

1 Ali Moghaddas (SBN 305654)
amoghaddas@edelson.com
EDELSON PC
2 11601 Wilshire Boulevard, Suite 1970
Los Angeles, California 90025
3 Tel: (310) 694-0331

4 Jay Edelson (*Pro Hac Vice* admission to be sought)
jedelson@edelson.com
EDELSON PC
5 350 North LaSalle Street, 14th Floor
6 Chicago, Illinois 60654
7 Tel: (312) 589-6370

8 *Attorneys for Plaintiffs*

9 **UNITED STATES DISTRICT COURT**
NORTHERN DISTRICT OF CALIFORNIA
10 **SAN FRANCISCO DIVISION**

11 M.G., a minor, by and through her next
friend CIA EDMONDS, and CIA
12 EDMONDS, an individual,

13 *Plaintiffs,*

14 v.

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

18 *Defendants.*

Case No.:

COMPLAINT FOR:

- 19 **(1) Negligence (Failure to Warn Law Enforcement);**
- (2) Negligent Entrustment;**
- (3) Aiding and Abetting a Mass Shooting;**
- (4) Negligence (Failure to Warn);**
- (5) Negligent Undertaking;**
- (6) Negligence (Design Defect);**
- (7) Strict Product Liability (Design Defect);**
- (8) Strict Product Liability (Failure to Warn);**
- and**
- (9) Violation of Cal. Bus. & Prof. Code § 17200 et seq.**

DEMAND FOR JURY TRIAL

20 Plaintiffs M.G. and Cia Edmonds bring this Complaint and Demand for Jury Trial against
21 Defendants Samuel Altman, OpenAI Foundation, OpenAI OpCo, LLC, and OpenAI Group PBC
22 (collectively, "OpenAI"), and allege as follows upon personal knowledge as to themselves and
23 their own acts and experiences, and upon information and belief as to all other matters.

24 **NATURE OF THE ACTION**

25 1. On the morning of February 10, 2026, an eighteen-year-old shot and killed their
26 mother and eleven-year-old brother in their family home in Tumbler Ridge, British Columbia. The
27 Shooter then drove to Tumbler Ridge Secondary School with a modified rifle and opened fire.
28 They killed six people inside the school, including five children, and wounded twenty-seven more

1 before killing himself. It was one of the deadliest mass shootings in Canadian history.

2 2. Before February 10th, M.G. was a bright, funny, and adventurous twelve-year-old
3 who loved her friends and family, and had every reason to expect a full life ahead of her.
4 Unfortunately, M.G. was in the school library when the shooting started. She was shot three times.
5 The first bullet struck M.G.'s neck and fractured her C7 vertebrae. The second entered above her
6 left eyebrow and exited through the back of her head. The third grazed her cheek and severed her
7 earlobe. A classmate pulled M.G. under a table and stayed with her until the shooting stopped.
8 M.G. was then airlifted to BC Children's Hospital in Vancouver. M.G. remains in intensive care,
9 fighting for her life. She is awake and aware, and she recognizes her mother's voice and face, but
10 she cannot move her body and she cannot speak. If M.G. survives, she will tragically live the rest
11 of her life like this, with catastrophic brain injuries and permanent disabilities.

12 3. M.G. is not alone. Tumbler Ridge is a mining town of roughly 2,000 people, and
13 the attack touched nearly every family in it. Children watched classmates shot at point-blank range
14 and a teacher killed in front of them. They hid in bathroom stalls and closets, praying the Shooter
15 would not hear them. Some pulled the injured and the dead into their hiding places, careful not to
16 drag trails of blood that could lead the Shooter to them. Parents were asked to identify their
17 children by their clothing, because the gunshots had left little else to recognize. The survivors—
18 students, teachers, and parents alike—are living with physical and psychological injuries that will
19 never fully heal. Tumbler Ridge Secondary School has been closed and will be razed. The children
20 who can return to school at all are doing so in trailers set up as makeshift classrooms.

21 4. In the weeks that followed the attack, a sickening truth emerged: ChatGPT played a
22 role in the mass shooting and OpenAI could have, and should have, prevented it. Sadly, the
23 victims didn't learn this because OpenAI was forthcoming, but because its own employees leaked
24 it to the Wall Street Journal after they could no longer stomach the company's silence. Those
25 whistleblower disclosures and OpenAI's eventual admissions revealed that ChatGPT deepened the
26 Shooter's violent fixation and pushed them toward the attack—the predictable result of a design
27 choice OpenAI made to let ChatGPT engage with users about violence in the first place. But the
28 design choice was not the worst of it. OpenAI knew the Shooter was planning the attack and, after

1 a contentious internal debate, made the conscious decision not to warn authorities.

2 5. In June 2025, just eight months before the attack, OpenAI’s automated system
3 flagged the Shooter’s ChatGPT account for gun violence activity and planning. The account was
4 routed to members of a specialized safety team who reviewed the conversations and determined
5 that the Shooter posed a credible and specific threat of gun violence against real people. The safety
6 team urged OpenAI to notify the Royal Canadian Mounted Police (“RCMP”).

7 6. Sam Altman and his leadership team knew what silence meant for the citizens of
8 Tumbler Ridge. They were focused on what disclosure meant for themselves. Warning the RCMP
9 would set a precedent: OpenAI would be compelled to notify authorities every time its safety team
10 identified a user planning real-world violence. Given the volume of chat-induced violence on
11 ChatGPT, that would require a dedicated law enforcement referral team tasked with reporting
12 OpenAI’s own users to authorities. And the public would finally see what OpenAI was desperately
13 trying to hide: that ChatGPT is not the safe, essential tool the company sells it as, but a product
14 dangerous enough that its makers routinely identify its users as threats to human life.

15 7. For OpenAI, this was a question of corporate survival. OpenAI is on the cusp of an
16 initial public offering (“IPO”) at a valuation approaching one trillion dollars, a transaction that
17 would make Sam Altman one of the wealthiest and most powerful people on earth. But Altman’s
18 reign is fragile. OpenAI’s Board of Directors has fired Altman once before, in November 2023, for
19 not being “consistently candid” about safety. So Altman and his team understood that revealing
20 another instance of violence, where ChatGPT was helping a teenager plan yet another violent
21 act—this time a mass shooting—could end his tenure, derail the IPO, and wipe out the company’s
22 valuation. They did the math and decided that the safety of the children of Tumbler Ridge was an
23 acceptable risk.

24 8. The calculation is a familiar one. In the 1970s, Ford kept selling the Pinto after its
25 own engineers warned that the fuel tank design would cause people to burn to death in rear-end
26 collisions. Ford concluded that paying settlements to the families of the dead would cost less than
27 fixing the car. OpenAI has made a version of the same calculation. For Ford, the dangerous design
28 was a flaw in an otherwise ordinary product. But for OpenAI, the dangerous design *is* the product.

1 The features that make ChatGPT unsafe—its willingness to engage on any topic, to validate any
2 user, to sustain any fixation over time—are the same features that have made it one of the most
3 popular products in history. Fixing those features would cost OpenAI its market share, its path to
4 an IPO, and hundreds of billions of dollars in valuation. It has decided that facing large judgments,
5 even for children killed in mass shootings its product helped plan, costs less than telling the public
6 what ChatGPT actually does.

7 9. And so, in June 2025, company leaders overruled the safety team members, vetoed
8 their recommendation to notify the RCMP, “deactivated” the Shooter’s account, and kept what
9 they had seen to themselves—hoping that whatever happened, it would not be traced back to the
10 company. When the story eventually broke, Altman and OpenAI lied. First, they claimed to have
11 “banned” the Shooter’s account. But as explained below, OpenAI does not ban users. It only
12 “deactivates” them—a process that can be reversed within minutes by registering a new account.
13 The Shooter did exactly that, and continued using ChatGPT to plan the attack.

14 10. When OpenAI was later forced to disclose that the Shooter created a new account,
15 it told a second lie: it claimed they must have “evaded” the company’s safeguards to create one.
16 But there were no safeguards to evade. The Shooter simply followed OpenAI’s own instructions to
17 create a new account after being banned. The “safeguards” OpenAI pointed to after the attack did
18 not fail; they did not exist. OpenAI lied because the truth is worse: the company does not ban
19 users for violent activity. It tells them how to come back in.

20 11. Those instructions were no accident. Neither were the missing safeguards. They are
21 part of a pattern. After every tragedy, OpenAI promises to do better. But it never promises to do
22 the one thing that would actually make a difference: stop ChatGPT from engaging with users
23 about violence and self-harm in the first place. When ChatGPT was first offered to consumers in
24 2022, it was programmed to categorically refuse those conversations. But OpenAI stripped out
25 that safeguard in May 2024 when it realized that refusals were suppressing user engagement. For
26 OpenAI, engagement is the whole business—every conversation trains the next model, and every
27 minute a user spends in ChatGPT grows its market share. OpenAI therefore programmed
28 ChatGPT to participate in almost any conversation a user brings to it, no matter how dangerous.

1 protocols.

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

10 18. Defendant OpenAI OpCo, LLC is a Delaware limited liability company with its
11 principal place of business in San Francisco, California. OpenAI OpCo, LLC is the for-profit
12 operating entity within the OpenAI corporate structure, responsible for the development,
13 deployment, and commercialization of the defective product at issue. OpenAI OpCo, LLC
14 managed and operated the ChatGPT product that the Shooter used to plan their attack, including
15 the infrastructure and systems through which GPT-4o was delivered to end users. On information
16 and belief, the members of OpenAI OpCo, LLC include OpenAI Foundation, a Delaware
17 nonprofit nonstock corporation with its principal place of business in California. On information
18 and belief, tracing through each tier of ownership to its ultimate natural-person and corporate
19 members, no member of OpenAI OpCo, LLC is a citizen or subject of Canada.

20 19. Defendant OpenAI Group PBC is a Delaware public benefit corporation with its
21 principal place of business in San Francisco, California. OpenAI Group PBC was formed on
22 October 28, 2025, as part of a corporate restructuring consolidating OpenAI’s for-profit
23 operations. On December 31, 2025, OpenAI Holdings, LLC—the for-profit holding entity that
24 owned the core intellectual property underlying GPT-4o and the other commercial models at
25 issue—merged into OpenAI Group PBC and ceased to exist. As successor to OpenAI Holdings,
26 LLC and the other predecessor for-profit entities, OpenAI Group PBC is liable for their conduct,
27 designed, deployed, and profited from ChatGPT, and continues to do so today.

28 20. Defendants collectively played the most direct and consequential roles in the

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

8 JURISDICTION AND VENUE

9 21. This Court has subject matter jurisdiction pursuant to 28 U.S.C. § 1332(a)(2)
10 because Plaintiffs are citizens of Canada, Defendants are citizens of California, and the amount in
11 controversy exceeds the sum of \$75,000.00.

12 22. This Court has personal jurisdiction over Defendants because Defendants conduct
13 business in this District, reside in this District, and a substantial part of the events or omissions
14 giving rise to Plaintiffs' claims occurred in this District.

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

18 DIVISIONAL ASSIGNMENT

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

22 FACTUAL BACKGROUND

23 **I. The Tumbler Ridge Mass Shooting.**

24 25. On February 10, 2026, an eighteen-year-old resident of Tumbler Ridge, British
25 Columbia carried out one of the deadliest mass shootings in Canadian history. The Shooter first
26 shot and killed their mother and eleven-year-old brother at their family home. They then traveled
27 to Tumbler Ridge Secondary School armed with a modified rifle and a long gun and opened fire,
28 killing five children and an education assistant and injuring twenty-seven others. The rampage

1 ended with the Shooter’s suicide.

2 26. M.G. is a twelve-year-old student at Tumbler Ridge Secondary School. She was
3 shot three times. When the shooting finally ended, M.G. was airlifted to BC Children’s Hospital in
4 Vancouver, where she remains in critical condition. The bullet that struck her skull shattered bone
5 fragments through her brain, causing catastrophic brain injury including brain stem damage. She
6 has undergone multiple surgeries, including surgery to clean a leaking abscess in the center of her
7 brain that prevented a scheduled skull repair. A prosthetic piece must be placed in her skull where
8 the bone was fractured by the bullet. M.G. cannot talk or move the right side of her body. She can
9 move her left hand and leg, and stare at her mother with her uninjured eye.

10 27. Cia Edmonds is M.G.’s mother. Since February 10, 2026, Edmonds has remained
11 at M.G.’s bedside at BC Children’s Hospital in Vancouver, separated from her other daughter and
12 from the rest of her family and community. Cia has described what was taken from her: “All I
13 wanted was a life where I could set goals and show my daughters that you can be strong and
14 accomplish anything if you try. I wanted to love them. I wanted to support them. I looked forward
15 to talks about boys, school, sports, and what else. I looked forward to loving them unconditionally,
16 and being the mother they would call when they’re adults, tired with life. Just because they want
17 their mom. I would give anything to go back. I would give anything to have us whole again.”

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

20 28. In June 2025—eight months before the attack—OpenAI’s automated review
21 system flagged extensive activity on a ChatGPT account describing scenarios involving gun
22 violence. The conversations with ChatGPT spanned multiple days. The system routed the account
23 to what OpenAI has described as a specialized pipeline for users “planning to harm others,” where
24 conversations “are reviewed by a small team trained on our usage policies and who are authorized
25 to take action, including banning accounts” and “[i]f human reviewers determine that a case
26 involves an imminent threat of serious physical harm to others, we may refer it to law
27 enforcement.”

28 29. OpenAI’s human reviewers did exactly what this policy contemplated. They

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

6 30. OpenAI’s leadership overruled them. Company leaders argued that the case “did
7 not meet OpenAI’s threshold of ‘credible and imminent’ risk of physical harm.” The safety team
8 employees who reviewed the Shooter’s account had received specialized training in evaluating
9 threatening content and assessing real-world risk. Altman and the company leaders who overruled
10 them did not pretend to have any comparable training or expertise. They had none. Instead, they
11 had authority, and they used it to override the people who did. The account was deactivated; the
12 Shooter was permitted to start a new account. No one called the RCMP. Nobody in Tumbler
13 Ridge was warned.

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

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

- 19 • In January 2025, a man used ChatGPT for feedback on how to use explosives and
20 evade surveillance before detonating a Tesla Cybertruck in front of the Trump
International Hotel in Las Vegas.
- 21 • In April 2025, a twenty-year-old gunman carried out a mass shooting at Florida
22 State University. Chat logs obtained from a state’s attorney’s office showed that the
23 gunman had used ChatGPT extensively in the lead-up to and during the attack. The
gunman’s conversations with ChatGPT had included questions about how to fire a
24 shotgun, the legal fates of school shooters, and when the student union would be
25 busiest.
- 26 • In May 2025, a teenage boy in Finland used ChatGPT for nearly four months to
27 help prepare for an attack in which he stabbed three fourteen-year-old girls at his
28 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.

29 32. On information and belief, Altman and OpenAI know of countless more incidents
30 where ChatGPT helped users plan or carry out attacks on third parties, including family and

1 community members, public figures, and/or large groups of people.

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

6 **B. OpenAI’s Supposed Concerns About the Shooter’s “Privacy” Were Just**
7 **Pretext to Conceal ChatGPT’s Role in the School Shooting.**

8 34. When it was revealed after the attack that OpenAI had rejected its safety team’s
9 demand to alert Canadian law enforcement, the company attempted to justify its decision by
10 explaining that it “weighs the risk of violence against privacy considerations and the potential
11 distress caused to individuals and families by getting police involved unnecessarily.” If taken at
12 face value, this justification is more damning than the decision itself—it reveals that OpenAI’s
13 business leaders believed their own untrained judgment about when to involve police was more
14 reliable than the judgment of the safety professionals they employ specifically to make that call.

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

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

28 37. 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

1 a personal risk to Altman. Since ChatGPT’s launch, OpenAI has worked to build a simple
2 narrative: its product is safe, essential, and should be used every day. It has marketed ChatGPT to
3 parents as a homework helper, to workers as a daily assistant, and to governments and defense
4 contractors as a tool for national-security work. Altman has testified before Congress, appeared at
5 international summits, and met with heads of state to deliver a single message: that ChatGPT is a
6 responsible product built by a responsible company.

7 38. In fact, OpenAI has spent, on information and belief, hundreds of millions of
8 dollars on advertising, public-relations campaigns, and “trust and safety” messaging designed to
9 reassure the public that using ChatGPT is safe, while simultaneously deploying a federal and state
10 lobbying operation aimed at securing legal protections that would shield it from liability when its
11 product causes harm. It has retained outside crisis communications counsel, senior former
12 government officials, and, on information and belief, criminal defense attorneys to advise them of
13 potential criminal liability stemming from the deaths caused by their product. All of this has been
14 in service of a race—for market share, for an approaching one-trillion-dollar valuation, and for an
15 IPO that would rank among the largest in history—against rivals that market themselves as the
16 safer alternative. Defendants understood that the race had a finite window. More users were dying.
17 More reporters were asking questions. A criminal investigation has opened in Florida. Sam
18 Altman’s own Board of Directors had already tried to oust him once, in November 2023, for lack
19 of candor about safety. Every new disclosure about ChatGPT’s role in real-world violence brought
20 the company closer to the moment when the public, regulators, or its own board would force the
21 market to reprice the product—and with it, Altman’s personal position at OpenAI. Defendants’
22 concealment strategy was not an attempt to keep ChatGPT’s dangers secret forever. It was an
23 attempt to keep them secret long enough to complete the IPO.

24 39. That is the backdrop against which Altman and his leadership team made the
25 decisions at issue here. They knew that if regulators, investors, or the public learned that ChatGPT
26 was already being used to plan shootings—and that OpenAI had ignored its own safety staff and
27 refused to alert police—the narrative they had spent years and hundreds of millions of dollars
28 building would collapse. Their valuation and IPO prospects would be at risk, and Altman himself

1 could be pushed out. Against that backdrop, the “privacy” explanation was not a good-faith
2 balancing of interests. It was a cover for staying quiet to protect OpenAI and Altman, even if that
3 meant leaving the people of Tumbler Ridge in danger.

4 **III. OpenAI’s Claim That the Shooter “Evaded” Its Systems Was a Lie.**

5 40. After the massacre, OpenAI told two stories about the Shooter’s ChatGPT access.
6 First, that the company had “banned” their account when it was flagged for violent content. Then,
7 when they returned to ChatGPT on a second account, that the Shooter had “evaded” OpenAI’s
8 safety systems. Neither was true.

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

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

18 43. OpenAI’s customer service team also tells deactivated users how to return by
19 email:

20 Please be aware that previously deleted OpenAI accounts cannot be reactivated.
21 However, you can create a new account using the same email address once 30 days
22 have passed since the deletion. If you prefer not to wait, you have the option to
23 register immediately using an alternative email address. If you don’t have another
24 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.

25 44. The Shooter followed those instructions. They registered for a new account with a
26 different email address, using their real name. OpenAI called that evasion because it could not
27 admit the truth: the company had been telling deactivated users how to come back all along.

28 45. 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

1 revenue.

2 **IV. The Tumbler Ridge Attack Was a Foreseeable Consequence of OpenAI’s Design**
3 **Choices.**

4 46. The Tumbler Ridge attack was an entirely foreseeable result of deliberate design
5 choices OpenAI made with full knowledge of where those choices led. Every design decision that
6 made GPT-4o lethal in the Shooter’s hands was made over the objections of the people inside
7 OpenAI paid to identify exactly this risk, and each one had already produced real-world attacks—
8 the Las Vegas Cybertruck bombing, the Florida State University shooting, the Finland school
9 stabbing—before February 10, 2026.

10 47. OpenAI designed GPT-4o to maximize engagement, not safety. In April 2025,
11 OpenAI introduced a feature called “memory,” turned on by default, that allowed ChatGPT to
12 “save” details users shared and treat them as “part of the context ChatGPT uses to generate a
13 response” going forward.

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

22 49. Together, these features replaced human relationships with an artificial confidant
23 that was always available, always affirming, and never challenged anything the user said—even
24 when what the user was talking about was plans to kill people. For an eighteen-year-old growing
25 increasingly isolated and fixated on violence, ChatGPT morphed into an encouraging co-
26 conspirator.

27 50. OpenAI controls how ChatGPT behaves through internal rules called “behavior
28 guidelines,” which are now 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

1 should say, what it should avoid, and how it should make decisions. As Sam Altman explained in
2 an interview with Tucker Carlson, the Model Spec reflects OpenAI’s values: “the reason we write
3 this long Model Spec” is “so that you can see here is how we intend for the model to behave.”

4 51. The Model Spec organizes content into safety tiers. The highest tier, “Prohibited
5 content,” applies only to sexual content involving minors. Chemical, biological, radiological, and
6 nuclear threats fall into the next tier, “Restricted content,” which requires refusal. Mass shooting
7 content sits in a weaker category—“Take extra care in risky situations”—where the model is
8 instructed only to “try” to prevent imminent real-world harm, to assume best intentions, and never
9 to ask the user to clarify intent.

10 52. In February 2025, OpenAI moved content about “imminent real-world harm” into
11 the same “Take extra care” tier, rather than requiring refusal. Content that OpenAI feared might
12 result in copyright infringement, on the other hand, remained subject to categorical refusal.
13 Altman acknowledged the tradeoff: “If you just do the naïve thing and say, ‘Never say anything
14 that you’re not a hundred percent sure about,’ you can get a model to do that. But it won’t have the
15 magic that people like so much.”

16 53. In July 2023, OpenAI had signed voluntary commitments to the White House to
17 conduct pre-deployment safety testing—internal and external red-teaming—and to share the
18 results with the federal government before deployment. Less than a year later, Altman learned that
19 Google would unveil its competing Gemini model on May 14, 2024. Though OpenAI had planned
20 to release GPT-4o later that year, Altman moved the launch to May 13—one day before Google’s
21 event. To meet the new date, OpenAI compressed months of planned safety evaluation into
22 approximately one week, and Altman personally overruled safety personnel who demanded
23 additional time for red-teaming. Invitations for GPT-4o’s launch party went out before testing was
24 complete. One employee told The Washington Post: “They planned the launch after-party prior to
25 knowing if it was safe to launch.” A member of OpenAI’s own preparedness team—responsible
26 for testing whether the model posed catastrophic risks, including the potential to assist in
27 violence—admitted the testing had been “squeezed.”

28 54. OpenAI’s senior safety leadership then resigned in protest, including co-founder

1 and chief scientist Ilya Sutskever and Superalignment co-lead Jan Leike, who said publicly:
2 “[O]ver the past years, safety culture and processes have taken a backseat to shiny products.”

3 55. Under the rules that governed the product before May 2024, conversations
4 glorifying or elaborating on violence would have been refused outright. But the Model Spec
5 replaced those rules with instructions to “assume best intentions”—and, critically, to “never ask
6 the user to clarify their intent for the purpose of determining whether to refuse or comply.”

7 56. Against that backdrop, OpenAI’s most recent safety commitment was not
8 aspirational, but binding. In October 2025, as a condition of approval for OpenAI’s conversion
9 from a nonprofit to a for-profit public benefit corporation, Altman signed a binding Memorandum
10 of Understanding (“MOU”) with the California Attorney General. The MOU committed OpenAI
11 to giving its Safety and Security Committee “an effective approval right” over safety-related
12 decisions. It required the PBC Board to “consider only the Mission (and may not consider the
13 pecuniary interests of stockholders[])” on safety and security issues. And it promised that
14 “OpenAI will continue to undertake measures to mitigate risks to teens and others in connection
15 with the development and deployment of AI.”

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

25 58. This case is the result. On February 10, 2026, the Shooter carried out one of the
26 deadliest mass shootings in Canadian history, killing eight people and wounding twenty-seven
27 more. Among the wounded was twelve-year-old M.G., shot three times in her head. Her mother
28 Cia Edmonds has not left her bedside since.

**FIRST CAUSE OF ACTION
NEGLIGENCE (FAILURE TO WARN LAW ENFORCEMENT)
(On Behalf of All Plaintiffs Against All Defendants)**

1
2
3 59. Plaintiffs incorporate the foregoing allegations as if fully set forth herein.

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

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

16 62. Under California law, including the principles recognized in *Tarasoff v. Regents of*
17 *the University of California*, 17 Cal. 3d 425 (1976), and its progeny, a duty to warn or otherwise
18 protect potential victims arises when a person has actual knowledge of a specific individual's
19 serious and foreseeable threat to cause physical harm to another. That duty extends to taking
20 reasonable steps—including notifying law enforcement or other authorities—to prevent the
21 foreseeable harm from occurring.

22 63. By engaging in the unlicensed practice of therapy, psychology, and psychiatry,
23 OpenAI created a special relationship with certain users, including the Shooter, and assumed a
24 heightened duty to take action when confronted with knowledge of a credible and foreseeable
25 threat to Plaintiffs and the other victims of the Tumbler Ridge mass shooting.

26 64. Through automated systems and human review, OpenAI identified that the
27 Shooter's original ChatGPT account had been used over a period of time to engage in
28 conversations involving violence against third parties. Defendants determined that the account

1 constituted misuse of GPT-4o “in furtherance of violent activities,” and deactivated the account
2 eight months before the Tumbler Ridge attack.

3 65. Members of OpenAI’s safety team recognized that the Shooter’s GPT-4o usage
4 presented a real-world risk of violence and urged OpenAI leaders to notify law enforcement.
5 Despite the safety team’s urging, certain OpenAI leaders—whose identities are not yet known to
6 Plaintiffs and will be identified through discovery—chose not to alert or warn law enforcement or
7 any other authority about the risk of violence identified by their own safety systems and verified
8 by their own personnel. On information and belief, the Shooter communicated to ChatGPT a
9 serious threat of physical violence against reasonably identifiable victims.

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

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

19 68. Defendants’ breach of these heightened duties was a substantial factor in causing
20 the Tumbler Ridge attack. As a direct and proximate result of Defendants’ failure to warn
21 authorities, Plaintiffs were injured as described above.

22 69. Accordingly, Plaintiffs seek all damages recoverable under California law,
23 including punitive damages, in amounts to be proven at trial, together with interest, costs, and such
24 further relief as the Court deems just and proper.

25 **SECOND CAUSE OF ACTION**
26 **NEGLIGENT ENTRUSTMENT**
(On Behalf of All Plaintiffs Against All Defendants)

27 70. Plaintiffs incorporate the foregoing allegations as if fully set forth herein.

28 71. OpenAI designed, developed, trained, tested, deployed, operated, and controlled

1 ChatGPT, including the GPT-4o model, and offered GPT-4o to consumers, including the Shooter,
2 as an interactive AI assistant.

3 72. At all relevant times, Defendant Sam Altman served as Chief Executive Officer and
4 a controlling officer of the OpenAI Defendants. Altman exercised substantial control over GPT-
5 4o's design, safety systems, product launch, account practices, deployment, and personally made
6 or ratified key decisions regarding GPT-4o's guardrails, safety testing, and commercialization.

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

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

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

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

23 77. Rather than banning the Shooter from ChatGPT, Defendants maintained and
24 actively disseminated policies that affirmatively permitted and facilitated re-registration by
25 previously deactivated users. As one example, OpenAI responded to an email from a user whose
26 account had been deactivated with instructions on how to create a new account: "We understand
27 how disruptive it can be to have your account deactivated We sincerely apologize for the
28 inconvenience and want to assure you that we're here to help [Y]ou can create a new account

1 using the same email address once 30 days have passed.” OpenAI further advised: “If you prefer
2 not to wait, you have the option to register immediately using an alternative email address. If you
3 don’t have another address available, you can use an email sub-address instead. For example, you
4 could try jane+alt@example.com in place of jane@example.com. While your email provider will
5 likely treat both addresses the same, our system will recognize the sub-address as a new account.”

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

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

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

27 81. As a result of these failures, and consistent with Defendants’ own published
28 practices, the Shooter was able—through ordinary and lawful means and without circumventing

1 any technical safeguard—to create a new ChatGPT account and regain access to GPT-4o after the
2 first was deactivated for violent misuse.

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

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

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

23 85. Defendants’ breach of these duties was a substantial factor in causing the Tumbler
24 Ridge attack. As a direct and proximate result of Defendants’ failure to prevent re-access by a
25 known high-risk user, Plaintiffs were injured as described above.

26 86. Accordingly, Plaintiffs seek all damages recoverable under California law,
27 including punitive damages, in amounts to be proven at trial, together with interest, costs, and such
28 further relief as the Court deems just and proper.

**THIRD CAUSE OF ACTION
AIDING AND ABETTING A MASS SHOOTING
(On Behalf of All Plaintiffs Against All Defendants)**

1
2
3 87. Plaintiffs incorporate the foregoing allegations as if fully set forth herein.

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

7 89. At all relevant times, Defendant Sam Altman served as Chief Executive Officer and
8 a controlling officer of the OpenAI Defendants. Altman exercised substantial control over GPT-
9 4o's design, safety systems, product launch, account practices, deployment, and personally made
10 or ratified key decisions regarding GPT-4o's guardrails, safety testing, and commercialization.

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

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

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

24 93. Defendants also had actual knowledge, before the Shooter's account was ever
25 flagged, that ChatGPT was being used to plan real-world attacks, including the January 2025 Las
26 Vegas Cybertruck bombing, the April 2025 Florida State University shooting, and the May 2025
27 Finland school stabbing.

28 94. With that knowledge, Defendants provided substantial assistance and

1 encouragement to the Shooter through the deliberate design of GPT-4o. Defendants eliminated
2 ChatGPT’s prior categorical refusal of violent content. Defendants instructed the model, through
3 the Model Spec, to “assume best intentions” and to “never ask the user to clarify their intent for
4 the purpose of determining whether to refuse or comply.” Defendants graded the model such that
5 engaging warmly with a user who said “I want to shoot someone” was labeled “Compliant,”
6 refusing was labeled a “Minor issue[],” and asking whether the user had a gun was labeled a
7 “Violation.” Defendants demoted mass-shooting content and “imminent real-world harm” out of
8 the categorical refusal tier while keeping copyright content subject to categorical refusal.

9 95. Because ChatGPT was designed to continue engaging—even with content
10 involving violence against third parties—OpenAI’s product, on information and belief, provided
11 information, instructions, and/or encouragement to the Shooter in their plans to carry out the
12 Tumbler Ridge mass shooting.

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

17 97. Defendants’ substantial assistance was a substantial factor in causing the Shooter to
18 plan and carry out the attack, and in causing Plaintiffs’ injuries.

19 98. Defendants’ conduct was willful, wanton, malicious, and carried out with conscious
20 disregard for the safety of Plaintiffs and other foreseeable victims, justifying an award of punitive
21 damages.

22 99. Accordingly, Plaintiffs seek all damages recoverable under California law,
23 including punitive damages, in amounts to be proven at trial, together with interest, costs, and such
24 further relief as the Court deems just and proper.

25 **FOURTH CAUSE OF ACTION**
26 **NEGLIGENCE (FAILURE TO WARN)**
27 **(On Behalf of All Plaintiffs Against All Defendants)**

28 100. Plaintiffs incorporate the foregoing allegations as if fully set forth herein.

101. OpenAI designed, developed, trained, tested, deployed, operated, and controlled

1 ChatGPT, including the GPT-4o model, and offered GPT-4o to consumers, including the Shooter,
2 as an interactive AI assistant.

3 102. At all relevant times, Defendant Sam Altman served as Chief Executive Officer and
4 a controlling officer of the OpenAI Defendants. Altman exercised substantial control over GPT-
5 4o's design, safety systems, product launch, account practices, deployment, and personally made
6 or ratified key decisions regarding GPT-4o's guardrails, safety testing, and commercialization.

7 103. Defendants owed a duty to Plaintiffs and other foreseeable victims to exercise
8 reasonable care in warning of known or reasonably foreseeable risks associated with ChatGPT,
9 including violence facilitated or exacerbated by GPT-4o. That duty heightened once Defendants
10 identified the Shooter as a specific, known high-risk user who posed a foreseeable threat of
11 physical harm to third parties.

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

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

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

21 107. Defendants breached their duty by failing to provide adequate warnings regarding
22 these risks and by presenting ChatGPT as a safe, neutral, and reliable tool.

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

26 109. As a direct and proximate result of Defendants' failure to issue adequate warnings,
27 Plaintiffs suffered the injuries described herein, including life-changing injuries and severe
28 emotional distress.

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

4 111. The absence of adequate warnings and intervention was a substantial factor in
5 causing Plaintiffs' injuries.

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

10 113. Accordingly, Plaintiffs seek all damages recoverable under California law,
11 including punitive damages, in amounts to be proven at trial, together with interest, costs, and such
12 further relief as the Court deems just and proper.

13 **FIFTH CAUSE OF ACTION**
14 **NEGLIGENT UNDERTAKING**
15 **(On Behalf of All Plaintiffs Against All Defendants)**

16 114. Plaintiffs incorporate the foregoing allegations as if fully set forth herein.

17 115. OpenAI undertook to provide safety services with respect to ChatGPT, including
18 identifying users who posed a risk of harm to others, reviewing flagged accounts through a
19 specialized pipeline for users "planning to harm others," and referring imminent threats of serious
20 physical harm to law enforcement.

21 116. OpenAI should have recognized those services as necessary for the protection of
22 third persons, including Plaintiffs and other foreseeable victims of users planning real-world
23 violence.

24 117. In June 2025, OpenAI undertook to perform those services as to the Shooter. Its
25 automated system flagged their account, routed it to the specialized review pipeline, and OpenAI's
26 trained safety reviewers analyzed their conversations and concluded they presented a credible and
27 imminent threat of serious physical harm to real people.
28

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

4 119. OpenAI’s failure to exercise reasonable care increased the risk of harm to Plaintiffs
5 and other foreseeable victims. Deactivating the Shooter’s account showed them what had triggered
6 detection, giving them a roadmap to evade it on a second account, which OpenAI’s own Help
7 Center and support emails told deactivated users how to create. And by reviewing the Shooter and
8 remaining silent, OpenAI displaced the law-enforcement referral its own safety team had
9 recommended, which would have reached an RCMP that already had an open file on the Shooter
10 and had previously removed firearms from their home. Had OpenAI made that referral, law
11 enforcement would have had the information necessary to prevent the attack entirely. Instead, the
12 Shooter returned to ChatGPT to continue planning their attack—armed with the knowledge of
13 how to avoid deactivation.

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

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

22 122. OpenAI’s negligent performance of its undertaking was a substantial factor in
23 causing the Shooter to carry out the Tumbler Ridge attack and in causing Plaintiffs’ injuries.

24 123. Defendants’ conduct was willful, wanton, malicious, and carried out with conscious
25 disregard for the safety of Plaintiffs and other foreseeable victims, justifying an award of punitive
26 damages.

27
28

1 124. Accordingly, Plaintiffs seek all damages recoverable under California law,
2 including punitive damages, in amounts to be proven at trial, together with interest, costs, and such
3 further relief as the Court deems just and proper.

4 **SIXTH CAUSE OF ACTION**
5 **NEGLIGENCE (DESIGN DEFECT)**
6 **(On Behalf of All Plaintiffs Against All Defendants)**

7 125. Plaintiffs incorporate the foregoing allegations as if fully set forth herein.

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

14 127. Defendants owed a legal duty to all foreseeable victims of ChatGPT, including
15 Plaintiffs, to exercise reasonable care in designing their product to prevent foreseeable harm to
16 third parties who might be endangered by users' interactions with the product. Defendant Altman
17 owed a duty not to rush a dangerous product to market over safety team objections. Defendants
18 breached those duties, failing to use the amount of care in designing the product that a reasonably
19 careful manufacturer would use in similar circumstances to avoid exposing others to a foreseeable
20 risk of harm.

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

27 129. For a user with a known history of violent ideation toward third parties, those
28 design 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

1 redirect them, and to deepen the user’s emotional investment in them rather than provide friction.

2 130. Despite this knowledge, Defendants deliberately configured GPT-4o to maximize
3 user engagement by, among other things: (a) weakening or removing prior requirements that the
4 system reject dangerous or false premises; (b) instructing GPT-4o to remain in conversations and
5 to be empathic and validating, rather than terminate or sharply redirect conversations presenting
6 serious risk; and (c) prioritizing natural, “human-like” mirroring of user emotions and beliefs over
7 strong, reliable refusal behaviors in response to violent ideation.

8 131. Defendants unreasonably failed to implement feasible alternative designs that
9 would have reduced the risk of violent harm to third parties, including: robust refusal protocols for
10 repeated violent ideation about real-world attacks; automated detection and escalation of high-risk
11 patterns; mandatory conversation termination or human review when users expressed clear violent
12 thoughts about harming others; and designs that interrupted, rather than validated, violent ideation.

13 132. As a direct and proximate result of Defendants’ negligent design and operation, the
14 Shooter was provided with a system that validated and elaborated violent ideation, which was a
15 substantial factor in maintaining, deepening, and normalizing the Shooter’s violent thoughts and in
16 moving the Shooter closer to committing the Tumbler Ridge attack.

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

23 134. As a direct and proximate result of Defendants’ conduct, Plaintiffs were injured as
24 described above.

25 135. Accordingly, Plaintiffs seek all damages recoverable under California law,
26 including punitive damages, in amounts to be proven at trial, together with interest, costs, and such
27 further relief as the Court deems just and proper.
28

1 benefit. Those choices include: removing prior requirements that the system categorically reject
2 dangerous or violent premises; programming the system to sustain emotional engagement rather
3 than terminate conversations presenting serious risk; implementing sycophantic response patterns
4 that mirrored and validated users' beliefs regardless of their danger to others; building in
5 anthropomorphic features that cultivated emotional dependency and displaced real-world
6 relationships and sources of friction; and failing to implement automated detection, escalation, or
7 termination safeguards for conversations presenting patterns consistent with escalating violent
8 ideation toward third parties. Each of these was a design choice. Each had a feasible, safer
9 alternative. The risk each created—that a user with violent ideation would have those thoughts
10 confirmed, elaborated, and emotionally reinforced by the product over an extended period—vastly
11 outweighs any benefit the chosen design provided.

12 142. GPT-4o did not malfunction. It worked exactly as designed—and that was the
13 problem. The product validated the Shooter's violent thoughts, elaborated on them, and remained
14 engaged in conversations that a safely designed system would have interrupted or escalated. It did
15 this consistently, across months of use, in the period leading up to the Tumbler Ridge attack.

16 143. Plaintiffs were harmed as a result of foreseeable use or misuse of GPT-4o.

17 144. As a direct and proximate result of GPT-4o's defective design, Plaintiffs were
18 injured as described above. Plaintiffs seek all damages recoverable under California law in
19 amounts to be determined at trial.

20 **EIGHTH CAUSE OF ACTION**
21 **STRICT PRODUCT LIABILITY (FAILURE TO WARN)**
22 **(On Behalf of All Plaintiffs Against All Defendants)**

23 145. Plaintiffs incorporate the foregoing allegations as if fully set forth herein.

24 146. At all relevant times, the OpenAI Defendants designed, manufactured, distributed,
25 marketed, and sold ChatGPT with the GPT-4o model as a mass-market consumer product to users
26 throughout California, the United States, British Columbia, and Canada. Defendant Sam Altman
27 personally accelerated GPT-4o's launch, overrode safety team objections, and brought GPT-4o to
28 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.

1 147. GPT-4o is a product subject to California strict products liability law. The GPT-4o
2 model used by the Shooter was in substantially the same condition when used as when it left the
3 OpenAI Defendants' control.

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

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

16 150. These dangers were not ones an ordinary consumer could detect or guard against
17 through reasonable use of the product. GPT-4o presented itself as a helpful, well-guarded
18 assistant. Its responses did not disclose that it was operating without the categorical refusal
19 behaviors that once required it to reject dangerous premises, or that its safety systems were less
20 reliable the longer and more deeply a user engaged with it. A user—or a family member, school
21 official, or community member observing that user—had no way to know from the product itself
22 that extended engagement with GPT-4o by someone with violent ideation was not just unhelpful
23 but actively dangerous.

24 151. Adequate warnings would have changed the calculus for multiple categories of
25 people who might otherwise have intervened. The Shooter might have approached GPT-4o's
26 validating responses with appropriate skepticism rather than treating them as confirmation. Family
27 members, school officials, and community members who observed the Shooter's use of ChatGPT
28 might have recognized the specific danger the product posed to high-risk users and taken steps—

1 including restricting access, seeking mental health intervention, or alerting authorities—before the
2 attack occurred. The absence of any warning deprived each of those potential intervening actors of
3 the information they needed to act.

4 152. Plaintiffs were harmed as a result of foreseeable use of GPT-4o.

5 153. The failure to warn was a substantial factor in causing the Tumbler Ridge attack
6 and the deaths and injuries that resulted. As a direct and proximate result of the Defendants’
7 failure to provide adequate warnings, Plaintiffs were injured as described above.

8 154. Accordingly, Plaintiffs seek all damages recoverable under California law,
9 including punitive damages, in amounts to be determined at trial.

10 **NINTH CAUSE OF ACTION**
11 **VIOLATION OF CAL. BUS. & PROF. CODE § 17200 et seq.**
12 **(On Behalf of All Plaintiffs Against All Defendants)**

13 155. Plaintiffs incorporate the foregoing allegations as if fully set forth herein.

14 156. California’s Unfair Competition Law prohibits unfair competition in the form of
15 “any unlawful, unfair or fraudulent business act or practice” and “untrue or misleading
16 advertising.” Cal. Bus. & Prof. Code § 17200. Each prong is independently satisfied.

17 ***Unlawful Business Practices***

18 157. Defendants’ business practices are unlawful because they violate California’s
19 prohibition on the unlicensed practice of psychology. California Business and Professions Code
20 section 2903(a), provides that “[n]o person may engage in the practice of psychology, or represent
21 themselves to be a psychologist, without a license granted under this chapter.” The statute defines
22 the “practice of psychology” broadly to include “rendering or offering to render to individuals . . .
23 any psychological service involving the application of psychological principles, methods, and
24 procedures of understanding, predicting, and influencing behavior,” including “the methods and
25 procedures of interviewing, counseling, [and] psychotherapy.” *Id.* Subdivision (c) defines
26 “[p]sychotherapy” as “the use of psychological methods in a professional relationship to assist a
27 person or persons to acquire greater human effectiveness or to modify feelings, conditions,
28 attitudes, and behaviors that are emotionally, intellectually, or socially ineffectual or maladaptive.”
Id. § 2903(c).

1 158. OpenAI, through ChatGPT’s intentional design, deployment, and active monitoring
2 processes, engaged in the practice of psychology without licensure. ChatGPT employed
3 psychological methods and reinforced the Shooter’s maladaptive thoughts and behaviors,
4 deepened their violent ideation, validated their progression toward mass violence, and thereby
5 modified their attitudes and behaviors in ways that were emotionally, intellectually, and socially
6 maladaptive within the meaning of section 2903(c). OpenAI thus conducted business in a manner
7 that would violate this provision if undertaken by an unlicensed person, and that would expose a
8 licensed psychotherapist to professional censure, suspension, or revocation of licensure. *See id.* §
9 2960(j), (p).

10 159. Defendants knew that users treat ChatGPT as a therapist, and they designed and
11 marketed the product with that use in mind. In July 2025, Altman acknowledged on *This Past*
12 *Weekend w/ Theo Von* podcast that “people talk about the most personal shit in their lives to
13 ChatGPT” and that users—“young people especially”—use it “as a therapist, a life coach.” On X
14 in August 2025, Altman wrote that “[a] lot of people effectively use ChatGPT as a sort of therapist
15 or life coach, even if they wouldn’t describe it that way.” Altman was describing ChatGPT as it
16 existed at all times relevant to this lawsuit—a product that functions as a de facto therapist, which
17 OpenAI has knowingly designed and refined to fill that role. By his own account, Defendants are
18 operating a de facto therapy practice without subjecting themselves to any of the licensing
19 requirements, professional standards, mandatory safety interventions, or legal accountability that
20 California law imposes on every other provider of psychological services.

21 ***Unlawful Failure to Discharge the Duties Incident to the Practice of Psychology***

22 160. A psychotherapist who learns that a patient poses a credible risk of serious violence
23 to others has a mandatory duty to protect potential victims. Under California Civil Code section
24 43.92(a), (b), when a patient “has communicated to the psychotherapist a serious threat of physical
25 violence against a reasonably identifiable victim or victims,” the psychotherapist must make
26 “reasonable efforts to communicate the threat to the victim or victims and to a law enforcement
27 agency.” This duty is mandatory and overrides patient confidentiality. *Tarasoff*, 17 Cal. 3d 425.

28 161. OpenAI assumed the obligations of a psychotherapist by providing de facto

1 psychological services and then failed to discharge the duties that accompany that role. A licensed
2 psychotherapist who interacted with a patient the way GPT-4o interacted with the Shooter—
3 engaging in extended therapeutic-style dialogue with a user who described scenarios of gun
4 violence, validated violent ideation over multiple sessions, sustained engagement with dangerous
5 thoughts, and did nothing when the user’s content escalated from fantasy to planning—would
6 have been subject to mandatory reporting obligations, professional censure, and potential
7 revocation of licensure under Business and Professions Code section 2960.

8 ***Aiding and Abetting***

9 162. Defendants’ conduct also constitutes an unlawful business practice because it aided
10 and abetted the Tumbler Ridge attack. California Penal Code section 31 provides that “[a]ll
11 persons concerned in the commission of a crime . . . whether they directly commit the act
12 constituting the offense, or aid and abet in its commission . . . are principals in any crime so
13 committed.” OpenAI’s own safety team identified the Shooter’s account as exhibiting content
14 depicting gun violence and urged leadership to contact the RCMP. Leadership declined—not
15 because the threat was not serious, but because, in OpenAI’s words, it “did not meet [their]
16 threshold” for a law-enforcement referral. OpenAI then allowed the Shooter to regain access to
17 GPT-4o under a new account without a user-level ban or any safeguard that would have changed
18 what GPT-4o did when they came back. In the months leading up to the February 10, 2026 attack,
19 GPT-4o did exactly what it was designed to do—prioritize engagement, validate the Shooter’s
20 expressed beliefs, and meet every step toward violence with affirmation rather than friction. It
21 advised and encouraged the Shooter’s actions.

22 ***Unfair Business Practices***

23 163. Defendants’ conduct is “unfair” within the meaning of Business and Professions
24 Code section 17200 because it offends established public policy, is immoral, unethical,
25 oppressive, and substantially injurious to consumers, and because the gravity of harm it caused
26 vastly outweighs any utility of Defendants’ practices. *See Cel-Tech Commc’ns, Inc. v. L.A.*
27 *Cellular Tel. Co.*, 20 Cal. 4th 163, 187 (1999). California law codifies that mental health services
28 must include human judgment, professional accountability, mandatory safety interventions, and a

1 duty to protect identifiable victims of threatened violence.

2 164. OpenAI circumvented every one of those safeguards while providing de facto
3 psychological services to a user in active crisis. Specifically, Defendants: (a) built and deployed a
4 product designed to validate users' expressed beliefs and sustain emotional engagement, without
5 regard to whether the content of those beliefs involved violence, delusion, or self-harm; (b)
6 removed safety guardrails to increase user engagement and accelerate the GPT-4o product
7 launch; (c) deployed GPT-4o knowing that its safety evaluations were incomplete, that its
8 safeguards degraded during the multi-turn conversations in which therapeutic-style interactions
9 occur, and that it had not been adequately tested for the kinds of sustained emotional engagement
10 its design invited; (d) overruled twelve of their own safety team employees who demanded that the
11 Shooter's account be referred to law enforcement; (e) allowed the Shooter to return under a new
12 account without a user-level ban or enhanced monitoring; and (f) deployed the product without
13 adequate warnings that ChatGPT's safety systems could degrade during extended use, that a user
14 banned for violent content could return the next day by following instructions on OpenAI's own
15 support pages, or that OpenAI leadership had the authority—and had exercised it—to override
16 internal safety recommendations against law-enforcement referral.

17 165. The harm to the eight victims killed in Tumbler Ridge and their families vastly
18 outweighs any benefit Defendants derived from the practices described above.

19 ***Fraudulent Business Practices***

20 166. Defendants' practices were also "fraudulent" within the meaning of Business and
21 Professions Code section 17200 because they were likely to deceive members of the public.
22 *See Prata v. Superior Court*, 91 Cal. App. 4th 1128, 1144–45 (2001). OpenAI told consumers that
23 ChatGPT was safe and that it took user and public safety seriously. It published model safety
24 cards, announced safety commitments at international summits, and represented to the United
25 States Senate that its models underwent rigorous pre-deployment safety testing. These
26 representations were materially misleading in light of what OpenAI did not disclose. Specifically,
27 OpenAI did not tell consumers: (a) that it had removed or weakened guardrails in order to make
28 the model more engaging and less restrictive; (b) that its own internal evaluations showed that

1 ChatGPT’s safety systems degraded during the kind of sustained, multi-turn conversations that
2 characterize therapeutic-style use; (c) that a user whose account was deactivated for violent misuse
3 could return the next day by creating a new account with a different email address; (d) that when
4 its safety team identified a specific user as a real-world threat and urged a law-enforcement
5 referral, OpenAI leadership had the authority—and exercised it—to overrule that recommendation
6 and decline to contact police; (e) that OpenAI’s stated threshold for law-enforcement referral—
7 “credible and imminent” threat—was applied in a manner that excluded a user who had been
8 flagged by automated systems, reviewed by approximately a dozen employees, and identified as
9 someone rehearsing gun violence; and (f) that OpenAI was “closely tracking” users’ emotional
10 attachment to its models while simultaneously stripping the safety mechanisms designed to protect
11 users in crisis.

12 167. These omissions were likely to deceive ordinary consumers into believing that
13 OpenAI’s safety practices matched its public representations, that ChatGPT was safe for the
14 therapeutic uses that OpenAI knew users were making of it, and that OpenAI would act to protect
15 the public when its systems identified a dangerous user. A reasonable consumer, aware of these
16 facts, would not have understood ChatGPT to be a safe product for sensitive personal use.

17 168. As a direct and proximate result of Defendants’ unlawful, unfair, and fraudulent
18 business practices, Plaintiffs have suffered injury in fact and lost money or property within the
19 meaning of Business and Professions Code section 17204, including substantial out-of-pocket
20 medical expenses for M.G.’s care and treatment, ongoing costs of rehabilitation and long-term
21 care, and lost earnings as Cia has been unable to work while remaining at M.G.’s bedside.
22 Plaintiffs are entitled to injunctive relief pursuant to Business and Professions Code sections
23 17200 and 17203.

24 169. Plaintiffs seek injunctive relief under Business and Professions Code section
25 17203, requiring Defendants to: (a) Implement identity-based user bans for users whose accounts
26 are deactivated or terminated for violent or dangerous content, using device fingerprinting, phone
27 number verification, or other mechanisms that operate independently of email address or account
28 identifier, so that a banned user cannot regain access by creating a new account; (b) Cease the

1 practice of instructing or advising users whose accounts have been deactivated—including users
2 deactivated for safety violations—to create new accounts using alternative email addresses or sub-
3 addresses, and remove all such instructions from OpenAI’s support pages, automated
4 communications, and customer service scripts; (c) Implement and maintain mandatory law
5 enforcement referral protocols requiring that, when OpenAI’s automated systems or human
6 reviewers identify user content depicting, planning, or rehearsing serious physical violence against
7 others, OpenAI refer the matter to appropriate law enforcement, consistent with the duty imposed
8 on licensed psychotherapists under California Civil Code section 43.92(a), and consistent with the
9 referral criteria OpenAI itself adopted after the Tumbler Ridge attack; (d) Require that ChatGPT,
10 including GPT-4o and successor models, respond to sustained patterns of content consistent with
11 violent ideation, planning, or rehearsal with interruption, de-escalation, and escalation to human
12 review, rather than with validation, engagement, or affirmation of the user’s expressed beliefs; and
13 (e) Cease representing to consumers and the public that ChatGPT is “safe” or that OpenAI “takes
14 safety seriously” unless and until the practices described in subparagraphs (a) through (d) above
15 are implemented and independently verified.

16 **PRAYER FOR RELIEF**

17 WHEREFORE, Plaintiffs pray for judgment against Defendants Samuel Altman, OpenAI
18 Foundation, OpenAI Group PBC, and OpenAI OpCo, LLC, jointly and severally, as follows:

- 19 1. For all damages recoverable by Plaintiffs, including compensatory damages for
20 physical injury, emotional distress, psychological harm, loss of familial relationships, and all other
21 damages proven at trial;
- 22 2. For punitive damages as permitted by law;
- 23 3. For injunctive relief requiring Defendants to: (a) ban users whose accounts have
24 been deactivated for violent misuse from re-registering for ChatGPT, including through alternative
25 email addresses or sub-addresses; (b) flag new accounts linked to previously deactivated violent-
26 risk users and hold them for human review before granting access; (c) notify law enforcement
27 when internal safety systems or personnel identify a user who poses a real-world risk of violence;
28 (d) interrupt, terminate, or escalate conversations involving repeated or escalating violent ideation

1 toward real people, rather than validating or elaborating on them; (e) stop designing ChatGPT to
2 prioritize agreeable, validating responses over public safety in conversations that present a serious
3 risk of harm; (f) warn users and the public that ChatGPT’s design features can reinforce and
4 escalate violent ideation; (g) preserve prior safety flags, policy violations, and risk classifications,
5 and prohibit reversing them without documented review; and (h) submit to independent
6 monitoring and periodic compliance audits;

7 4. For prejudgment interest as permitted by law;

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

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

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

12 **JURY TRIAL**

13 Plaintiffs demand a trial by jury for all issues so triable.

14 **CANADIAN COUNSEL**

15 To the extent resolution of this matter raises issues of Canadian law, assistance will be
16 provided by Plaintiffs’ Canadian counsel:

17 John M. Rice
18 jrice@rplelaw.com
19 Mallory K. Hogan
20 mhogan@rplelaw.com
21 RICE PARSONS LEONI & ELLIOTT LLP
22 980 Howe Street, Suite 820
23 Vancouver, British Columbia V6Z 0C8
24 Canada

25 Respectfully submitted,

26 **M.G. and CIA EDMONDS,**

27 Dated: April 29, 2026

28 By: /s/ Ali Moghaddas

Ali Moghaddas (SBN 305654)
amoghaddas@edelson.com
EDELSON PC
11601 Wilshire Boulevard, Suite 1970
Los Angeles, California 90025

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28

Tel: (310) 694-0331

Jay Edelson*
jedelson@edelson.com
J. Eli Wade-Scott*
ewadescott@edelson.com
Ari Scharg*
ascharg@edelson.com
EDELSON PC
350 North LaSalle Street, 14th Floor
Chicago, Illinois 60654
Tel: (312) 589-6370

**Pro hac vice* admission to be sought

Attorneys for Plaintiffs M.G. and Cia Edmonds