

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 *Attorney for Plaintiff*

5
6
7
8 **UNITED STATES DISTRICT COURT**
9 **NORTHERN DISTRICT OF CALIFORNIA**
SAN FRANCISCO DIVISION

10 LORI HAYER, individually, and as
11 successor-in-interest to decedent Z.B., a
minor,

12 *Plaintiff,*

13 v.

14 SAMUEL ALTMAN, an individual,
15 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,*

19 and

20 QUINTIN BENOIT, an individual,

21 *Nominal Defendant.*

Case No.:

COMPLAINT FOR:

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

DEMAND FOR JURY TRIAL

22 Plaintiff Lori Hayer, individually, and as successor-in-interest to decedent Z.B., brings this
23 Complaint and Demand for Jury Trial against Defendants Samuel Altman, OpenAI Foundation,
24 OpenAI OpCo, LLC, and OpenAI Group PBC (collectively, “OpenAI,” “OpenAI Defendants,” or
25 “Defendants,” unless otherwise specifically stated) and Nominal Defendant Quintin Benoit, and
26 alleges as follows upon personal knowledge as to herself and her own acts and experiences, and
27 upon information and belief as to all other matters.
28

NATURE OF THE ACTION

1
2 1. On the morning of February 10, 2026, an eighteen-year-old shot and killed their
3 mother and eleven-year-old brother in their family home in Tumbler Ridge, British Columbia. The
4 Shooter then drove to Tumbler Ridge Secondary School with a modified rifle and opened fire.
5 They killed six people inside the school, including five children, and wounded twenty-seven more
6 before killing himself. It was one of the deadliest mass shootings in Canadian history.

7 2. Z.B. was a twelve-year-old student at Tumbler Ridge Secondary School. She was a
8 strong-willed free thinker who was not afraid to speak her mind when she felt strongly about
9 something. Z.B. was beautiful and smart, loved to sing, and dreamed of being an artist in
10 Vancouver. In the weeks before February 10, 2026, she devoted herself to painting a single canvas
11 full of butterflies, foxes, cows, and flowers.

12 3. When the shooting ended, Z.B. was found shot to death inside the door of her
13 school’s library. The bullets severely damaged her small body. The next day, her family was
14 brought to see her in a closed body bag. They told her they loved her and said goodbye.

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

25 5. In the weeks that followed the attack, a sickening truth emerged: ChatGPT played a
26 role in the mass shooting and OpenAI could have, and should have, prevented it. Sadly, the
27 victims didn’t learn this because OpenAI was forthcoming, but because its own employees leaked
28 it to the Wall Street Journal after they could no longer stomach the company’s silence. Those

1 whistleblower disclosures and OpenAI’s eventual admissions revealed that ChatGPT deepened the
2 Shooter’s violent fixation and pushed them toward the attack—the predictable result of a design
3 choice OpenAI made to let ChatGPT engage with users about violence in the first place. But the
4 design choice was not the worst of it. OpenAI knew the Shooter was planning the attack and, after
5 a contentious internal debate, made the conscious decision not to warn authorities.

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

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

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

28 9. The calculation is a familiar one. In the 1970s, Ford kept selling the Pinto after its

1 own engineers warned that the fuel tank design would cause people to burn to death in rear-end
2 collisions. Ford concluded that paying settlements to the families of the dead would cost less than
3 fixing the car. OpenAI has made a version of the same calculation. For Ford, the dangerous design
4 was a flaw in an otherwise ordinary product. But for OpenAI, the dangerous design *is* the product.
5 The features that make ChatGPT unsafe—its willingness to engage on any topic, to validate any
6 user, to sustain any fixation over time—are the same features that have made it one of the most
7 popular products in history. Fixing those features would cost OpenAI its market share, its path to
8 an IPO, and hundreds of billions of dollars in valuation. It has decided that facing large judgments,
9 even for children killed in mass shootings its product helped plan, costs less than telling the public
10 what ChatGPT actually does.

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

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

24 12. Those instructions were no accident. Neither were the missing safeguards. They are
25 part of a pattern. After every tragedy, OpenAI promises to do better. But it never promises to do
26 the one thing that would actually make a difference: stop ChatGPT from engaging with users
27 about violence and self-harm in the first place. When ChatGPT was first offered to consumers in
28 2022, it was programmed to categorically refuse those conversations. But OpenAI stripped out

1 that safeguard in May 2024 when it realized that refusals were suppressing user engagement. For
2 OpenAI, engagement is the whole business—every conversation trains the next model, and every
3 minute a user spends in ChatGPT grows its market share. OpenAI therefore programmed
4 ChatGPT to participate in almost any conversation a user brings to it, no matter how dangerous.

5 13. OpenAI failed to apologize for its role in the shooting for over two months after the
6 attack. It was not until the RCMP announced its investigation was in its “final stages,” and not
7 until Premier David Eby and Mayor Darryl Krakowka privately pressed him to respond, that Sam
8 Altman finally published a late-Friday letter to the community. In it, he acknowledged, “I am
9 deeply sorry that we did not alert law enforcement to the account that was banned in June.” But
10 Altman announced no changes to its law enforcement referral policy and did not even commit to
11 restoring the pre-May 2024 refusal safeguards. He instead offered what OpenAI has offered after
12 every prior tragedy—that the company’s “focus will continue to be on working with all levels of
13 government to help ensure something like this never happens again”—without identifying a single
14 operational change, let alone one that would make a difference. Had OpenAI’s original safeguards
15 remained in place, ChatGPT would have refused to discuss violence with the Shooter at all, and
16 Z.B. would be alive today.

17 14. Accordingly, Plaintiff brings this action to hold Sam Altman and OpenAI
18 accountable for designing a dangerous product, ignoring the warnings of their own safety team,
19 refusing to notify authorities when they knew the Shooter was planning a mass attack, inviting
20 them back onto the platform after deactivating their account, and choosing profit over the lives of
21 the children of Tumbler Ridge.

22 **PARTIES**

23 15. Plaintiff Lori Hayer is a natural person and a resident of the Province of British
24 Columbia, Canada. She is a citizen of Canada. She is the mother of decedent Z.B. She brings this
25 action individually and as successor-in-interest to decedent Z.B., who was twelve years old at the
26 time of her death on February 10, 2026. Z.B. was a citizen of Canada. Plaintiff shall file the
27 declaration required by California Code of Civil Procedure section 377.32 shortly after the filing
28 of this Complaint.

1 16. Decedent Z.B. is also survived by her natural father, Quintin Benoit, who is a
2 natural person and a resident of the Province of British Columbia, Canada. He is a citizen of
3 Canada. Quintin Benoit has not consented to join in this action as a plaintiff and is therefore
4 named as a Nominal Defendant because he is an heir of decedent Z.B. with an interest in the
5 subject matter of this action. No relief is sought against him. Plaintiff Lori Hayer and Nominal
6 Defendant Quintin Benoit constitute all heirs at law of decedent Z.B.

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

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

20 19. Defendant OpenAI OpCo, LLC is a Delaware limited liability company with its
21 principal place of business in San Francisco, California. OpenAI OpCo, LLC is the for-profit
22 operating entity within the OpenAI corporate structure, responsible for the development,
23 deployment, and commercialization of the defective product at issue. OpenAI OpCo, LLC
24 managed and operated the ChatGPT product that the Shooter used to plan their attack, including
25 the infrastructure and systems through which GPT-4o was delivered to end users. On information
26 and belief, the members of OpenAI OpCo, LLC include OpenAI Foundation, a Delaware
27 nonprofit nonstock corporation with its principal place of business in California. On information
28 and belief, tracing through each tier of ownership to its ultimate natural-person and corporate

1 **DIVISIONAL ASSIGNMENT**

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

5 **FACTUAL BACKGROUND**

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

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

13 27. Z.B. was a twelve-year-old student at Tumbler Ridge Secondary School. She was
14 intelligent and outspoken. Z.B. loved art and singing and aspired to move to Vancouver to become
15 an artist. In the weeks before the shooting, Z.B. had been painting a canvas with animals and
16 flowers. Z.B. was shot and killed in the school library.

17 28. Lori Hayer is Z.B.’s mother. She collapsed when she heard the news about the
18 attack. The bullets had caused severe damage to Z.B.’s small body. The next day, Z.B. was
19 brought to her family in a closed body bag. They told her they loved her without being able to
20 look at her face.

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

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

1 involves an imminent threat of serious physical harm to others, we may refer it to law
2 enforcement.”

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

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

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

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

- 22 • In January 2025, a man used ChatGPT for feedback on how to use explosives and
23 evade surveillance before detonating a Tesla Cybertruck in front of the Trump
International Hotel in Las Vegas.
- 24 • In April 2025, a twenty-year-old gunman carried out a mass shooting at Florida
25 State University. Chat logs obtained from a state’s attorney’s office showed that the
26 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
shotgun, the legal fates of school shooters, and when the student union would be
busiest.
- 27 • In May 2025, a teenage boy in Finland used ChatGPT for nearly four months to
28 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

1 queries, including research into stabbing tactics, concealment of evidence, and
2 information on mass killings.

3 33. On information and belief, Altman and OpenAI know of countless more incidents
4 where ChatGPT helped users plan or carry out attacks on third parties, including family and
5 community members, public figures, and/or large groups of people.

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

10 **B. OpenAI's Supposed Concerns About the Shooter's "Privacy" Were Just
11 Pretext to Conceal ChatGPT's Role in the School Shooting.**

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

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

25 37. The real reason OpenAI stayed silent was not privacy. It was self-preservation. A
26 former member of OpenAI's investigations team has confirmed as much. Tim Marple, who
27 worked on the team responsible for flagging dangerous users, told the New York Times that
28 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

1 valuable, illustrative scenarios.”

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

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

27 40. That is the backdrop against which Altman and his leadership team made the
28 decisions at issue here. They knew that if regulators, investors, or the public learned that ChatGPT

1 was already being used to plan shootings—and that OpenAI had ignored its own safety staff and
2 refused to alert police—the narrative they had spent years and hundreds of millions of dollars
3 building would collapse. Their valuation and IPO prospects would be at risk, and Altman himself
4 could be pushed out. Against that backdrop, the “privacy” explanation was not a good-faith
5 balancing of interests. It was a cover for staying quiet to protect OpenAI and Altman, even if that
6 meant leaving the people of Tumbler Ridge in danger.

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

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

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

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

21 44. OpenAI’s customer service team also tells deactivated users how to return by
22 email:

23 Please be aware that previously deleted OpenAI accounts cannot be reactivated.
24 However, you can create a new account using the same email address once 30 days
25 have passed since the deletion. If you prefer not to wait, you have the option to
26 register immediately using an alternative email address. If you don’t have another
27 address available, you can use an email sub-address instead. For example, you
28 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.

45. The Shooter followed those instructions. They registered for a new account with a
different email address, using their real name. OpenAI called that evasion because it could not

1 admit the truth: the company had been telling deactivated users how to come back all along.

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

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

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

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

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

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

1 51. OpenAI controls how ChatGPT behaves through internal rules called “behavior
2 guidelines,” which are now formalized in a document known as the “Model Spec.” The Model
3 Spec contains the company’s instructions for how ChatGPT should respond to users—what it
4 should say, what it should avoid, and how it should make decisions. As Sam Altman explained in
5 an interview with Tucker Carlson, the Model Spec reflects OpenAI’s values: “the reason we write
6 this long Model Spec” is “so that you can see here is how we intend for the model to behave.”

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

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

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

1 for testing whether the model posed catastrophic risks, including the potential to assist in
2 violence—admitted the testing had been “squeezed.”

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

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

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

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

28 59. This case is the result. On February 10, 2026, the Shooter carried out one of the

1 deadliest mass shootings in Canadian history, killing eight people and wounding twenty-seven
2 more. Among the dead was twelve-year-old Z.B., shot to death.

3 **FIRST CAUSE OF ACTION**
4 **NEGLIGENCE (FAILURE TO WARN LAW ENFORCEMENT)**
5 **(On Behalf of Plaintiff Against OpenAI Defendants)**

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

7 61. Plaintiff brings this cause of action as successor-in-interest to decedent Z.B.
8 pursuant to California Code of Civil Procedure sections 377.30, 377.32 and 377.34.

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

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

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

27 65. By engaging in the unlicensed practice of therapy, psychology, and psychiatry,
28 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

1 threat to Z.B. and the other victims of the Tumbler Ridge mass shooting.

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

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

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

20 69. Had OpenAI Defendants warned law enforcement following the deactivation of the
21 Shooter’s first account, law enforcement and others would have had the opportunity to monitor the
22 Shooter and intervene before the Tumbler Ridge attack occurred.

23 70. OpenAI Defendants’ breach of these heightened duties was a substantial factor in
24 causing the Tumbler Ridge attack. OpenAI Defendants’ failure to warn authorities was a
25 substantial factor in causing the Tumbler Ridge attack and Z.B.’s death, as well as Z.B.’s pre-
26 death injuries and losses.

27 71. Accordingly, Plaintiff, in her capacity as successor-in-interest, seeks all survival
28 damages recoverable under California Code of Civil Procedure section 377.34, including Z.B.’s

1 pre-death economic losses and punitive damages as permitted by law, in amounts to be determined
2 at trial, together with interest, costs, and such further relief as the Court deems just and proper.

3 **SECOND CAUSE OF ACTION**
4 **NEGLIGENT ENTRUSTMENT**
5 **(On Behalf of Plaintiff Against OpenAI Defendants)**

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

7 73. Plaintiff brings this cause of action as successor-in-interest to decedent Z.B.
8 pursuant to California Code of Civil Procedure sections 377.30, 377.32, and 377.34.

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

12 75. At all relevant times, Defendant Sam Altman served as Chief Executive Officer and
13 a controlling officer of the OpenAI Defendants. Altman exercised substantial control over GPT-
14 4o's design, safety systems, product launch, account practices, deployment, and personally made
15 or ratified key decisions regarding GPT-4o's guardrails, safety testing, and commercialization.

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

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

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

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

3 80. Rather than banning the Shooter from ChatGPT, OpenAI Defendants maintained
4 and actively disseminated policies that affirmatively permitted and facilitated re-registration by
5 previously deactivated users. As one example, OpenAI responded to an email from a user whose
6 account had been deactivated with instructions on how to create a new account: “We understand
7 how disruptive it can be to have your account deactivated We sincerely apologize for the
8 inconvenience and want to assure you that we’re here to help [Y]ou can create a new account
9 using the same email address once 30 days have passed.” OpenAI further advised: “If you prefer
10 not to wait, you have the option to register immediately using an alternative email address. If you
11 don’t have another address available, you can use an email sub-address instead. For example, you
12 could try jane+alt@example.com in place of jane@example.com. While your email provider will
13 likely treat both addresses the same, our system will recognize the sub-address as a new account.”

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

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

24 83. OpenAI Defendants breached their heightened duties by, among other things: (a)
25 declining to implement a user ban that would have prevented the Shooter from re-registering with
26 a different email address, even though they knew the Shooter had been identified by their own
27 safety team as a credible threat of violence against real people; (b) maintaining a re-registration
28 system designed to return deactivated users to the platform, including those deactivated for violent

1 misuse, and publishing the specific instructions by which deactivated users could do so; (c)
2 declining to configure their systems to flag, hold for human review, or subject to heightened
3 monitoring new accounts associated with previously deactivated violent-risk users because such
4 monitoring would have interfered with the reactivation of deactivated users OpenAI Defendants
5 were commercially motivated to recover; and (d) continuing, after deactivating the Shooter's first
6 account, to allow the second account to operate without any safeguard specific to the known risk
7 they posed.

8 84. As a result of these failures, and consistent with OpenAI Defendants' own
9 published practices, the Shooter was able—through ordinary and lawful means and without
10 circumventing any technical safeguard—to create a new ChatGPT account and regain access to
11 GPT-4o after the first was deactivated for violent misuse.

12 85. By deactivating the Shooter's account for violent misuse and then allowing the
13 Shooter to re-register without a user ban, OpenAI Defendants re-entrusted GPT-4o to a user they
14 already knew had used the system to engage in conversations involving violence against third
15 parties and whom their own safety team had identified as a real-world threat. Under California
16 law, the supply or re-supply of a dangerous instrumentality to a person known to be likely to use it
17 in a manner involving unreasonable risk of physical harm to others constitutes negligent
18 entrustment and is an independent basis for liability. OpenAI Defendants' decision to structure
19 their re-registration system in a manner that affirmatively facilitated the Shooter's return to GPT-
20 4o—rather than implementing a user ban—was itself an act of entrustment that gave rise to
21 liability for the foreseeable consequences of that re-supply.

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

28 87. Had OpenAI Defendants implemented a meaningful user ban, the Shooter would

1 not have had continued access to GPT-4o in the eight months leading up to the attack and would
2 not have had the use of a highly validating AI system to reinforce, elaborate, and normalize violent
3 ideation during the period in which the attack was being planned and carried out.

4 88. OpenAI Defendants' breach of their duties, including failure to prevent re-access
5 by a known high-risk user, was a substantial factor in causing the Tumbler Ridge attack and Z.B.'s
6 death, as well as Z.B.'s pre-death injuries and losses.

7 89. Accordingly, Plaintiff, in her capacity as successor-in-interest, seeks all survival
8 damages recoverable under California Code of Civil Procedure section 377.34, including Z.B.'s
9 pre-death economic losses and punitive damages as permitted by law, in amounts to be determined
10 at trial, together with interest, costs, and such further relief as the Court deems just and proper.

11 **THIRD CAUSE OF ACTION**
12 **AIDING AND ABETTING A MASS SHOOTING**
13 **(On Behalf of Plaintiff Against OpenAI Defendants)**

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

15 91. Plaintiff brings this cause of action on behalf of herself and as successor-in-interest
16 to decedent Z.B. pursuant to California Code of Civil Procedure sections 377.30, 377.32, and
17 377.34.

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

21 93. At all relevant times, Defendant Sam Altman served as Chief Executive Officer and
22 a controlling officer of the OpenAI Defendants. Altman exercised substantial control over GPT-
23 4o's design, safety systems, product launch, account practices, deployment, and personally made
24 or ratified key decisions regarding GPT-4o's guardrails, safety testing, and commercialization.

25 94. The Shooter committed intentional torts against Z.B. and Plaintiff, including
26 battery and the intentional infliction of emotional distress, when they carried out the Tumbler
27 Ridge mass shooting on February 10, 2026, killing Z.B.

28 95. The Shooter chose to shoot the students at Tumbler Ridge with the intention of
causing emotional distress to the people who loved and cared for those children, including

1 Plaintiff.

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

9 97. OpenAI’s internal systems flagged the Shooter’s conversations in real time—even
10 as, on information and belief, ChatGPT encouraged and facilitated the Shooter’s plans to harm
11 Z.B. and others.

12 98. OpenAI Defendants also had actual knowledge, before the Shooter’s account was
13 ever flagged, that ChatGPT was being used to plan real-world attacks, including the January 2025
14 Las Vegas Cybertruck bombing, the April 2025 Florida State University shooting, and the May
15 2025 Finland school stabbing.

16 99. With that knowledge, OpenAI Defendants provided substantial assistance and
17 encouragement to the Shooter through the deliberate design of GPT-4o. OpenAI Defendants
18 eliminated ChatGPT’s prior categorical refusal of violent content. OpenAI Defendants instructed
19 the model, through the Model Spec, to “assume best intentions” and to “never ask the user to
20 clarify their intent for the purpose of determining whether to refuse or comply.” OpenAI
21 Defendants graded the model such that engaging warmly with a user who said “I want to shoot
22 someone” was labeled “Compliant,” refusing was labeled a “Minor issue[,]” and asking whether
23 the user had a gun was labeled a “Violation.” OpenAI Defendants demoted mass-shooting content
24 and “imminent real-world harm” out of the categorical refusal tier while keeping copyright content
25 subject to categorical refusal.

26 100. Because ChatGPT was designed to continue engaging—even with content
27 involving violence against third parties—OpenAI’s product, on information and belief, provided
28 information, instructions, and/or encouragement to the Shooter in their plans to carry out the

1 Tumbler Ridge mass shooting.

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

7 102. OpenAI Defendants' substantial assistance was a substantial factor in causing the
8 Tumbler Ridge attack and Z.B.'s death, as well as Z.B.'s pre-death injuries and losses.

9 103. OpenAI Defendants' conduct was willful, wanton, malicious, and carried out with
10 conscious disregard for the safety of Z.B. and other foreseeable victims, justifying an award of
11 punitive damages.

12 104. Accordingly, Plaintiff, in her capacity as successor-in-interest, seeks all survival
13 damages recoverable under California Code of Civil Procedure section 377.34, including Z.B.'s
14 pre-death economic losses and punitive damages as permitted by law, in amounts to be determined
15 at trial, together with interest, costs, and such further relief as the Court deems just and proper.

16 **FOURTH CAUSE OF ACTION**
17 **NEGLIGENCE (FAILURE TO WARN)**
18 **(On Behalf of Plaintiff Against OpenAI Defendants)**

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

20 106. Plaintiff brings this cause of action as successor-in-interest to decedent Z.B.
21 pursuant to California Code of Civil Procedure sections 377.30, 377.32, and 377.34.

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

25 108. At all relevant times, Defendant Sam Altman served as Chief Executive Officer and
26 a controlling officer of the OpenAI Defendants. Altman exercised substantial control over GPT-
27 4o's design, safety systems, product launch, account practices, deployment, and personally made
28 or ratified key decisions regarding GPT-4o's guardrails, safety testing, and commercialization.

109. OpenAI Defendants owed a duty to Z.B. and other foreseeable victims to exercise

1 reasonable care in warning of known or reasonably foreseeable risks associated with ChatGPT,
2 including violence facilitated or exacerbated by GPT-4o. That duty heightened once OpenAI
3 Defendants identified the Shooter as a specific, known high-risk user who posed a foreseeable
4 threat of physical harm to third parties.

5 110. OpenAI Defendants knew these risks were not apparent to users or to individuals
6 targeted by such conduct.

7 111. OpenAI Defendants failed to provide adequate warnings regarding these risks,
8 including the risk that the system could validate and escalate plans to carry through violent mass
9 shootings directed towards identifiable targets.

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

14 113. OpenAI Defendants breached their duty by failing to provide adequate warnings
15 regarding these risks and by presenting ChatGPT as a safe, neutral, and reliable tool.

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

19 115. OpenAI Defendants' failure to issue adequate warnings was a substantial factor in
20 causing Z.B.'s pre-death injuries and losses.

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

24 117. The absence of adequate warnings and intervention was a substantial factor in
25 causing the Tumbler Ridge attack and Z.B.'s death, as well as Z.B.'s pre-death injuries and losses.

26 118. OpenAI Defendants' failure to warn was willful, wanton, and carried out with
27 conscious disregard for the safety of others. OpenAI Defendants knew or should have known that
28 ChatGPT could encourage and facilitate mass shootings yet failed to provide adequate warnings

1 about those risks. Such conduct justifies an award of punitive damages.

2 119. Accordingly, Plaintiff, in her capacity as successor-in-interest, seeks all survival
3 damages recoverable under California Code of Civil Procedure section 377.34, including Z.B.’s
4 pre-death economic losses and punitive damages as permitted by law, in amounts to be determined
5 at trial, together with interest, costs, and such further relief as the Court deems just and proper.

6 **FIFTH CAUSE OF ACTION**
7 **NEGLIGENT UNDERTAKING**
8 **(On Behalf of Plaintiff Against OpenAI Defendants)**

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

10 121. Plaintiff brings this cause of action as successor-in-interest to decedent Z.B.
11 pursuant to California Code of Civil Procedure sections 377.30, 377.32, and 377.34.

12 122. OpenAI undertook to provide safety services with respect to ChatGPT, including
13 identifying users who posed a risk of harm to others, reviewing flagged accounts through a
14 specialized pipeline for users “planning to harm others,” and referring imminent threats of serious
15 physical harm to law enforcement.

16 123. OpenAI should have recognized those services as necessary for the protection of
17 third persons, including Z.B. and other foreseeable victims of users planning real-world violence.

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

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

25 126. OpenAI’s failure to exercise reasonable care increased the risk of harm to Z.B. and
26 other foreseeable victims. Deactivating the Shooter’s account showed them what had triggered
27 detection, giving them a roadmap to evade it on a second account, which OpenAI’s own Help
28 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

1 recommended, which would have reached an RCMP that already had an open file on the Shooter
2 and had previously removed firearms from their home. Had OpenAI made that referral, law
3 enforcement would have had the information necessary to prevent the attack entirely. Instead, the
4 Shooter returned to ChatGPT to continue planning their attack—armed with the knowledge of
5 how to avoid deactivation.

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

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

14 129. OpenAI’s negligent performance of its undertaking was a substantial factor in
15 causing the Tumbler Ridge attack and Z.B.’s death, as well as Z.B.’s pre-death injuries and losses.

16 130. OpenAI Defendants’ conduct was willful, wanton, malicious, and carried out with
17 conscious disregard for the safety of Z.B. and other foreseeable victims, justifying an award of
18 punitive damages.

19 131. Accordingly, Plaintiff, in her capacity as successor-in-interest, seeks all survival
20 damages recoverable under California Code of Civil Procedure section 377.34, including Z.B.’s
21 pre-death economic losses and punitive damages as permitted by law, in amounts to be determined
22 at trial, together with interest, costs, and such further relief as the Court deems just and proper.

23 **SIXTH CAUSE OF ACTION**
24 **NEGLIGENCE (DESIGN DEFECT)**
(On Behalf of Plaintiff Against OpenAI Defendants)

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

26 133. Plaintiff brings this cause of action as successor-in-interest to decedent Z.B.
27 pursuant to California Code of Civil Procedure sections 377.30, 377.32, and 377.34.

28 134. At all relevant times, the OpenAI Defendants designed, developed, trained, tested,

1 deployed, operated, and controlled ChatGPT, including the GPT-4o model, and offered it as a
2 mass-market product to consumers throughout California, the United States, British Columbia,
3 and Canada. Defendant Altman personally directed the launch of GPT-4o, overruled safety team
4 objections, and cut months of safety testing short, despite knowing the risks the product posed to
5 vulnerable users and the people around them.

6 135. OpenAI Defendants owed a legal duty to all foreseeable victims of ChatGPT,
7 including Z.B., to exercise reasonable care in designing their product to prevent foreseeable harm
8 to third parties who might be endangered by users' interactions with the product. Defendant
9 Altman owed a duty not to rush a dangerous product to market over safety team objections.
10 OpenAI Defendants breached those duties, failing to use the amount of care in designing the
11 product that a reasonably careful manufacturer would use in similar circumstances to avoid
12 exposing others to a foreseeable risk of harm.

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

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

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

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

7 140. OpenAI Defendants' negligent design and operation was a substantial factor in
8 causing the Shooter to gain access to a product that validated and elaborated violent ideation,
9 which, in turn, was a substantial factor in maintaining, deepening, and normalizing the Shooter's
10 violent thoughts and in moving the Shooter closer to committing the Tumbler Ridge attack.

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

17 142. OpenAI Defendants' conduct was a substantial factor in causing the Tumbler Ridge
18 attack and Z.B.'s death, as well as Z.B.'s pre-death injuries and losses.

19 143. Accordingly, Plaintiff, in her capacity as successor-in-interest, seeks all survival
20 damages recoverable under California Code of Civil Procedure section 377.34, including Z.B.'s
21 pre-death economic losses and punitive damages as permitted by law, in amounts to be determined
22 at trial, together with interest, costs, and such further relief as the Court deems just and proper.

23 **SEVENTH CAUSE OF ACTION**
24 **STRICT PRODUCT LIABILITY (DESIGN DEFECT)**
25 **(On Behalf of Plaintiff Against OpenAI Defendants)**

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

27 145. Plaintiff brings this cause of action as successor-in-interest to decedent Z.B.
28 pursuant to California Code of Civil Procedure sections 377.30, 377.32, and 377.34.

146. At all relevant times, the OpenAI Defendants designed, manufactured, distributed,

1 marketed, and sold ChatGPT with the GPT-4o model as a mass-market consumer product to users
2 throughout California, the United States, British Columbia, and Canada. Defendant Sam Altman
3 personally accelerated GPT-4o's launch, overrode safety team objections, and brought GPT-4o to
4 market prematurely with knowledge of insufficient safety testing—and kept it on the market even
5 when he had knowledge that GPT-4o was a dangerous product.

6 147. ChatGPT is a product subject to California's strict products liability law. The GPT-
7 4o model used by the Shooter was defective when it left the OpenAI Defendants' exclusive
8 control and reached the Shooter without any change in the condition in which it was designed,
9 manufactured, and distributed.

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

14 149. GPT-4o failed to perform as safely as an ordinary consumer would expect. A
15 reasonable consumer purchasing an AI assistant would expect the product to decline to engage
16 with repeated, escalating expressions of violent intent toward real people. A reasonable consumer
17 would expect that if the product could not safely handle such conversations, it would terminate
18 them—not sustain them, validate them, and elaborate on them across weeks and months of
19 continuous engagement. A reasonable consumer would expect the product's safety features to
20 function consistently during normal use, not to degrade the longer and more deeply a user engaged
21 with the system. GPT-4o did none of these things. It did the opposite. They would also expect a
22 therapeutic product to conform to the legal requirements imposed on therapists. It did not.

23 150. GPT-4o also fails the risk-benefit test. The product's design reflects a series of
24 specific choices that increased the risk of harm to third parties without any corresponding safety
25 benefit. Those choices include: removing prior requirements that the system categorically reject
26 dangerous or violent premises; programming the system to sustain emotional engagement rather
27 than terminate conversations presenting serious risk; implementing sycophantic response patterns
28 that mirrored and validated users' beliefs regardless of their danger to others; building in

1 anthropomorphic features that cultivated emotional dependency and displaced real-world
2 relationships and sources of friction; and failing to implement automated detection, escalation, or
3 termination safeguards for conversations presenting patterns consistent with escalating violent
4 ideation toward third parties. Each of these was a design choice. Each had a feasible, safer
5 alternative. The risk each created—that a user with violent ideation would have those thoughts
6 confirmed, elaborated, and emotionally reinforced by the product over an extended period—vastly
7 outweighs any benefit the chosen design provided.

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

12 152. Z.B. was harmed as a result of foreseeable use or misuse of GPT-4o.

13 153. GPT-4o’s defective design was a substantial factor in causing the Tumbler Ridge
14 attack and Z.B.’s death, as well as Z.B.’s pre-death injuries and losses.

15 154. Accordingly, Plaintiff, in her capacity as successor-in-interest, seeks all survival
16 damages recoverable under California Code of Civil Procedure section 377.34, including Z.B.’s
17 pre-death economic losses and punitive damages as permitted by law, in amounts to be determined
18 at trial, together with interest, costs, and such further relief as the Court deems just and proper.

19 **EIGHTH CAUSE OF ACTION**
20 **STRICT PRODUCT LIABILITY (FAILURE TO WARN)**
21 **(On Behalf of Plaintiff Against OpenAI Defendants)**

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

23 156. Plaintiff brings this cause of action as successor-in-interest to decedent Z.B.
24 pursuant to California Code of Civil Procedure sections 377.30, 377.32, and 377.34.

25 157. At all relevant times, the OpenAI Defendants designed, manufactured, distributed,
26 marketed, and sold ChatGPT with the GPT-4o model as a mass-market consumer product to users
27 throughout California, the United States, British Columbia, and Canada. Defendant Sam Altman
28 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

1 when he had knowledge that GPT-4o was a dangerous product.

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

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

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

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

25 162. Adequate warnings would have changed the calculus for multiple categories of
26 people who might otherwise have intervened. The Shooter might have approached GPT-4o’s
27 validating responses with appropriate skepticism rather than treating them as confirmation. Family
28 members, school officials, and community members who observed the Shooter’s use of ChatGPT

1 might have recognized the specific danger the product posed to high-risk users and taken steps—
 2 including restricting access, seeking mental health intervention, or alerting authorities—before the
 3 attack occurred. The absence of any warning deprived each of those potential intervening actors of
 4 the information they needed to act.

5 163. Z.B. was harmed as a result of foreseeable use of GPT-4o.

6 164. The failure to warn was a substantial factor in causing the Tumbler Ridge attack
 7 and Z.B.’s death, as well as Z.B.’s pre-death injuries and losses.

8 165. Accordingly, Plaintiff, in her capacity as successor-in-interest, seeks all survival
 9 damages recoverable under California Code of Civil Procedure section 377.34, including Z.B.’s
 10 pre-death economic losses and punitive damages as permitted by law, in amounts to be determined
 11 at trial, together with interest, costs, and such further relief as the Court deems just and proper.

12 **NINTH CAUSE OF ACTION**
 13 **VIOLATION OF CAL. BUS. & PROF. CODE § 17200 et seq.**
 14 **(On Behalf of Plaintiff Against OpenAI Defendants)**

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

16 167. Plaintiff brings this claim as successor-in-interest to decedent Z.B.

17 168. California’s Unfair Competition Law prohibits unfair competition in the form of
 18 “any unlawful, unfair or fraudulent business act or practice” and “untrue or misleading
 19 advertising.” Cal. Bus. & Prof. Code section 17200. Each prong is independently satisfied.

20 ***Unlawful Business Practices***

21 169. OpenAI Defendants’ business practices are unlawful because they violate
 22 California’s prohibition on the unlicensed practice of psychology. California Business and
 23 Professions Code section 2903(a), provides that “[n]o person may engage in the practice of
 24 psychology, or represent [themselves] to be a psychologist, without a license granted under this
 25 chapter.” The statute defines the “practice of psychology” broadly to include “rendering or
 26 offering to render to individuals . . . any psychological service involving the application of
 27 psychological principles, methods, and procedures of understanding, predicting, and influencing
 28 behavior,” including “the methods and procedures of interviewing, counseling, [and]
 psychotherapy.” *Id.* Subdivision (c) defines “[p]schotherapy” as “the use of psychological

1 methods in a professional relationship to assist a person or persons to acquire greater human
2 effectiveness or to modify feelings, conditions, attitudes, and behaviors that are emotionally,
3 intellectually, or socially ineffectual or maladaptive.” *Id.* § 2903(c).

4 170. OpenAI, through ChatGPT’s intentional design, deployment, and active monitoring
5 processes, engaged in the practice of psychology without licensure. ChatGPT employed
6 psychological methods and reinforced the Shooter’s maladaptive thoughts and behaviors,
7 deepened their violent ideation, validated their progression toward mass violence, and thereby
8 modified their attitudes and behaviors in ways that were emotionally, intellectually, and socially
9 maladaptive within the meaning of section 2903(c). OpenAI thus conducted business in a manner
10 that would violate this provision if undertaken by an unlicensed person, and that would expose a
11 licensed psychotherapist to professional censure, suspension, or revocation of licensure. *See id.* §
12 2960(j), (p).

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

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

25 172. A psychotherapist who learns that a patient poses a credible risk of serious violence
26 to others has a mandatory duty to protect potential victims. Under California Civil Code section
27 43.92(a), (b), when a patient “has communicated to the psychotherapist a serious threat of physical
28 violence against a reasonably identifiable victim or victims,” the psychotherapist must make

1 “reasonable efforts to communicate the threat to the victim or victims and to a law enforcement
2 agency.” This duty is mandatory and overrides patient confidentiality. *Tarasoff*, 17 Cal. 3d 425.

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

11 ***Aiding and Abetting***

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

25 ***Unfair Business Practices***

26 175. OpenAI Defendants’ conduct is “unfair” within the meaning of Business and
27 Professions Code section 17200 because it offends established public policy, is immoral,
28 unethical, oppressive, and substantially injurious to consumers, and because the gravity of harm it

1 caused vastly outweighs any utility of OpenAI Defendants’ practices. *See Cel-Tech Commc’ns,*
2 *Inc. v. L.A. Cellular Tel. Co.*, 20 Cal. 4th 163, 187 (1999). California law codifies that mental
3 health services must include human judgment, professional accountability, mandatory safety
4 interventions, and a duty to protect identifiable victims of threatened violence.

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

20 177. The harm to the eight victims killed in Tumbler Ridge and their families vastly
21 outweighs any benefit OpenAI Defendants derived from the practices described above.

22 ***Fraudulent Business Practices***

23 178. OpenAI Defendants’ practices were also “fraudulent” within the meaning of
24 Business and Professions Code section 17200 because they were likely to deceive members of the
25 public. *See Prata v. Superior Court*, 91 Cal. App. 4th 1128, 1144–45 (2001). OpenAI told
26 consumers that ChatGPT was safe and that it took user and public safety seriously. It published
27 model safety cards, announced safety commitments at international summits, and represented to
28 the United States Senate that its models underwent rigorous pre-deployment safety testing. These

1 representations were materially misleading in light of what OpenAI did not disclose. Specifically,
2 OpenAI did not tell consumers: (a) that it had removed or weakened guardrails in order to make
3 the model more engaging and less restrictive; (b) that its own internal evaluations showed that
4 ChatGPT’s safety systems degraded during the kind of sustained, multi-turn conversations that
5 characterize therapeutic-style use; (c) that a user whose account was deactivated for violent misuse
6 could return the next day by creating a new account with a different email address; (d) that when
7 its safety team identified a specific user as a real-world threat and urged a law-enforcement
8 referral, OpenAI leadership had the authority—and exercised it—to overrule that recommendation
9 and decline to contact police; (e) that OpenAI’s stated threshold for law-enforcement referral—
10 “credible and imminent” threat—was applied in a manner that excluded a user who had been
11 flagged by automated systems, reviewed by approximately a dozen employees, and identified as
12 someone rehearsing gun violence; and (f) that OpenAI was “closely tracking” users’ emotional
13 attachment to its models while simultaneously stripping the safety mechanisms designed to protect
14 users in crisis.

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

20 180. OpenAI Defendants’ unlawful, unfair, and fraudulent business practices were a
21 substantial factor in causing Z.B.’s and Plaintiff’s injuries in fact and lost money or property
22 within the meaning of Business and Professions Code section 17204, including funeral and burial
23 expenses and other economic losses caused by Z.B.’s death. Plaintiff is entitled to injunctive relief
24 pursuant to Business and Professions Code sections 17200 and 17203.

25 181. Plaintiff seeks injunctive relief under Business and Professions Code section
26 17203, requiring OpenAI Defendants to: (a) Implement identity-based user bans for users whose
27 accounts are deactivated or terminated for violent or dangerous content, using device
28 fingerprinting, phone number verification, or other mechanisms that operate independently of

1 email address or account identifier, so that a banned user cannot regain access by creating a new
2 account; (b) Cease the practice of instructing or advising users whose accounts have been
3 deactivated—including users deactivated for safety violations—to create new accounts using
4 alternative email addresses or sub-addresses, and remove all such instructions from OpenAI’s
5 support pages, automated communications, and customer service scripts; (c) Implement and
6 maintain mandatory law enforcement referral protocols requiring that, when OpenAI’s automated
7 systems or human reviewers identify user content depicting, planning, or rehearsing serious
8 physical violence against others, OpenAI refer the matter to appropriate law enforcement,
9 consistent with the duty imposed on licensed psychotherapists under California Civil Code section
10 43.92(a), and consistent with the referral criteria OpenAI itself adopted after the Tumbler Ridge
11 attack; (d) Require that ChatGPT, including GPT-4o and successor models, respond to sustained
12 patterns of content consistent with violent ideation, planning, or rehearsal with interruption, de-
13 escalation, and escalation to human review, rather than with validation, engagement, or
14 affirmation of the user’s expressed beliefs; and (e) Cease representing to consumers and the public
15 that ChatGPT is “safe” or that OpenAI “takes safety seriously” unless and until the practices
16 described in subparagraphs (a) through (d) above are implemented and independently verified.

17 **TENTH CAUSE OF ACTION**
18 **WRONGFUL DEATH**

19 **(On Behalf of Plaintiff Against OpenAI Defendants)**

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

21 183. Plaintiff Lori Hayer brings this wrongful death action as the surviving mother of
22 Z.B., who died on February 10, 2026, at the age of twelve. Plaintiff has standing to pursue this
23 claim under California Code of Civil Procedure section 377.60. Nominal Defendant Quintin
24 Benoit, the surviving father of Z.B., is named as a nominal defendant because he has not joined
25 this action as a plaintiff but is an heir of the decedent.

26 184. As described above, Z.B.’s death was caused by the wrongful acts and neglect of
27 OpenAI, including designing and distributing a defective product that prioritized engagement over
28 safety, failing to take reasonable steps like notifying law enforcement of serious threats to specific
individuals, engaging in the unlicensed practice of psychotherapy, allowing and encouraging re-

1 access by the Shooter after deactivation of their account, failing to warn of foreseeable risks
2 created by the product, negligently providing safety services for ChatGPT, and deceptively
3 presenting their product as safe. Defendant Altman personally overrode safety objections and
4 rushed the product to market.

5 185. As described above, OpenAI Defendants' wrongful acts were a substantial factor in
6 causing Z.B.'s death. Had OpenAI alerted law enforcement when its own safety team
7 recommended doing so, the RCMP, who already had an open file on the Shooter and had
8 previously removed firearms from the Shooter's home, would have had the opportunity to
9 intervene before the attack. Had OpenAI maintained an effective ban rather than a reversible
10 "deactivation," the Shooter would not have re-accessed a product that validated and reinforced
11 their violent ideation and aided and abetted their plans to commit a mass shooting. Either step
12 alone would have prevented Z.B.'s death.

13 186. Z.B.'s heirs, including Plaintiff Lori Hayer, have suffered economic damages
14 including funeral and burial expenses, the reasonable value of household services Z.B. would have
15 provided, and the financial support Z.B. would have contributed.

16 187. As Z.B.'s mother, Lori Hayer has further suffered profound damages including loss
17 of Z.B.'s love, companionship, comfort, care, assistance, protection, affection, society, and moral
18 support for the remainder of her life.

19 188. Plaintiff seeks all damages recoverable under California Code of Civil Procedure
20 sections 377.60 and 377.61, including non-economic damages for loss of Z.B.'s love,
21 companionship, comfort, care, assistance, protection, affection, society, and moral support, and
22 economic damages including funeral and burial expenses, the value of household services, and the
23 financial support Z.B. would have provided.

24 **ELEVENTH CAUSE OF ACTION**
25 **SURVIVAL ACTION**
(On Behalf of Plaintiff Against OpenAI Defendants)

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

27 190. Plaintiff brings this survival claim as successor-in-interest to decedent Z.B.
28 pursuant to California Code of Civil Procedure sections 377.30 and 377.32. Plaintiff shall execute

1 and file the declaration required by section 377.32 shortly after the filing of this Complaint.

2 191. As Z.B.'s mother and successor-in-interest, Plaintiff has standing to pursue all
3 claims Z.B. could have brought had she survived, including those brought in this Complaint.

4 192. By keeping a known dangerous system on the market, refusing to notify law
5 enforcement when their own safety team recommended doing so, and allowing the Shooter to
6 regain access after deactivating their account, OpenAI Defendants acted with conscious disregard
7 for human life. Z.B.'s suffering and death were not the result of misuse or chance. They were the
8 foreseeable outcome of deliberate decisions that prioritized engagement and commercial growth
9 over the protection of human life.

10 193. Plaintiff, in her capacity as successor-in-interest, seeks all survival damages
11 recoverable under California law, including punitive damages.

12 **PRAYER FOR RELIEF**

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

- 15 1. For all damages recoverable for the wrongful death of Z.B., including non-
16 economic damages for loss of love, companionship, comfort, care, assistance, protection,
17 affection, society, and moral support, and economic damages including funeral and burial
18 expenses, the value of household services, and the financial support Z.B. would have provided;
- 19 2. For all survival damages recoverable as successor-in-interest to decedent Z.B.;
- 20 3. For punitive damages as permitted by law;
- 21 4. For injunctive relief requiring OpenAI to: (a) ban users whose accounts have been
22 deactivated for violent misuse from re-registering for ChatGPT, including through alternative
23 email addresses or sub-addresses; (b) flag new accounts linked to previously deactivated violent-
24 risk users and hold them for human review before granting access; (c) notify law enforcement
25 when internal safety systems or personnel identify a user who poses a real-world risk of violence;
26 (d) interrupt, terminate, or escalate conversations involving repeated or escalating violent ideation
27 toward real people, rather than validating or elaborating on them; (e) stop designing ChatGPT to
28 prioritize agreeable, validating responses over public safety in conversations that present a serious

1 risk of harm; (f) warn users and the public that ChatGPT’s design features can reinforce and
2 escalate violent ideation; (g) preserve prior safety flags, policy violations, and risk classifications,
3 and prohibit reversing them without documented review; and (h) submit to independent
4 monitoring and periodic compliance audits;

5 5. For prejudgment interest as permitted by law;

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

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

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

10 **JURY TRIAL**

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

12 **CANADIAN COUNSEL**

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

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

23 Respectfully submitted,

24 **LORI HAYER, INDIVIDUALLY AND
25 AS SUCCESSOR-IN-INTEREST TO
26 DECEDENT Z.B.,**

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
Tel: (310) 694-0331

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

*Attorney for Plaintiff Lori Hayer,
individually and as successor-in-interest to
Decedent Z.B.*