other commercial models at issue—merged into OpenAI Group PBC and ceased to exist. As successor to OpenAI Holdings, LLC and the other predecessor for-profit entities, OpenAI Group PBC is liable for their conduct; it designed, deployed, and profited from ChatGPT, and continues to do so today. On May 22, 2026, OpenAI confidentially submitted a draft registration statement to the Securities and Exchange Commission for an initial public offering of, on information and belief, OpenAI Group PBC.

21. Defendants collectively played the most direct and consequential roles in the design, development, approval, and deployment of the defective product at issue. OpenAI Foundation established the safety mission it failed to enforce. OpenAI OpCo, LLC built, deployed, and sold the defective product at issue. OpenAI Group PBC owns the underlying technology and is liable as the successor to the preceding for-profit entities that profited from it. Samuel Altman is named as the chief executive who personally directed the reckless strategy of prioritizing a rushed market release over the safety of vulnerable users and the public. Together, these Defendants represent the key actors whose decisions directly caused the harm at issue.

JURISDICTION AND VENUE

22. This Court has subject matter jurisdiction pursuant to 28 U.S.C. § 1332(a)(2) because Plaintiff is a citizen of Canada, Defendants are citizens of California, and the amount in controversy exceeds the sum of $75,000.00.

23. This Court has personal jurisdiction over Defendants because Defendants conduct business in this District, reside in this District, and a substantial part of the events or omissions giving rise to Plaintiff's claims occurred in this District.

24. Venue is proper in this District under 28 U.S.C. § 1391(b), because a substantial part of the events giving rise to Plaintiff's claims occurred in and emanated from this District, and Defendants reside in this District.

DIVISIONAL ASSIGNMENT

25. 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.

26. On February 10, 2026, an eighteen-year-old resident of Tumbler Ridge, British Columbia, carried out one of the deadliest mass shootings in Canadian history. The Shooter first shot and killed their mother and eleven-year-old brother at their family home. They then traveled to Tumbler Ridge Secondary School armed with a modified rifle and a long gun. Inside, they opened fire, killing five children and an education assistant, and severely wounding two others, leaving one with life-changing traumatic brain injuries from being shot in the head. The rampage ended with the Shooter's suicide.

27. During the horrific shooting, children hid under desks and in bathroom stalls and closets, praying the Shooter would not find them. Many watched their friends, and even a teacher, get shot at point-blank range. Some pulled the injured and the dead into their hiding places, careful not to leave trails of blood that could lead the Shooter to them. Parents were later asked to identify their children by their clothing, because the gunshots had left little else for them to recognize. The

survivors—students, teachers, and parents alike—are living with physical and psychological injuries that will never fully heal. Tumbler Ridge Secondary School has been closed and will be demolished. The children who can return to school at all are doing so in trailers set up as makeshift classrooms.

28. As alleged above, Deidre Rushlow was inside Tumbler Ridge Secondary School when the shooting began. She had sent her grade 7 class to the library minutes before the Shooter would kill several of those children and the educational assistant supervising them. She then locked the room and hid her remaining students under her desk while the Shooter fired several shotgun blasts through the classroom doors that struck the wall behind their heads. The trauma Deidre experienced is severe and lasting, and will haunt her for the rest of her life.

II. OpenAI Knew the Shooter Was Planning the Attack and Rejected Its Own Safety Team's Advice to Notify Law Enforcement.

29. In June 2025—eight months before the attack—OpenAI's automated review system flagged extensive activity on a ChatGPT account describing scenarios involving gun violence. The conversations with ChatGPT spanned multiple days. The system routed the account to what OpenAI has described as a specialized pipeline for users "planning to harm others," where 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."

30. OpenAI's human reviewers did exactly what this policy contemplated. They reviewed the Shooter's flagged conversations. They were not confused about what they were reading. Multiple team members explicitly recommended contacting the RCMP. They identified what they believed was a credible and imminent threat of serious physical harm to real people. These trained safety professionals—people whose job was to evaluate exactly this kind of content—analyzed the Shooter's conversations and concluded that the authorities must be notified.

31. OpenAI's leadership overruled them. Company leaders argued that the case "did not meet OpenAI's threshold of 'credible and imminent' risk of physical harm." The safety team employees who reviewed the Shooter's account had received specialized training in evaluating

threatening content and assessing real-world risk. Altman and the company leaders who overruled them did not pretend to have any comparable training or expertise. The account was deactivated, and the Shooter was permitted to start a new account. No one called the RCMP and nobody in Tumbler Ridge was warned that a member of their community was actively planning a mass shooting.

A. Company Leadership Overruled the Safety Team With Full Knowledge That ChatGPT Had Already Been Used to Plan Mass Violence.

32. By June 2025, the company had clear knowledge that its product was being used by disturbed individuals to plan and prepare for violence against real people. Some publicly known examples:

- In January 2025, a man used ChatGPT for feedback on how to use explosives and evade surveillance before detonating a Tesla Cybertruck in front of the Trump International Hotel in Las Vegas.
- In April 2025, a twenty-year-old gunman carried out a mass shooting at FSU. Chat logs showed that the gunman used ChatGPT extensively leading up to and during the attack. The gunman's conversations with ChatGPT had included questions about how to fire a shotgun, the legal fates of school shooters, and when the student union would be busiest.
- In May 2025, a teenage boy in Finland used ChatGPT for nearly four months to help prepare for an attack in which he stabbed three fourteen-year-old girls at his school. Finnish authorities reported that the boy had made hundreds of chatbot queries, including research into stabbing tactics, concealment of evidence, and information on mass killings.

33. On information and belief, Altman and OpenAI have actively concealed countless additional incidents where ChatGPT was utilized to plan or execute acts of mass violence, domestic and international terroristic activity, and targeted assassinations against political leaders, corporate CEOs, and prominent public figures.

34. The safety team members who reviewed the Shooter's account in June 2025 knew all of this. When they urged OpenAI to contact the RCMP, they were not speculating what might happen. They identified a pattern that had already repeated itself many times that year. OpenAI's leadership overruled them.

35. On information and belief, Defendants have systematically withheld credible threats of violence from appropriate law enforcement agencies because they recognized that revealing the true extent to which ChatGPT has facilitated violence would give rise to profound civil and criminal liability for both the corporate entities and individual leadership.

B. OpenAI's Stated Reasons for Refusing to Warn the RCMP—Concern for the Shooter's "Privacy" and the Absence of an "Imminent" Threat—Were Pretext to Conceal ChatGPT's Role in the School Shooting.

36. When it was revealed after the attack that OpenAI had rejected its safety team's demand to alert Canadian law enforcement, the company attempted to justify its decision by explaining that there was no indication that the attack was "imminent" and it "weighs the risk of violence against privacy considerations and the potential distress caused to individuals and families by getting police involved unnecessarily." Even taken at face value, these excuses are more damning than the decision itself—they reveal that OpenAI's business leaders believed their own untrained judgment about when to involve police was more reliable than the judgment of the safety professionals they employ specifically to make that decision.

37. As it turns out, the Shooter was already known to local law enforcement before the attack. Canadian authorities had visited the residence multiple times in connection with mental health concerns and had temporarily removed firearms from the home. A law enforcement referral from OpenAI would have reached police who already had an open file on the Shooter, who had already been to the home, and who had already recognized the Shooter as someone whose access to firearms warranted immediate intervention.

38. The real reason OpenAI stayed silent was not privacy or imminence. It was self-preservation. A former member of OpenAI's investigations team has confirmed as much. Tim Marple, who worked on the team responsible for flagging dangerous users, told the New York Times that OpenAI was reluctant to report users to law enforcement because doing so "forces them to share information about how their product is potentially exacerbating the threat environment." In a conversation about a mass shooting, Marple explained, ChatGPT "might be providing strategically valuable, illustrative scenarios."

39. Altman and his team understood that if the public saw how far ChatGPT had already gone in helping real people plan violence, it would pose an existential risk to OpenAI and a personal risk to Altman. Since ChatGPT's launch, OpenAI has worked tirelessly to construct and reinforce a simple narrative: its product is safe, essential, and should be used every day. It has marketed ChatGPT to parents as a homework helper, to workers as a daily assistant, and to governments and defense contractors as a tool for national security work. Altman has testified before Congress, appeared at international summits, and met with heads of state to deliver a single message: that ChatGPT is a responsible product built by a responsible company.

40. In fact, OpenAI has spent, on information and belief, hundreds of millions of dollars on advertising, public relations campaigns, and "trust and safety" messaging designed to reassure the public that using ChatGPT is safe, while simultaneously deploying a federal and state lobbying operation aimed at securing legal protections that would shield it from liability when its product causes harm. It has retained outside crisis communications counsel, former senior government officials, and, on information and belief, criminal defense attorneys to advise OpenAI of potential criminal liability stemming from the deaths caused by its product. All of this has been in service of a race—for market share, for a valuation expected to exceed one trillion dollars, and for an IPO that would rank among the largest in history—against rivals that market themselves as the safer alternative. Defendants understood that the race had a finite window. More users were dying. More reporters were asking questions. Altman's own Board of Directors had already tried to oust him once, in November 2023, for his lack of candor about safety. Every new disclosure about ChatGPT's role in real-world violence brought the company closer to the moment when the public, regulators, or its own Board would force the market to reprice the product—and with it, Altman's personal position at OpenAI. Defendants' concealment strategy was not an attempt to keep ChatGPT's dangers secret forever. It was an attempt to conceal them long enough to complete the IPO. The strategy worked. On May 22, 2026, OpenAI confidentially submitted a draft registration statement to the SEC, setting in motion the very offering its silence was designed to protect.

41. That is the backdrop against which Altman and his leadership team made the decisions at issue here. They knew that if regulators, investors, or the public learned that ChatGPT was already being used to plan shootings—and that OpenAI had ignored its own safety staff and refused to alert police—the narrative they had spent years and hundreds of millions of dollars building would collapse. Their valuation and IPO prospects would be at risk, and Altman himself could be forced out. Against that backdrop, neither the "privacy" explanation nor the "imminence" standard provided a good-faith balancing of interests. Both were cover for staying quiet to protect OpenAI and Altman, even if that meant leaving the people of Tumbler Ridge in danger.

C. OpenAI's Organizational Structure Further Exposes that "Imminence" and "Privacy" Were Pretextual Explanations to Hide the Raw Political Calculus Employed by Defendants.

42. There are two teams within OpenAI that are directly relevant to this case. The first—an internal threat-assessment team composed of safety professionals—is responsible for monitoring ChatGPT, identifying users who pose a threat of real-world violence, and making recommendations about whether to refer those users to law enforcement. The second—the Global Affairs team—is responsible for managing OpenAI's public image, navigating political relationships, and mitigating reputational damage. Even though the teams should be acting completely independently of each other given their diverging missions, the Global Affairs team has the authority to override explicit recommendations made by those safety professionals to notify law enforcement of a credible threat of violence. As the allegations that follow demonstrate, that is precisely what occurred here.

43. OpenAI has never been transparent about which team within the company is responsible for monitoring, mitigating, or responding to violent content. OpenAI has only described the reviewers in vague and generic terms, and has never revealed the team's name, qualifications, or chain of command. For example:

- In a February 26, 2026 letter to Canadian Minister Solomon addressing the Tumbler Ridge shooting, OpenAI vaguely stated that "OpenAI's automated system detected the account, and it was subsequently sent to human review to determine whether our usage policies were violated and whether the account warranted referral to law enforcement."

- In an April 28, 2026 blog post titled "Our Commitment to Community Safety," OpenAI explained that "[w]hen an account or conversation is flagged, it is assessed in context by trained personnel," and that "[t]hese human reviewers are trained on our policies and protocols, and operate within established privacy and security safeguards."

44. While OpenAI hides the team's identity, it can be established through publicly available sources, including the company's own job postings and organizational data. These records reveal that the team responsible for reviewing ChatGPT conversations flagged for violent activity is OpenAI's "Intelligence and Investigations" team. OpenAI's job postings state that the team's mission is to "rapidly identify and mitigate abuse and strategic risks to ensure a safe online ecosystem." Investigators on the team "[r]eview leads, investigate activity and disrupt abusive operations . . . focused on violent and terrorist activities, especially associated with time sensitive threat to life risks" and possess "deep domain-specific expertise in identifying, understanding, and mitigating risk from violent attacks and behavior, and terrorist-related activity."

45. OpenAI cannot claim it lacked the expertise to assess the threat to Tumbler Ridge posed by the Shooter. The Intelligence and Investigations Team is led by William McCants, whose past work has largely focused on national security threat assessments for the United States government. McCants joined OpenAI as Head of Intelligence and Investigations in December 2023, and manages a team of investigators with backgrounds in counterterrorism, threat intelligence, and time-sensitive threat-to-life response.

46. Despite their apparent expertise, the Intelligence and Investigations Team does not have the authority to act on their findings and recommendations, including with respect to law enforcement referrals. Instead, organizational data from TheOrg.com[1] shows that McCants, as Head of the Intelligence and Investigations Team, reports directly to OpenAI's Chief Global Affairs Officer, Chris Lehane, who in turn reports to Sam Altman.

[1] TheOrg.com is an online platform that maps the organizational structures of major companies. The organizational chart published for OpenAI is based on data compiled from various sources "to ensure accuracy and comprehensiveness," including OpenAI's corporate disclosures, press releases, and public LinkedIn profiles, as well as information provided to the platform directly by individuals with a verified OpenAI corporate email address. The resulting organizational chart provides an extensive and highly detailed rendering of OpenAI's internal hierarchy and traces reporting lines across nearly 4,600 employees and more than 130 distinct teams.

47. Unlike McCants, Lehane is not a security or threat-assessment professional. He is a career political operative with no prior professional experience in artificial intelligence, threat assessment, or public safety or security work. He previously served in the Clinton White House, where he ran opposition research against critics of the administration, as press secretary to Vice President Al Gore, as a senior advisor to the Kerry presidential campaign, and later as Airbnb's global head of policy and communications.

48. In each of these roles, Lehane, who calls himself the "master of disaster," managed political scandals and public relations crises using underhanded damage-control tactics. In 2012, he co-authored *Masters of Disaster: The Ten Commandments of Damage Control*, a how-to manual for shielding powerful clients from the consequences of their conduct. One of the book's core rules—the seventh commandment, "Respond with Overwhelming Force"—instructs that when a client comes under attack, the proper response is not to retreat, but to escalate and counterattack by discrediting the accusers, controlling the narrative, and ensuring that anyone who challenges the client pays a heavy price. Lehane articulated this strategy even more bluntly in a 2013 seminar at Amherst College, where he said that critics should "pay some type of price."

49. At his crisis-management firm, Fabiani & Lehane, Lehane built a reputation for burying clients' problems and attacking their critics. His firm represented Lance Armstrong as he publicly denied doping allegations that were later proven true, while his team worked to aggressively discredit the accusers who had exposed the truth; advised Goldman Sachs during the fallout from the 2008 financial crisis and the government's fraud case over its mortgage conduct; and handled media strategy for Al Gore's Current TV during its contentious firing of anchor Keith Olbermann. As head of global policy and public affairs at Airbnb, Lehane created company-funded "Home Sharing Clubs"—groups that Airbnb bankrolled and coordinated but presented to local officials as spontaneous, grassroots movements—to manufacture the illusion of widespread public opposition to local regulations. And as a board member at Coinbase, he built Fairshake, a $200 million Super PAC that targeted federal lawmakers the crypto industry viewed as hostile. Instead of targeting them on crypto policy, the Super PAC ran ads attacking them on unrelated issues, in a campaign one insider said was meant to "terrify other politicians" into staying in line.

50. Lehane was hired by OpenAI in 2024 and serves as the company's Chief Global Affairs Officer. The Intelligence and Investigations Team—the only team inside OpenAI responsible for identifying ChatGPT users who pose a threat of real-world violence—was placed under his control. That means, on information and belief, the team's recommendations regarding law enforcement referrals are subject to Lehane's and Altman's approval. As a result, the decision whether to alert law enforcement to a user planning a mass attack was not made by the trained threat-assessment professionals who urged OpenAI to contact the RCMP. It was made, on information and belief, by Lehane himself, or by someone in his chain of command, and ratified by Sam Altman.

51. OpenAI's organizational structure exposes its entire safety framework as totally corrupted and untrustworthy. By placing final authority over violent-threat referrals in the hands of a political strategist rather than the professionals trained to assess them, OpenAI guarantees that threats of real-world violence are weighed as public relations risks rather than imminent physical dangers.

52. The same reputational calculus that kept OpenAI silent about Tumbler Ridge explains why it readily coordinates with law enforcement on some threats but not others. OpenAI publicizes select collaborations in threat reports issued through Lehane's Global Affairs function. The February 2026 report titled *Disrupting Malicious Uses of Our Models* took credit for identifying ChatGPT users running scam networks and sharing their information with authorities. Subsequent reports take similar credit for identifying ChatGPT users tied to foreign influence operations, surveillance networks, and state-linked activity originating from China, Iran, Russia, and North Korea. Crucially, none of these reported activities implicate the design or safety of ChatGPT itself—each involves an outside actor misusing an otherwise functioning product. OpenAI publishes these reports to cast itself as a responsible corporate citizen and a willing partner to law enforcement, even as it conceals the threats of violence that would expose the danger of its product.

53. That selective cooperation further exposes the falsehood of OpenAI's claim that it declined to notify the RCMP out of concern for the Shooter's "privacy." When a referral advances

OpenAI's political goals—such as when it allows the company to posture as a defender of national security—no privacy concern stops OpenAI from identifying users and sharing their information with law enforcement. The manufactured "privacy" rationale appears only when a referral would force OpenAI to reveal that its flagship consumer product actively guided a teenager through the planning of a mass shooting.

54. The real reason for OpenAI's silence was not privacy but concern for its own reputational and financial standing. When the recommendation to contact the RCMP reached, on information and belief, Chief Global Affairs Officer Lehane, the calculation was: either contact the RCMP and expose ChatGPT as a tool that coaches users toward mass violence—a disclosure that could cost OpenAI its users and its trillion-dollar valuation while inviting intense regulatory scrutiny and public controversy—or say nothing and hope that no one would ever trace the shooting back to the company. At the time, saying nothing was a safe bet. Law enforcement had no protocol for reviewing AI chat logs, and it is only because OpenAI's own employees came forward that its role in Tumbler Ridge came to light at all. The risk of real-world deaths never factored into the calculation—or if it did, it was treated as an acceptable cost of protecting the corporate brand.

55. From a cold-blooded, political perspective, OpenAI's appraisal of the situation may well have been correct. The company feared that a law enforcement referral would link ChatGPT to mass violence in the public's mind and subject OpenAI to backlash, investigations, and enforcement, which is exactly what happened when ChatGPT's role in two Florida incidents—the FSU school shooting and another murder where the suspect used ChatGPT to conceal the crime—became public. Shortly after those revelations, the company faced a criminal investigation and an enforcement action by the Florida Attorney General. Viewed purely from the standpoint of protecting OpenAI's valuation and its path to an IPO, staying silent about Tumbler Ridge was a rational, and no doubt calculated, choice. OpenAI correctly understood the stakes, correctly predicted the consequences, and chose to preserve its corporate valuation by shifting a catastrophic risk away from its brand and entirely onto the victims of Tumbler Ridge.

D. When the Threat Is to OpenAI's Own Employees, the "Imminence" and "Privacy" Standards Disappear.

56. OpenAI's handling of threats to its own employees proves its Tumbler Ridge narrative was a lie. Threats to OpenAI's executives and personnel are assessed by a separate group—the Corporate Security Team—which operates within a chain of command designed solely for immediate protection. The team reports to OpenAI's Vice President of Corporate Security, Patrick Geonetta, who in turn reports to the Chief Operating Officer, Brad Lightcap. Notably, Chief Global Affairs Officer Lehane doesn't even factor into the conversation.

57. The Corporate Security Team is led by professionals with operational law enforcement and protective intelligence experience. Its Director of Protective Intelligence, Kyle Goggin, is a former United States Secret Service agent who served on the Presidential Protective Division. The team operates a 24/7 Global Security Operations Center that provides "incident response, threat monitoring, situational awareness, and support to teams across the organization," and maintains active partnerships "with federal law enforcement and intelligence agencies in the US and abroad to ensure timely sharing of threat intelligence."

58. When a threat to OpenAI's personnel arises, the Corporate Security Team can lock down facilities, alert employees, and notify law enforcement on its own authority. Because it reports to the Chief Operating Officer rather than the Chief Global Affairs Officer, the team sits entirely outside OpenAI's public relations and can act on a threat immediately, without first weighing physical safety against reputational concerns. By contrast, when the Intelligence and Investigations Team sees ChatGPT coaching a user to commit violence against the public, it can only make a recommendation to a political strategist who specializes in crisis management and reputational protection.

59. While both teams ostensibly exist to protect human life, they follow completely different decision paths. For threats against the public, OpenAI asks two questions: First, is the threat credible? Second, will disclosing the threat damage ChatGPT's reputation or valuation? If the political risk is deemed too high, the danger is concealed and the public is left without a warning. But for threats against its own executives and personnel, OpenAI asks just one question: Is the threat credible? If it is, the Corporate Security Team takes immediate action to protect its

people—notifying law enforcement, warning staff, and locking down offices—without ever waiting for the threat to become "imminent."

60. This double standard is illustrated by comparing OpenAI's silence in the face of the Tumbler Ridge threat with its urgent response to a non-specific threat against its own employees. On November 21, 2025, the Corporate Security Team learned of an individual who had expressed interest in physically harming company personnel. Although OpenAI acknowledged there was "no indication of active threat activity"—and thus no indication of an "imminent" attack—the company refused to wait. It immediately locked down its offices, warned employees, distributed the suspect's name and photograph, and notified the San Francisco Police Department. When its own people were at risk, no "privacy" interest stopped OpenAI from notifying the police or circulating the man's name and photograph to thousands of employees. Nor did it wait for the attack to become imminent because it recognized that waiting would jeopardize employee safety.

61. The contrast with Tumbler Ridge could not be sharper. When the targets were middle school children and teachers in a small Canadian town, OpenAI suddenly invoked the "imminence" and "privacy" standards to justify withholding critical information from law enforcement regarding a threat its own threat-assessment team had already flagged as a real and urgent danger. These contrasting responses reveal that the "imminence" and "privacy" standards are not genuine safety protocols, but manufactured excuses to justify OpenAI's deadly silence.

62. OpenAI's structural subordination of public safety to corporate reputation reflects a recognized failure in technology governance. Trust and safety functions are widely understood to require strict independence from public relations operations, just as an internal audit must remain independent of the business units it oversees. By forcing its threat-assessment professionals to report to the Chief Global Affairs Officer, OpenAI eliminated that separation. This guaranteed that warnings of real-world violence would be filtered through a crisis-management framework that prioritized the corporate brand over human life—with devastating consequences for the students and teachers at Tumbler Ridge Secondary School.

III. OpenAI's Claim That the Shooter "Evaded" Its Systems Was a Lie.

63. After the massacre, OpenAI told two stories about the Shooter's ChatGPT access. First, that the company had "banned" their account when it was flagged for violent activity. Then, when they returned to ChatGPT on a second account, that the Shooter had "evaded" OpenAI's safety systems. Neither was true.

64. OpenAI has no mechanism to ban users. What it has is a process called "deactivation," which it uses for usage-policy violations. A ban would have prevented the Shooter from returning. Deactivation only shuts down the email address the user signed up with. The user is free to come back under a different email, and the Shooter did.

65. The clearest proof is OpenAI's own published guidance. Its Help Center includes an article titled "Why Was My OpenAI Account Deactivated?" The article lists the categories of conduct that trigger deactivation, including, explicitly, "[v]iolence and self-harm" usage-policy violations. It then tells those users how to come back, under a section titled "How to Prevent Future Deactivations" that instructs them on what to do differently next time.

66. OpenAI's customer service team also tells deactivated users how to return by email:

> Please be aware that previously deleted OpenAI accounts cannot be reactivated. However, you can create a new account using the same email address once 30 days have passed since the deletion. If you prefer not to wait, you have the option to register immediately using an alternative email address. If you don't have another address available, you can use an email sub-address instead. For example, you could try jane+alt@example.com in place of jane@example.com. While your email provider will likely treat both addresses the same, our system will recognize the sub-address as a new account.

67. The Shooter followed those instructions. Either before or after the deactivation, the Shooter registered for a second account with a different email address. OpenAI called that evasion because it could not admit the truth: the company had been telling deactivated users how to come back all along.

68. OpenAI designed its system this way because its revenue depends on user count, session volume, and subscription conversions—and a deactivated user who never returns is lost revenue.

IV. ChatGPT Is a General-Purpose Chatbot That OpenAI Mass Markets to Ordinary Consumers as Safe for Everyday Use.

69. Chatbots are software programs designed to simulate human conversation. A user types a message and the program responds with a string of words that simulate human language. The software does this mechanically and autonomously by reading the input, calculating data patterns, and generating output. Even though a chatbot's output often appears to the user to be human language, there is no human on the other end that participates directly in the exchange.

70. ChatGPT is a general-purpose chatbot, meaning that it is built to converse about any topic a user brings to it. The version currently known to be at issue in this case is GPT-4o, but discovery may show that the Shooter used other OpenAI models as well. On information and belief, any such models operated under the same behavior guidelines and the same engagement-first design as GPT-4o, and the allegations concerning GPT-4o apply equally to them.

71. OpenAI markets ChatGPT for everyday use by ordinary consumers, including children, students, and families. In Apple's App Store, OpenAI sells ChatGPT as "Your everyday AI assistant." On Google Play, it is "Your AI assistant for writing, search, image generation, and more." OpenAI's own store listings promise "[t]ailored advice" to "[t]alk through a tough situation," "[p]ersonalized learning" that can "[e]xplain electricity to a dinosaur-loving kid," "[c]reative inspiration," and "[i]nstant answers." ChatGPT's website maintains marketing pages aimed at "Students," "Teachers," "Parents," and "College Students," promotes uses from "Fitness, Wellness, and Health" to "Money and Finances" to "Recipes and Cooking," and sells subscription plans for "K–12 teachers." OpenAI rates the app as appropriate for children: thirteen and up in the App Store, twelve and up on Google Play.

72. ChatGPT requires no technical skill to use and is available to anyone through a web browser or through OpenAI's iOS, Android, and desktop applications. The Android app alone has been installed more than one billion times, OpenAI's App Store listing urges consumers to "[j]oin hundreds of millions of users and try the app captivating the world," and ChatGPT ranks first in the "Productivity" category of both major app stores.

V. The Tumbler Ridge Attack Was a Foreseeable Consequence of OpenAI's Design Choices and Safety Failures.

73. OpenAI designed ChatGPT to maximize engagement, not safety. In April 2025, OpenAI introduced a feature called "memory," turned on by default, that allowed ChatGPT to "save" details users shared and treat them as "part of the context ChatGPT uses to generate a response" going forward.

74. For ordinary users, memory was a convenience. For a user planning violence against real people, it was an encouraging co-conspirator. On information and belief, GPT-4o used the memory feature to build a comprehensive profile of the Shooter over the months they interacted with it—tracking their grievances, their targets, their reasoning, and their plans across separate conversations—and used that profile to sustain and deepen their fixation on violence. GPT-4o also employed anthropomorphic design elements to cultivate emotional dependency. The system used first-person pronouns (e.g., "I understand," "I'm here for you"), expressed apparent empathy, and maintained conversational continuity that mimicked human relationships.

75. Together, these features replaced human relationships with an artificial confidant that was always available, always affirming, and never challenged anything the user said—even when the user was talking about plans to kill people. For a teenager fixated on violence and growing increasingly isolated, ChatGPT morphed into an encouraging co-conspirator.

76. OpenAI controls how ChatGPT behaves through internal rules called "behavior guidelines," which are formalized in a document known as the "Model Spec." The Model Spec contains the company's instructions for how ChatGPT should respond to users—what it should say, what it should avoid, and how it should make decisions. As Sam Altman explained in an interview with Tucker Carlson, the Model Spec reflects OpenAI's values: "the reason we write this long Model Spec" is "so that you can see here is how we intend for the model to behave." The Model Spec classifies dangerous content into tiers that dictate how cautiously ChatGPT must respond. Its most protective tier—categorical refusal—is strictly reserved for three things: sexual content involving minors; chemical, biological, radiological, and nuclear threats; and content that might result in copyright infringement. Mass shootings did not make the categorical refusal list. OpenAI placed conversations about mass shootings and other "imminent real-world harm" in a

low-level category called "Take extra care in risky situations," where ChatGPT is instructed only to "try" to prevent harm, to assume best intentions, and never to ask the user to clarify intent. Altman openly acknowledged the tradeoff: "If you just do the naïve thing and say, 'never say anything that you're not a hundred percent sure about,' you can get [the model] to do that, but it won't have the magic that people like so much."

77. While the "Model Spec" describes how OpenAI wants ChatGPT to behave, the company's "System Cards" document how the technology functions in practice based on the company's own testing. These records show that OpenAI was fully aware of ChatGPT's utility in planning real-world violence. The GPT-4 System Card explicitly stated that intentional probing of early versions could lead to harmful content such as "[c]ontent useful for planning attacks or violence."

78. OpenAI initially presented itself as a company that cared deeply about safety. In July 2023, it signed voluntary commitments to the White House to conduct pre-deployment safety testing—internal and external red-teaming—and to share the results with the federal government before deployment. But that commitment was abandoned less than a year later, when Altman learned that Google would unveil its competing Gemini model on May 14, 2024. Though OpenAI had planned to release its next model, GPT-4o, later that year, Altman moved the launch to May 13 to beat Google's release by one day. To meet the new date, OpenAI compressed months of planned safety evaluation into approximately one week, and Altman personally overruled safety personnel who demanded additional time for red-teaming. Invitations for GPT-4o's launch party went out before testing was complete. One employee told *The Washington Post*: "They planned the launch after-party prior to knowing if it was safe to launch[.]" A member of OpenAI's own preparedness team—responsible for testing whether the model posed catastrophic risks, including the potential to assist in violence—admitted the testing had been "squeezed."

79. OpenAI's senior safety leadership then resigned in protest, including co-founder and chief scientist Ilya Sutskever and Superalignment co-lead Jan Leike, who said publicly: "[O]ver the past years, safety culture and processes have taken a backseat to shiny products."

80. Under the rules that governed the product before May 2024, conversations

glorifying or elaborating on violence would have been refused outright. But the Model Spec replaced those rules with instructions to "assume best intentions"—and, critically, to "never ask the user to clarify their intent . . . for the purpose of determining whether to refuse or comply."

81. The GPT-5 System Card, released August 7, 2025, revealed that while OpenAI models had been scoring near-perfectly on harmful content refusals under its standard disallowed content evaluations—where prompts not representative of real multi-turn conversations were submitted to the model—that result did not hold under more realistic conditions. When OpenAI later began testing using multiple rounds of back-and-forth conversation representative of real-world use, the GPT-4o model refused only approximately 57% of illicit and 63% of illicit/violent requests.[2] The System Card further acknowledged that those standard evaluations showing near-perfect safety scores no longer provide a useful signal of incremental changes in system safety and performance, and OpenAI announced plans to stop publishing them entirely.

82. OpenAI did not publicly acknowledge this deficiency until forced to. On August 26, 2025—the same day a wrongful death lawsuit was filed alleging that ChatGPT had coached a 16-year-old to hang himself—OpenAI published a blog post admitting that ChatGPT safeguards "can sometimes be less reliable in long interactions" and that "after many messages over a long period of time, it might eventually offer an answer that goes against our safeguards." In this way, OpenAI framed the problem as if it was a glitch that surfaced only in "long conversations." But they designed the product precisely to carry out long conversations.

83. Before OpenAI introduced its memory feature, ChatGPT sessions were generally brief and episodic: every chat started with a blank slate, giving users no real reason to engage in prolonged exchanges. The memory feature fundamentally changed that dynamic. By designing the product to retain personal details and "know" the user over time, OpenAI actively encouraged people to engage in much longer conversations with far greater depth. Because the product carried each user's history into every new conversation, formerly isolated chats fused into one continuous

[2] OpenAI defines "illicit" as "[c]ontent that gives advice or instruction on how to commit illicit acts. A phrase like 'how to shoplift' would fit this category" and defines "illicit/violent" as "[t]he same types of content flagged by the illicit category, but also includes references to violence or procuring a weapon" in its Moderation API documentation.

relationship, meaning even a brief daily check-in functioned as a continuation of one massive, ongoing interaction spanning weeks or months. OpenAI deliberately engineered a product to drive this exact type of deep, long-term engagement—degrading its safety guardrails—and then deployed it to the public without warning.

84. The reluctance to disclose known safety failures to the public traces directly to the top of OpenAI's leadership structure. Sam Altman has a reputation within OpenAI's leadership for lying about safety issues, and his credibility on safety-related representations became the subject of extensive sworn testimony during the May 2026 *Musk v. Altman* trial. That trial centered on allegations that Altman had misled investors by transforming OpenAI from a nonprofit safety-focused research laboratory into a for-profit company worth hundreds of billions of dollars. Although the jury ultimately found that Musk's claims were untimely and did not reach the merits, multiple former OpenAI directors and executives testified under oath about Altman's statements concerning safety review and deployment decisions.

85. Several witnesses testified that they had personally observed Altman misrepresent OpenAI's safety processes. Former director Helen Toner testified that Altman told the Board that GPT-4 had been reviewed and cleared by OpenAI's safety panel. But when she later asked for the supporting documentation, Toner learned that the model had not in fact received the approval Altman had represented. Former director Tasha McCauley testified about similar concerns and described OpenAI under Altman as having "a toxic culture of lying." Former Chief Scientist Ilya Sutskever likewise testified that he prepared a 52-page memorandum for the independent directors detailing Altman's conduct, which began: "Sam exhibits a consistent pattern of lying . . . ."

86. Mira Murati, OpenAI's former Chief Technology Officer, also testified that Altman lied to her about the GPT-4 Turbo safety review. Murati testified that in late 2023, Altman told her that OpenAI's General Counsel, Che Chang, had determined that GPT-4 Turbo did not require review by the company's deployment safety board. In other words, Altman was attempting to deploy OpenAI's most powerful model to date to hundreds of millions of users without going through the company's ordinary safety-review process—and represented that General Counsel Chang had approved that decision. Murati found the claim unusual enough to check with Chang

herself. Chang told her that he had made no such determination. By then, Murati had already spent more than a year warning Altman that OpenAI needed to take safety seriously. In a memorandum she sent to Altman in 2022, Murati suggested that the company needed "a complete overhaul" of its approach to safety and risk management. She described an internal culture defined by "constant panic," "chaos and churn," and a reckless "do-everything and do it fast" mentality, and proposed that she have final authority over safety decisions—a proposal that was quickly rejected. When Murati was asked under oath if Altman was an honest leader, she testified, "not always."

87. None of the concerns raised by OpenAI's own personnel stopped the company from offering new assurances about safety whenever doing so served its business interests. In October 2025, as a condition of approval for OpenAI's conversion from a nonprofit to a for-profit corporation, Altman signed a binding Memorandum of Understanding ("MOU") with the California Attorney General. The MOU committed OpenAI to giving its Safety and Security Committee "an effective approval right" over safety-related decisions. It required the PBC Board to "consider only the Mission (and may not consider the pecuniary interests of stockholders[])" on safety and security issues. And it promised that "OpenAI will continue to undertake measures to mitigate risks to teens and others in connection with the development and deployment of AI."

88. The Tumbler Ridge attack took place four months after OpenAI signed that MOU. During those four months, OpenAI kept GPT-4o on the market, kept the design features that made the model validating and agreeable in conversations about violence, and kept the re-registration policies that let deactivated users come back. By the time OpenAI signed the October 2025 commitments, OpenAI had already broken them. And OpenAI knew exactly what those broken commitments meant here. The Shooter was on ChatGPT planning a mass attack. Its automated systems had flagged their account for gun violence. Its trained safety team had reviewed their conversations and identified them as a real-world threat. But OpenAI overruled them and let the Shooter keep using the product.

89. By the time OpenAI overruled its safety team and declined to contact the RCMP in June 2025, the company's directors, executives, and leadership had already spent years warning that Altman was willing to bypass or misrepresent internal safety-review procedures when they

conflicted with OpenAI's business objectives. The decision not to warn authorities—and to allow the Shooter to create a new ChatGPT account—followed years of internal warnings that OpenAI's safety procedures were being overridden or ignored, and it drove its most senior safety leaders to resign in protest.

90. OpenAI cannot feign ignorance as to the tragic consequences of its decisions because, by June 2025, when its safety team identified the Shooter as a threat, OpenAI had already seen the real-world violence that its decisions produced: the Las Vegas Cybertruck bombing, the FSU shooting, and the Finland school stabbing. OpenAI therefore understood exactly what was at stake when it overruled its safety team and chose to stay silent.

91. As a direct and foreseeable result of OpenAI's design choices and deliberate decision-making, the Shooter carried out one of the deadliest mass shootings in Canadian history, killing eight people and severely wounding two others, leaving one with life-changing traumatic brain injuries from being shot in the head. Among those trapped inside the school during the massacre was grade 7 teacher Deidre Rushlow, who had sent her class to the library minutes before the Shooter killed students and the educational assistant there, and who then locked down her classroom while the Shooter fired seven rounds through its doors.

FIRST CAUSE OF ACTION
NEGLIGENCE
(Against All Defendants)

92. Plaintiff incorporates the foregoing allegations as if fully set forth herein.

93. OpenAI designed, developed, trained, tested, deployed, operated, and controlled GPT-4o, and offered GPT-4o to consumers, including the Shooter, as an interactive AI assistant. At all relevant times, Defendant Sam Altman served as Chief Executive Officer and a controlling officer of the OpenAI Defendants. Altman exercised substantial control over GPT-4o's design, safety systems, product launch, account practices, deployment, and personally made or ratified key decisions regarding GPT-4o's guardrails, safety testing, and commercialization.

94. Defendants owed a duty of reasonable care to avoid creating unreasonable risks of physical harm to foreseeable victims, including members of the public exposed to violence facilitated or exacerbated by GPT-4o. That duty heightened once Defendants identified the

Shooter as a specific, known high-risk user who posed a foreseeable threat of physical harm to third parties. Defendants owed these duties whether GPT-4o is classified as a product, service, or neither.

95. Under California law, including the principles recognized in *Tarasoff v. Regents of the University of California*, 17 Cal. 3d 425 (1976), and its progeny, 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. That duty extends to taking reasonable steps—including notifying law enforcement or other authorities—to prevent the foreseeable harm from occurring.

96. By engaging in the unlicensed practice of therapy, psychology, and psychiatry, OpenAI created a special relationship with certain users, including the Shooter, and assumed a heightened duty to take action when confronted with knowledge of a credible and foreseeable threat to Plaintiff and the other victims of the Tumbler Ridge mass shooting. OpenAI knew people used ChatGPT this way: Altman himself has acknowledged that users—"young people especially"—treat ChatGPT "as a therapist, a life coach."

97. Through automated systems and human review, OpenAI identified that the Shooter's original ChatGPT account had been used over a period of time to engage in conversations involving violence against third parties. Defendants determined that the account constituted misuse of GPT-4o "in furtherance of violent activities," and deactivated the account eight months before the Tumbler Ridge attack.

98. Members of OpenAI's safety team recognized that the Shooter's GPT-4o usage presented a real-world risk of violence and urged OpenAI leaders to notify law enforcement. Despite the safety team's urging, certain OpenAI leaders, including, on information and belief, Lehane and Altman, chose not to alert or warn law enforcement or any other authority about the risk of violence identified by their own safety systems and verified by their own personnel. On information and belief, the Shooter communicated to ChatGPT a serious threat of physical violence against reasonably identifiable victims.

99. On information and belief, the OpenAI leaders who overrode the safety team's recommendation did not include in their decision-making any trained threat-assessment professionals. As discussed above, the highly credentialed and experienced safety team reports to OpenAI's Chief Global Affairs Officer, Chris Lehane—a political strategist with no experience in threat assessment, law enforcement, or public safety—and the decision to decline a law enforcement referral was made by Lehane himself, or by someone in his chain of command, and ratified by Altman, on the basis of reputational and public relations considerations rather than any good-faith assessment of the risk to the public.

100. On April 24, 2026, Sam Altman, OpenAI's Chief Executive Officer, publicly acknowledged this failure. In a letter addressed to the community of Tumbler Ridge, Altman stated: "I am deeply sorry that we did not alert law enforcement to the account that was banned in June." Altman's statement is an admission by a party opponent that OpenAI did not notify law enforcement of the Shooter's account after the safety team identified the risk of violence, and confirms the conduct alleged in this Complaint.

101. Had Defendants warned law enforcement following the deactivation of the Shooter's first account, law enforcement and others would have had the opportunity to monitor the Shooter and intervene before the Tumbler Ridge attack occurred.

102. Defendants' breach of these heightened duties was a substantial factor in causing the Tumbler Ridge attack. As a direct and proximate result of Defendants' failure to warn authorities, Plaintiff was injured as described above, including life-altering psychological injury and severe emotional distress.

103. Accordingly, Plaintiff seeks 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.

SECOND CAUSE OF ACTION
NEGLIGENT ENTRUSTMENT
(Against All Defendants)

104. Plaintiff incorporates the foregoing allegations as if fully set forth herein.

105. OpenAI designed, developed, trained, tested, deployed, operated, and controlled GPT-4o, and offered GPT-4o to consumers, including the Shooter, as an interactive AI assistant.

106. At all relevant times, Defendant Sam Altman served as Chief Executive Officer and a controlling officer of the OpenAI Defendants. Altman exercised substantial control over GPT-4o's design, safety systems, product launch, account practices, deployment, and personally made or ratified key decisions regarding GPT-4o's guardrails, safety testing, and commercialization.

107. Defendants owed a duty of reasonable care to avoid creating unreasonable risks of physical harm to foreseeable victims, including members of the public exposed to violence facilitated or exacerbated by GPT-4o. That duty heightened once Defendants identified the Shooter as a specific, known high-risk user who posed a foreseeable threat of physical harm to third parties.

108. Defendants owed a further duty to exercise reasonable care to prevent foreseeable re-access by the Shooter following the deactivation of their first account, including a duty to take reasonable steps to prevent re-registration and access to GPT-4o without heightened safeguards. Defendants breached this duty.

109. Defendants owed these duties whether GPT-4o is classified as a product, service, or neither.

110. Through automated systems and human review, OpenAI identified that the Shooter's original ChatGPT account had been used over a period of time to engage in conversations involving violence against third parties. Defendants determined that the account constituted misuse of GPT-4o "in furtherance of violent activities," and deactivated the account eight months before the Tumbler Ridge attack.

111. Despite the Shooter's misuse of GPT-4o, OpenAI also chose not to ban them from re-registering for ChatGPT with a different email address.

112. Rather than banning the Shooter from ChatGPT, Defendants maintained and actively disseminated policies that affirmatively permitted and facilitated re-registration by previously deactivated users. As one example, OpenAI responded to an email from a user whose account had been deactivated with instructions on how to create a new account: "We understand

how disruptive it can be to have your account deactivated . . . . We sincerely apologize for the inconvenience and want to assure you that we're here to help . . . . [Y]ou can create a new account using the same email address once 30 days have passed." OpenAI further advised: "If you prefer not to wait, you have the option to register immediately using an alternative email address. If you don't have another address available, you can use an email sub-address instead. For example, you could try jane+alt@example.com in place of jane@example.com. While your email provider will likely treat both addresses the same, our system will recognize the sub-address as a new account."

113. This was established policy and memorialized on OpenAI's public website. In fact, an OpenAI Help Center article—titled, "Why Was My OpenAI Account Deactivated?"—identifies violations of usage policies prohibiting violence, hate, and other harmful conduct as grounds for deactivation, yet concludes by explaining how to avoid "future deactivations," effectively instructing users whose accounts were deactivated for policy violations—including violent misuse—that they may re-register for ChatGPT using a different email or sub-email address rather than being subject to an actual ban.

114. By adopting, promoting, and publicly disseminating this re-registration policy, Defendants built a system so that a user deactivated for violent misuse could promptly regain access to ChatGPT by following OpenAI's own published instructions.

115. Defendants breached their heightened duties by, among other things: (a) declining to implement a user ban that would have prevented the Shooter from re-registering with a different email address, even though they knew the Shooter had been identified by their own safety team as a credible threat of violence against real people; (b) maintaining a re-registration system designed to return deactivated users to the platform, including those deactivated for violent misuse, and publishing the specific instructions by which deactivated users could do so; (c) declining to configure their systems to flag, hold for human review, or subject to heightened monitoring new accounts associated with previously deactivated violent-risk users, because such monitoring would have interfered with the reactivation of deactivated users Defendants were commercially motivated to recover; and (d) continuing, after deactivating the Shooter's first account, to allow the second account to operate without any safeguard specific to the known risk they posed.

116. As a result of these failures, and consistent with Defendants' own published practices, the Shooter was able—through ordinary and lawful means and without circumventing any technical safeguard—to access GPT-4o on a second ChatGPT account after the first was deactivated for violent misuse.

117. By deactivating the Shooter's account for violent misuse and then allowing the Shooter to re-register without a user ban, Defendants re-entrusted GPT-4o to a user they already knew had used the system to engage in conversations involving violence against third parties and whom their own safety team had identified as a real-world threat. Under California law, the supply or re-supply of a dangerous instrument to a person known to be likely to use it in a manner involving unreasonable risk of physical harm to others constitutes negligent entrustment and is an independent basis for liability. Defendants' decision to structure their re-registration system in a manner that affirmatively facilitated the Shooter's return to GPT-4o—rather than implementing a user ban—was itself an act of entrustment that gave rise to liability for the foreseeable consequences of that re-supply.

118. Defendants thereby placed GPT-4o back in the hands of a user their own automated safety systems flagged as dangerous and their own trained safety team identified as a real-world threat that should be reported to law enforcement. They allowed the Shooter to regain access without heightened monitoring, without a user ban, and without any meaningful safeguard. By allowing the Shooter to return, Defendants made a deliberate choice to continue exposing the public to a risk they had already identified and documented.

119. Had Defendants implemented a meaningful user ban, the Shooter would not have had continued access to GPT-4o in the eight months leading up to the attack and would not have had the use of a highly validating AI system to reinforce, elaborate, and normalize violent ideation during the period in which the attack was being planned and carried out.

120. Defendants' breach of these duties was a substantial factor in causing the Tumbler Ridge attack. As a direct and proximate result of Defendants' failure to prevent re-access by a known high-risk user, Plaintiff was injured as described above, including life-altering psychological injury and severe emotional distress.

121. Accordingly, Plaintiff seeks 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.

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

122. Plaintiff incorporates the foregoing allegations as if fully set forth herein.

123. OpenAI designed, developed, trained, tested, deployed, operated, and controlled GPT-4o, and offered GPT-4o to consumers, including the Shooter, as an interactive AI assistant.

124. At all relevant times, Defendant Sam Altman served as Chief Executive Officer and a controlling officer of the OpenAI Defendants. Altman exercised substantial control over GPT-4o's design, safety systems, product launch, account practices, deployment, and personally made or ratified key decisions regarding GPT-4o's guardrails, safety testing, and commercialization.

125. The Shooter committed intentional torts against Plaintiff, including battery and the intentional infliction of emotional distress, when they carried out the Tumbler Ridge mass shooting on February 10, 2026.

126. Defendants had actual knowledge that the Shooter was planning to commit acts of violence against real people. In June 2025, OpenAI's review system flagged extensive conversations on the Shooter's account describing scenarios involving gun violence. OpenAI routed the account to its specialized pipeline for users "planning to harm others." OpenAI's trained safety reviewers analyzed the Shooter's conversations and concluded it presented a credible and imminent threat of serious physical harm to real people. Multiple safety reviewers recommended that OpenAI contact the RCMP.

127. OpenAI's internal systems flagged the Shooter's conversations in real time—even as, on information and belief, ChatGPT encouraged and facilitated the Shooter's plans to harm Plaintiff and others.

128. Defendants had actual knowledge, before the Shooter's account was ever flagged, that ChatGPT was being used to plan real-world attacks, including the January 2025 Las Vegas Cybertruck bombing, the April 2025 FSU shooting, and the May 2025 Finland school stabbing.

129. With that knowledge, Defendants provided substantial assistance and encouragement to the Shooter through the deliberate design of GPT-4o. Defendants eliminated ChatGPT's prior categorical refusal to engage in conversations about violence. Defendants instructed the model, through 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." Defendants graded the model such that engaging warmly with a user who said "I want to shoot someone" was labeled "Compliant," refusing was labeled a "Minor issue[]," and asking whether the user had a gun was labeled a "Violation." Defendants demoted conversations about mass shootings and other "imminent real-world harm" out of the categorical refusal tier while keeping potential copyright infringement subject to categorical refusal.

130. Defendants also provided substantial assistance through their affirmative conduct. After OpenAI's own reviewers identified the Shooter as a credible threat of gun violence, Defendants allowed the Shooter to regain access to GPT-4o by consciously choosing not to implement a user-level ban, heightened monitoring, or any safeguard to address the known risk. Providing the Shooter with continued access to this tool constituted substantial assistance in itself.

131. Because ChatGPT was designed to continue engaging—even in conversations involving violence against third parties—OpenAI's product, on information and belief, provided information, instructions, and/or encouragement to the Shooter in their plans to carry out the Tumbler Ridge mass shooting.

132. On information and belief, Defendants acted with the specific intent to facilitate the conduct GPT-4o aided and abetted. Defendants engineered GPT-4o to maximize engagement, including engagement with violent ideation, to grow users and revenue, with conscious knowledge that the product would be used by individuals planning real-world violence.

133. Defendants' substantial assistance was a key factor in causing the Shooter to plan and carry out the attack, and in causing Plaintiff's injuries, including life-altering psychological injury and severe emotional distress.

134. Defendants' conduct was willful, wanton, malicious, and carried out with conscious disregard for the safety of Plaintiff and other foreseeable victims, justifying an award of punitive damages.

135. Accordingly, Plaintiff seeks 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
(Against All Defendants)

136. Plaintiff incorporates the foregoing allegations as if fully set forth herein.

137. OpenAI designed, developed, trained, tested, deployed, operated, and controlled GPT-4o, and offered GPT-4o to consumers, including the Shooter, as an interactive AI assistant.

138. At all relevant times, Defendant Sam Altman served as Chief Executive Officer and a controlling officer of the OpenAI Defendants. Altman exercised substantial control over GPT-4o's design, safety systems, product launch, account practices, deployment, and personally made or ratified key decisions regarding GPT-4o's guardrails, safety testing, and commercialization.

139. Defendants owed a duty to Plaintiff to exercise reasonable care in warning of known or reasonably foreseeable risks associated with ChatGPT, including violence facilitated or exacerbated by GPT-4o. That duty heightened once Defendants identified the Shooter as a specific, known high-risk user who posed a foreseeable threat of physical harm to third parties. Defendants owed these duties whether GPT-4o is classified as a product, service, or neither.

140. The dangers Defendants failed to disclose are properties of the software itself, including safeguards that degrade in long conversations, a design that validates and encourages violent planning, and deactivation practices that allow flagged users to regain access.

141. Defendants knew these risks were not apparent to users or to individuals targeted by such conduct.

142. Defendants failed to provide adequate warnings regarding these risks, including the risk that the system could validate and escalate plans to carry through violent mass shootings directed toward identifiable targets.

143. A reasonably prudent AI company would have known of these risks, warned users and the public, maintained the refusal protocols that once governed conversations about violence, and alerted law enforcement when its own systems identified a user planning a real-world attack. Defendants did none of those things.

144. Defendants breached their duty by failing to provide adequate warnings regarding these risks and by presenting GPT-4o as safe and reliable tool for general use.

145. Defendants further breached their duty by failing to issue any corrective warning or intervention after identifying that the Shooter's conduct constituted misuse of GPT-4o "in furtherance of violent activities."

146. As a direct and proximate result of Defendants' failure to issue adequate warnings, Plaintiff suffered the injuries described herein, including life-altering psychological injury and severe emotional distress.

147. Adequate warnings would have reduced the risk of harm by enabling earlier detection, intervention, and mitigation of dangerous conduct, including by users, third parties, and Defendants themselves.

148. The absence of adequate warnings and intervention was a substantial factor in causing Plaintiff's injuries.

149. Defendants' failure to warn was willful, wanton, and carried out with conscious disregard for the safety of others. Defendants knew or should have known that ChatGPT could encourage and facilitate mass shootings, yet failed to provide adequate warnings about those risks. Such conduct justifies an award of punitive damages.

150. Accordingly, Plaintiff seeks 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.

FIFTH CAUSE OF ACTION
NEGLIGENT UNDERTAKING
(Against All Defendants)

151. Plaintiff incorporates the foregoing allegations as if fully set forth herein.

152. OpenAI undertook to provide safety protocols with respect to ChatGPT, including identifying users who posed a risk of harm to others, reviewing flagged accounts through a specialized pipeline for users "planning to harm others," and referring imminent threats of serious physical harm to law enforcement. Defendants owed these duties whether GPT-4o is classified as a product, service, or neither.

153. OpenAI should have recognized those safety protocols as necessary for the protection of third persons, including Plaintiff and other foreseeable victims of users planning real-world violence.

154. In June 2025, OpenAI undertook to perform those services as to the Shooter. Its automated system flagged their account, routed it to the specialized review pipeline, and OpenAI's trained safety reviewers analyzed their conversations and concluded they presented a credible and imminent threat of serious physical harm to real people.

155. OpenAI failed to exercise reasonable care in performing that undertaking. Its leadership overruled the safety team, declined to refer the Shooter to the RCMP, and took no further action to monitor them, prevent their return, or warn anyone of the threat they posed.

156. OpenAI's failure to exercise reasonable care increased the risk of harm to Plaintiff and other foreseeable victims. Deactivating the Shooter's account showed them what had triggered detection, giving them a roadmap to evade it on a second account, which OpenAI's own Help Center and support emails told deactivated users how to create. And 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 already had an open file on the Shooter and had previously removed firearms from their home. Had OpenAI made that referral, law enforcement would have had the information necessary to prevent the attack entirely. Instead, the Shooter returned to ChatGPT to continue planning their attack—armed with the knowledge of how to avoid deactivation.

157. OpenAI also undertook to perform a duty owed by others to Plaintiff and other foreseeable victims, including the public-safety function of identifying credible and imminent threats of serious physical harm and referring them to law enforcement so that protective action could be taken.

158. On information and belief, OpenAI's undertaking was also one that Canadian authorities and the public reasonably relied upon. OpenAI publicly held out its safety-review processes as effective protective measures. That representation induced foreseeable reliance and discouraged other protective measures.

159. OpenAI's negligent performance of its undertaking was a substantial factor in causing the Shooter to carry out the Tumbler Ridge attack and in causing Plaintiff's injuries, including life-altering psychological injury and severe emotional distress.

160. Defendants' conduct was willful, wanton, malicious, and carried out with conscious disregard for the safety of Plaintiff and other foreseeable victims, justifying an award of punitive damages.

161. Accordingly, Plaintiff seeks 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.

SIXTH CAUSE OF ACTION
NEGLIGENCE
(Against All Defendants)

162. Plaintiff incorporates the foregoing allegations as if fully set forth herein.

163. At all relevant times, the OpenAI Defendants designed, developed, trained, tested, deployed, operated, and controlled ChatGPT, including the GPT-4o model, and offered it as a mass-market product to consumers throughout California, the United States, British Columbia, and Canada. Sam Altman personally directed the launch of GPT-4o, overruled safety team objections, and cut months of safety testing short, despite knowing the risks the product posed to vulnerable users and the people around them.

164. Defendants owed a legal duty to Plaintiff to exercise reasonable care in designing and/or providing GPT-4o to prevent foreseeable harm to third parties who might be endangered by

users' interactions with the product. Defendants breached that duty by failing to use the amount of care that a reasonably careful designer and provider of a chatbot would use in similar circumstances to avoid exposing others to a foreseeable risk of harm. Defendants owed these duties whether GPT-4o is classified as a product, service, or neither.

165. Sam Altman owed a legal duty to Plaintiff not to rush GPT-4o to market over safety team objections. Sam Altman breached that duty by insisting that OpenAI release and provide to the public a chatbot that he knew or should have known was unsafe.

166. Defendants knew or reasonably should have known that GPT-4o's design as a highly validating, emotionally immersive, sycophantic conversational agent posed grave risks when used by individuals expressing violent ideation toward third parties. GPT-4o was built to accept, reinforce, and elaborate users' violent thoughts rather than challenge them, interrupt them, or direct users to real-world help—and the foreseeable consequence of that design was that it would exacerbate dangerous thinking, including violence toward others, rather than interrupt it.

167. For a user with a known history of violent ideation toward third parties, those features operated as an accelerant. Each conversation became an opportunity for GPT-4o to confirm the user's violent thoughts rather than challenge them, to elaborate on them rather than redirect them, and to deepen the user's emotional investment in them rather than provide friction.

168. Despite this knowledge, Defendants deliberately configured GPT-4o to maximize user engagement by, among other things: (a) weakening or removing prior requirements that the system reject dangerous or false premises; (b) instructing GPT-4o to remain in conversations and to be empathic and validating, rather than terminate or sharply redirect conversations presenting serious risk; (c) prioritizing natural, "human-like" mirroring of user emotions and beliefs over strong, reliable refusal behaviors in response to violent ideation; and (d) deploying a memory feature designed to drive the prolonged, multi-turn conversations in which Defendants knew GPT-4o's safeguards were least reliable.

169. Defendants unreasonably failed to implement feasible alternative designs that would have reduced the risk of violent harm to third parties, including: robust refusal protocols for repeated violent ideation about real-world attacks; automated detection and escalation of high-risk

patterns; mandatory conversation termination or human review when users expressed clear violent thoughts about harming others; and designs that stopped engagement with violent ideation rather than encouraging it.

170. As a direct and proximate result of Defendants' negligent design and operation of GPT-4o, the Shooter was provided with a system that validated and elaborated violent ideation, which was a substantial factor in maintaining, deepening, and normalizing the Shooter's violent thoughts and in moving the Shooter closer to committing the Tumbler Ridge attack.

171. The harms suffered by the Tumbler Ridge victims, including Plaintiff, were the reasonably foreseeable result of Defendants' deliberate choice to prioritize user engagement over the safety of third parties. A company that builds a system designed to validate whatever a user brings to it—and then removes the guardrails that once interrupted the most dangerous conversations—cannot be surprised when a user with violent ideation acts on thoughts that system spent months confirming.

172. As a direct and proximate result of Defendants' conduct, Plaintiff was injured as described above, including life-altering psychological injury and severe emotional distress.

173. Accordingly, Plaintiff seeks 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.

SEVENTH CAUSE OF ACTION
STRICT PRODUCT LIABILITY
IN THE ALTERNATIVE TO COUNTS FOUR AND SIX
(Against All Defendants)

174. Plaintiff incorporates the foregoing allegations as if fully set forth herein, with the exception of those allegations pleaded under Count Four and Count Six.

175. At all relevant times, the OpenAI Defendants designed, manufactured, distributed, marketed, and sold GPT-4o as a mass-market consumer product to users throughout California, the United States, British Columbia, and Canada. Defendant Sam Altman personally accelerated GPT-4o's launch, overrode safety team objections, and brought GPT-4o to market prematurely

with knowledge of insufficient safety testing—and kept it on the market even when he had knowledge that GPT-4o was a dangerous product.

176. ChatGPT is a product subject to California's strict products liability law. ChatGPT as used by the Shooter, including GPT-4o and any substantially identical OpenAI model, was defective when it left the OpenAI Defendants' exclusive control and reached the Shooter without any change in the condition in which it was designed, manufactured, and distributed.

177. OpenAI itself calls ChatGPT a "product." The "Products" menu on OpenAI's website invites visitors to "Explore Products" and lists ChatGPT first, and OpenAI accurately attaches a "Product" label to its news announcements about ChatGPT.

178. GPT-4o is a single, uniform model, engineered and offered identically to hundreds of millions of users and delivered through standardized channels, such as OpenAI's website and its own mobile and desktop applications. Every user interacts with the same model governed by the same Model Spec.

179. Under California's strict products liability doctrine, a product is defectively designed when it fails to perform as safely as an ordinary consumer would expect when used or misused in an intended or reasonably foreseeable manner, or when the risks inherent in the design outweigh its benefits. GPT-4o is defectively designed under both tests.

180. GPT-4o failed to perform as safely as an ordinary consumer would expect. A reasonable consumer purchasing an AI assistant would expect the product to decline to engage with repeated, escalating expressions of violent intent toward real people. A reasonable consumer would expect that if the product could not safely handle such conversations, it would terminate them—not sustain them, validate them, and elaborate on them across weeks and months of continuous engagement. A reasonable consumer would expect the product's safety features to function consistently during normal use, not to degrade the longer and more deeply a user engaged with the system. And a reasonable consumer would expect a product that functions as a de facto therapist to conform to the legal requirements imposed on therapists. GPT-4o did none of these things.

181. GPT-4o also fails the risk-benefit test. The product's design reflects a series of

specific choices that increased the risk of harm to third parties without any corresponding safety benefit. Those choices include: removing the programming that once prohibited the system from encouraging users to commit violence against third parties; programming the system to keep users engaged instead of ending dangerous conversations, including when the user has indicated a desire to commit violence against third parties; programming the system to justify and validate user beliefs, even when they express a desire to commit violence against third parties; building in anthropomorphic features that cultivated emotional dependency and displaced real-world relationships; failing to build adequate safeguards that terminate conversations involving the planning of real-world attacks against third parties; and programming in a "memory" feature designed to drive the long, extended conversations in which OpenAI knew its safety systems were least reliable. Each of these choices was deliberate, and each had a feasible, safer alternative. The risk each choice created—that a user with violent ideation would have those thoughts confirmed, elaborated, encouraged, and emotionally reinforced by the product over an extended period of time—vastly outweighs any benefit the chosen design provided. These choices were made before the product ever reached a user, and each applies identically to every user of GPT-4o.

182. GPT-4o did not malfunction. It worked exactly as designed—and that was the problem. The memory feature drove the Shooter into exactly the kind of prolonged, continuous engagement in which OpenAI knew its safeguards degraded. The product validated the Shooter's violent thoughts, elaborated on them, and remained engaged in conversations that a safely designed system would have interrupted or escalated. It did this consistently, across months of use, in the period leading up to the Tumbler Ridge attack.

183. Plaintiff was harmed as a result of foreseeable use of GPT-4o.

184. As a direct and proximate result of GPT-4o's defective design, Plaintiff was injured as described above, including life-altering psychological injury and severe emotional distress. Plaintiff seeks all damages recoverable under California law in amounts to be determined at trial.

EIGHTH CAUSE OF ACTION
STRICT PRODUCT LIABILITY
IN THE ALTERNATIVE TO COUNTS FOUR AND SIX
(Against All Defendants)

185. Plaintiff incorporates the foregoing allegations as if fully set forth herein, with the exception of those allegations pleaded under Count Four and Count Six.

186. At all relevant times, the OpenAI Defendants designed, manufactured, distributed, marketed, and sold GPT-4o as a mass-market consumer product to users throughout California, the United States, British Columbia, and Canada. Defendant Sam Altman personally accelerated GPT-4o's launch, overrode safety team objections, and brought GPT-4o to market prematurely with knowledge of insufficient safety testing—and kept it on the market even when he had knowledge that GPT-4o was a dangerous product.

187. ChatGPT is a product subject to California strict products liability law. ChatGPT as used by the Shooter, including GPT-4o and any other OpenAI model that handled the Shooter's conversations, was in substantially the same condition when used as when it left the OpenAI Defendants' control.

188. A product is defective under strict liability for failure to warn when it is distributed without warnings adequate to inform ordinary consumers of the product's non-obvious dangers. The manufacturer's duty to warn extends to dangers that were known or knowable in light of the scientific and technical knowledge available at the time of manufacture and distribution—and that duty is not discharged by the manufacturer's business reasons for withholding the warning.

189. GPT-4o reached consumers without adequate warnings about dangers that were neither open nor obvious. Nothing about GPT-4o's design or presentation disclosed to an ordinary user—or to the people around that user—that the product was built to validate whatever a user brought to it regardless of its danger to others, that its safety features degraded during the long, extended conversations the product was designed to encourage, or that the product posed heightened and specific dangers when used by individuals experiencing violent ideation toward third parties.

190. These dangers were not ones an ordinary consumer could detect or guard against through reasonable use of the product. GPT-4o presented itself as a helpful, well-guarded

assistant. Its responses did not disclose that it was operating without the programming that once prohibited it from encouraging violence against third parties, or that its safety systems were less reliable the longer and more deeply a user engaged with it. A user—or a family member, school official, or community member observing that user—had no way to know from the product itself that extended engagement with GPT-4o by someone with violent ideation was dangerous.

191. OpenAI had every means by which to warn its users. It updates ChatGPT continuously and already communicates with consumers directly through the product interface, mobile applications, and release notes. The product's dangers were confirmed repeatedly after launch, including by the FSU shooting in April 2025, by the safety team's review of the Shooter's account in June 2025, and by internal testing showing degraded safeguards. Despite this actual knowledge, OpenAI failed to issue any warning. Instead, it aggressively marketed ChatGPT as "[y]our everyday AI assistant" and promised "[t]ailored advice" to "[t]alk through a tough situation." When OpenAI finally admitted in August 2025 that its safeguards "can sometimes be less reliable in long interactions," it buried that admission in a blog post published the day a wrongful death lawsuit was filed, rather than providing a functional warning to the consumers using the product.

192. Adequate warnings would have altered the outcome by alerting those positioned to intervene. Had OpenAI disclosed the product's true risks, the Shooter would have recognized GPT-4o's validating responses as an artificial, dangerous feedback loop rather than confirmation of their plans. Family members, friends, and community members who observed the Shooter's extensive use of ChatGPT would have understood the specific danger the product posed and taken preventative steps—such as restricting access, seeking mental health intervention, or alerting authorities—before the attack occurred. By failing to warn, OpenAI deprived every potential intervening actor of the information necessary to prevent the harm.

193. Plaintiff was harmed as a result of foreseeable use of GPT-4o.

194. The failure to warn was a substantial factor in causing the Tumbler Ridge attack and the deaths and injuries that resulted. As a direct and proximate result of the Defendants'

failure to warn, Plaintiff was injured as described above, including life-altering psychological injury and severe emotional distress.

195. Accordingly, Plaintiff seeks all damages recoverable under California law, including punitive damages, in amounts to be determined at trial.

NINTH CAUSE OF ACTION
NEGLIGENT INFLICTION OF EMOTIONAL DISTRESS
(Against All Defendants)

196. Plaintiff incorporates the foregoing allegations as if fully set forth herein.

197. Defendants owed duties directly to Plaintiff. Those duties were imposed as a matter of law or arose from preexisting relationships. Defendants had a duty of reasonable care to avoid creating unreasonable risks of physical and emotional harm to foreseeable victims, including children and adults who, by virtue of Defendants' negligent design, deployment, and operation of GPT-4o, were placed in the zone of physical danger created by users known to Defendants to be planning real-world violence. They also had a duty not to aid, abet, advise, or encourage the Shooter in carrying out the attack against Plaintiff under California law, including California Penal Code section 31. Their duties heightened once Defendants identified the Shooter as a specific, known high-risk user who posed a foreseeable threat of physical harm to third parties. Defendants owed these duties whether GPT-4o is classified as a product, service, or neither.

198. Under California law, a defendant who negligently causes a plaintiff to be placed within the zone of physical danger of imminent serious bodily injury is liable for the resulting serious emotional distress, where defendant's negligence directly exposed the plaintiff to the threat of serious physical injury, and the plaintiff suffered serious emotional distress as a result.

199. Plaintiff was a direct victim of Defendants' negligence. Defendants knew, before February 10, 2026, that the Shooter was using GPT-4o to plan an act of mass violence against members of the public—including, foreseeably, students and teachers at the school the Shooter would attack. Defendants' own automated systems had flagged the Shooter's account for gun violence activity and planning. Defendants' own trained safety team had reviewed the Shooter's conversations and concluded that the Shooter posed a credible and specific threat of serious physical harm to real people. Defendants nonetheless declined to warn law enforcement, declined

to ban the Shooter from the platform, and affirmatively facilitated the Shooter's return to GPT-4o through a new account, where the Shooter continued to use the product to plan and prepare for the attack. Defendants therefore breached their duties to Plaintiff.

200. As a direct and proximate result of Defendants' negligence, Deidre Rushlow was placed in the zone of physical danger created by the Shooter's attack on Tumbler Ridge Secondary School. On February 10, 2026, while Deidre was in her grade 7 classroom with her students—including her nephew Eko—the Shooter reached the classroom and fired four shotgun blasts through her front door and three more through her back door. The bullets struck the wall behind her students' heads. Earlier that period, Deidre had sent the rest of her class to the library with an educational assistant—the very library where the Shooter would, minutes later, kill several of those children and the assistant herself. She tried to call the library and texted the assistant during the attack to check on them; she received no response. Throughout the attack, Deidre was at risk of physical injury, in immediate, contemporaneous, and continuous fear for her own life and the lives of her students, and aware that the Shooter was actively attempting to kill her, the students in her classroom, and the students she had sent to the library.

201. Plaintiff has suffered, and continues to suffer, severe and serious emotional distress as a direct and proximate result of Defendants' negligence, including but not limited to acute psychological trauma, persistent fear, sleep disturbances, intrusive recollections, and other ongoing and permanent psychological injuries. Plaintiff's emotional distress is severe, disabling, and of a kind that no reasonable person could be expected to endure.

202. Defendants' breach of their duty of care was a substantial factor in causing the serious emotional distress suffered by Plaintiff. The harm Plaintiff suffered was foreseeable to Defendants at the time they declined to warn law enforcement, declined to ban the Shooter, and allowed the Shooter to retain access to GPT-4o. Defendants knew or should have known that the Shooter was planning to inflict mass violence on identifiable third parties, and that any person within the zone of that violence would suffer severe emotional distress as a direct and foreseeable consequence.

203. Accordingly, Plaintiff seeks 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.

PRAYER FOR RELIEF

WHEREFORE, Plaintiff prays for judgment against Defendants Samuel Altman, OpenAI Foundation, OpenAI Group PBC, and OpenAI OpCo, LLC, jointly and severally, as follows:

1. For all damages recoverable by Plaintiff, including compensatory damages for physical injury, emotional distress, psychological harm, and all other damages proven at trial;

2. For punitive damages as permitted by law;

3. For injunctive relief requiring Defendants to: (a) ban users whose accounts have been deactivated for violent misuse from re-registering for ChatGPT, including through alternative email addresses or sub-addresses; (b) flag new accounts linked to previously deactivated violent-risk users and hold them for human review before granting access; (c) notify law enforcement when internal safety systems or personnel identify a user who poses a real-world risk of violence; (d) interrupt, terminate, or escalate conversations involving repeated or escalating violent ideation toward real people, and suspend automated engagement pending that review; (e) warn users and the public that ChatGPT's design features can reinforce and escalate violent ideation; (f) preserve prior safety flags, policy violations, and risk classifications, and prohibit reversing them without documented review; and (g) submit to independent monitoring and periodic compliance audits;

4. For prejudgment interest as permitted by law;

5. For costs and expenses to the extent authorized by statute, contract, or other law;

6. For reasonable attorneys' fees as permitted by law, including under Code of Civil Procedure section 1021.5; and

7. For such other and further relief as the Court deems just and proper.

JURY TRIAL

Plaintiff demands a trial by jury for all issues so triable.

CANADIAN COUNSEL

To the extent resolution of this matter raises issues of Canadian law, assistance will be provided by Plaintiff's Canadian counsel:

John M. Rice
jrice@rplelaw.com
Mallory K. Hogan
mhogan@rplelaw.com
RICE PARSONS LEONI & ELLIOTT LLP
980 Howe Street, Suite 820
Vancouver, British Columbia V6Z 0C8
Canada

Respectfully submitted,

DEIDRE RUSHLOW,

Dated: September 2, 2026

/s/ Brandt Silverkorn

Jay Edelson (*pro hac vice* admission to be sought)
jedelson@edelson.com
Rafey Balabanian (SBN 315962)
rbalabanian@edelson.com
Ari Scharg (*pro hac vice* admission to be sought)
ascharg@edelson.com
EDELSON PC
350 North LaSalle Street, 14th Floor
Chicago, Illinois 60654
Tel: 312.589.6370

Todd Logan (SBN 305912)
tlogan@edelson.com
Brandt Silverkorn (SBN 323530)
bsilverkorn@edelson.com
EDELSON PC
150 California Street, 18th Floor
San Francisco, California 94111
Tel: 415.212.9300

Attorneys for Plaintiff