UNITED STATES DISTRICT COURT
SOUTHERN DISTRICT OF NEW YORK

IN RE: OPENAI, INC., COPYRIGHT INFRINGEMENT LITIGATION This Document Relates To: News Cases 1:23-cv-11195-SHS-OTW 1:24-cv-01515-SHS-OTW 1:24-cv-03285-SHS-OTW 1:24-cv-04872-SHS-OTW 1:25-cv-04315-SHS-OTW	25-md-3143 (SHS) (OTW)

NEWS PLAINTIFFS' REPLY TO OPENAI DEFENDANTS' BRIEF IN OPPOSITION TO NEWS PLAINTIFFS' MOTION FOR SANCTIONS

TABLE OF CONTENTS

I. INTRODUCTION 1

II. ARGUMENT 3

A. Sanctions are warranted under Rule 37(b) for OpenAI's violation of this Court's Preservation Order and orders to meet and confer. 3

B. Sanctions are warranted under Rule 37(e) for OpenAI's deletion of ChatGPT conversations, which has prejudiced News Plaintiffs. 6

C. Sanctions are appropriate under Rules 37(c) and 37(d) for OpenAI's repeated concealment of critical data repositories and evidence responsive to Plaintiffs' requests for discovery. 8

1. OpenAI failed to respond to discovery requests. 8

2. Sanctions are warranted under Rules 37(c) and 37(d). 10

D. Sanctions are appropriate under Rule 37(e)(2) for OpenAI's concealment and subsequent deletion of critical ChatGPT conversations 11

1. OpenAI repeatedly misrepresented to this Court and News Plaintiffs its preservation and search capabilities. 12

2. Sanctions are appropriate under Rule 37(e)(2). 13

E. The remedies News Plaintiffs request are appropriate and proportional to the harm suffered. 14

1. OpenAI should be precluded from relying on the 20 million output logs for any purpose 14

2. Fact sanctions are warranted. 15

3. Attorneys' fees are mandatory. 17

III. CONCLUSION 18

TABLE OF AUTHORITIES

Cases

Agiwal v. Mid Island Mortg. Corp.,
555 F.3d 298 (2d Cir. 2009) 14

Doubleline Capital LP v. Odebrecht Fin., Ltd.,
No. 17-cv-4576, 2021 WL 1191527 (S.D.N.Y. Mar. 30, 2021) 8

Fireman's Fund Ins. Co. v. Great Am. Ins. Co. of N.Y.,
284 F.R.D. 132 (S.D.N.Y. 2012) 10

Health & Happiness (H&H) US LLC v. Nutramax Labs., Inc.,
No. 23-cv-10849, 2026 WL 836132 (S.D.N.Y. Mar. 26, 2026) 7, 8

In re Keurig Green Mountain Single-Serve Coffee Antitrust Litig.,
341 F.R.D. 474 (S.D.N.Y. 2022) 4

Karsch v. Blink Health Ltd.,
17-cv-3880, 2019 WL 2708125 (S.D.N.Y. June 20, 2019) 8

N. Am. Photon Infotech, Ltd. v. ZoomInfo LLC,
No. 20-cv-2180, 2022 WL 4132941 (S.D.N.Y. Sept. 7, 2022) 10

Novak v. Wolpoff & Abramson LLP,
536 F.3d 175 (2d Cir. 2008) 17

Residential Funding Corp. v. DeGeorge Fin. Corp.,
306 F.3d 99 (2d Cir. 1997) 16

Wisser v. Vox Media, Inc.,
No. 19-cv-1445, 2020 WL 1547381 (S.D.N.Y. Apr. 1, 2020) 11

Rules

Fed. R. Civ. P. 26 10

Fed. R. Civ. P. 37 ..... passim

I. INTRODUCTION

The critical questions now before this Court are: (1) Did OpenAI mislead News Plaintiffs and this Court when it repeatedly said it did not have and could not provide evidence showing OpenAI trained on and output News Plaintiffs' content; (2) Did OpenAI fail to preserve this evidence as it was required to by Court order or otherwise; and (3) Do the answers to the above two questions justify the imposition of sanctions? OpenAI's brief makes clear that the answer to these long-brewing questions is "yes." Indeed, OpenAI admits the following in its Opposition, Dkt. 1672 ("Opp."):

(1) OpenAI created a Bloom filter under a project codenamed "Giraffe" to allow it to identify copyrighted content from entities that sued it in at least 97 million ChatGPT outputs (at 17);

(2) By May 2023, OpenAI searched its Common Crawl training datasets for New York Times content (at 18);

(3) Around the same time, OpenAI had a "rough-and-ready sample of 78 million [ChatGPT] conversations" created from sets of "5 million" conversations it drew "weekly" (at 18-19);

(4) Also around the same time, OpenAI de-identified the 78-million-conversation sample, a separate 10-million-conversation sample, and a separate 9-million-conversation sample (collectively the "Giraffe ChatGPT datasets") (at 36);

(5) OpenAI ran the "NYT Recall Pipeline," an evaluation of ChatGPT's "regurgitation of New York Times content," which identified how often New York Times content appeared in the 78 million conversations (at 17-18); and

(6) Within a few weeks of being sued by News Plaintiffs, OpenAI searched for News Plaintiffs' content in its Common Crawl training dataset and also ran several dozen evaluations searching for outputs containing News Plaintiffs' content in the Giraffe ChatGPT datasets (at 17-18); Ex. 2[1] at 220:5-221:24, 236:2-237:14, 248:9-249:8.[2]

[1] Unless otherwise noted, "Ex." refers to exhibits to News Plaintiffs' Motion for Sanctions ("Mot."), Dkt. 1618-1.

[2] While OpenAI minimizes the Common Crawl training dataset as "one of its many training datasets" (Opp. at 18), that dataset contains "billions of webpages," Dkt. 1749 (OpenAI Mot. Summ. J.) at 2, including [redacted] of News Plaintiffs' articles, Dkt. 1728 (News Pl.'s Mot. Summ. J.) at 39, 40, 43.

Faced with these undeniable facts, OpenAI's opposition takes two tacks. *First*, it recasts two years of obfuscation and misrepresentations as "routine discovery practice in a heavily contested litigation." Opp. at 21; *see also id.* at 2. It asks this Court to treat as routine its false representations, spoliation of evidence, disregard of express Court orders, and its disclosure of key evidence only days before the close of fact discovery and the night before a key OpenAI witness's deposition.

Second, OpenAI seeks to excuse its conduct on the ground that this case involved "novel, complex discovery." Opp. at 1. While true, the complexity of discovery was exacerbated by OpenAI's lack of candor. Rather than responding forthrightly to News Plaintiffs' discovery requests, OpenAI led Plaintiffs and this Court to believe that it could not do—and had not done—anything remotely like what Plaintiffs were seeking, even though OpenAI had already done something very close to what was requested. Put differently, what is "novel," or at least unusual, is not the discovery sought but rather OpenAI's failure to disclose information that it now admits it had, and its misrepresentations to the Court about the existence of that information and/or OpenAI's ability to produce it.

Insofar as OpenAI complains about the "time and resources" needed to respond to News Plaintiffs' discovery requests (Opp. at 3), OpenAI's 30(b)(6) witness on this issue, Mr. Monaco, testified: "[W]e had the compute." Ex. 2 at 223:14-15. Moreover, as OpenAI now admits, it had in hand the Giraffe ChatGPT datasets, as well as searches (evaluations) of those conversations for News Plaintiffs' content, that it hid from the parties and the Court until after the end of fact discovery. Finally, while OpenAI claims "fact discovery has revealed *no evidence* whatsoever to support Plaintiffs' allegations" (Opp. at 1) (emphasis in original), the Giraffe ChatGPT datasets OpenAI concealed show ChatGPT outputting News Plaintiffs' content. *See, e.g.*, Dkt. 1728 at 17-

29. It also bears repeating that OpenAI, a company on track to conduct an IPO that values the company at $1 trillion, retained at least five law firms and has had 90 attorneys enter appearances on its behalf.

As to prejudice, OpenAI is relying at summary judgment on its flawed 20 million log sample to argue that ChatGPT rarely regurgitates, Dkt. 1749 at 4—while having denied News Plaintiffs a fair opportunity to ask the over 70 OpenAI and Microsoft witnesses they deposed, or seek additional discovery, about the very project that shows their content was in fact regurgitated. Moreover, because these datasets were produced so late, News Plaintiffs' experts were not able to fully analyze the data within the Court's fact and expert discovery deadlines. After the Court ordered production, Dkt. 1419 at 1, OpenAI was able to produce the Giraffe ChatGPT conversation datasets within ten days. Had OpenAI done so much earlier, the parties would have conducted discovery quite differently.

While this case may be among the first to address discovery of generative AI companies, it will not be the last. Under the Federal Rules of Civil Procedure and this Court's inherent authority, sanctions should be imposed to ensure this conduct does not in fact become "routine," but serves to caution OpenAI (and others) that there are consequences for not following the basic rules of discovery.

II. ARGUMENT

A. Sanctions are warranted under Rule 37(b) for OpenAI's violation of this Court's Preservation Order and orders to meet and confer.

This Court's Preservation Order was unambiguous and unconditional: OpenAI had to "preserve and segregate all output log data that would otherwise be deleted on a going-forward basis." Dkt. 33 at 2. Nothing in the Order exempted "user requests to delete ChatGPT Consumer Conversation Data that is more than 30 days old" (Opp. at 11). The Court later emphasized the

Order's clarity and its expectation that OpenAI would comply: "we do have an ongoing issue about how and whether my May 13 order was complied with, given how explicit it was." Ex. 24 at 15:9–11. OpenAI admits that it failed to preserve "all output log data that would otherwise be deleted" and took no "additional steps" to preserve "chats generated before the Order's effective date." Opp. at 25. Its deliberate failure to comply with the Preservation Order warrants sanctions under Rule 37(b)(2)(A). *See In re Keurig Green Mountain Single-Serve Coffee Antitrust Litig.*, 341 F.R.D. 474, 493 (S.D.N.Y. 2022) (culpable state of mind for Rule 37(b) sanctions can be satisfied by finding a party acted knowingly in bad faith through ordinary negligence).

First, OpenAI's contemporaneous statements to this Court are inconsistent with its retelling. After the parties' repeated hearings, conferences, and filings on preservation, OpenAI told the Court on May 27, that it would "leave in-place the mechanisms it has implemented to comply with the Preservation Order." Dkt. 77, at 2. OpenAI claims this Court and the parties understood its statement to mean it would continue to delete conversations generated before the Order's effective date. But this Court's "wholesale preservation" comments at the May 27 hearing were made in response to an exchange regarding API data and enterprise ChatGPT data, and this Court subsequently modified the Order to exclude the latter. Ex. 31 at 52:8-17; Dkt. 79. Moreover, OpenAI's statement does not include any reference to such ongoing deletions, and at no other hearings, nor in any filings, did OpenAI ever state that it would continue to delete consumer ChatGPT conversations that predated the Preservation Order.[3] Dkt. 77; Ex. 3. Instead, OpenAI

[3]OpenAI acknowledges the Court's admonition "to not disclose 'protected settlement communications' that took place during these confidential and off-the-record 'settlement conferences,'" but then refers to those conferences to argue that the Court condoned its noncompliance, without ever citing specifics. Opp. at 11, n. 4. News Plaintiffs will not violate this Court's instruction, but OpenAI's characterization is entirely inconsistent with the undersigned attorneys' recollections and the record.

represented that it was "[w]illing to move forward on the consumer side." Ex. 31 at 46:21-22.

Second, OpenAI knew how to seek an amendment to the Court's orders; it thrice invoked that process and was thrice rejected. OpenAI first moved for reconsideration two days after the Preservation Order. Dkt. 40. This Court denied that motion without prejudice. Dkt. 42. A week later, OpenAI filed a supplemental opposition, renewing technical-infeasibility and burden objections (Dkt. 64), and it pressed those objections at the May 27, 2025 hearing (Ex. 31 at 26:24-27:17). This Court then rejected OpenAI's request for reconsideration. Dkt. 79. On June 4, OpenAI filed a Rule 72(a) objection to Judge Stein (Dkt. 92), who denied OpenAI's request and affirmed Magistrate Judge Wang's Order (Dkt. 271). On a docket with more than 1,900 entries, not one indicates that the Court modified the Preservation Order or granted OpenAI a reprieve to delete consumer ChatGPT conversations predating the Order.

Third, OpenAI's argument that it was permitted to delete consumer ChatGPT conversations predating the Order is inconsistent with the parties' ongoing conversations and conferences with the Court regarding discovery into those very same conversations. OpenAI was ordered to preserve those conversations identified in the sampling exercise but chose not to.

* * *

In addition to the Preservation Order, OpenAI violated the Court's numerous orders to meet and confer in good faith to resolve disputes over the production of the ChatGPT conversations. *See, e.g.,* Dkt. 42 (directing the parties to meet and confer on preservation issues). It should go without saying that OpenAI violated these Court orders (not to mention rules governing legal practice) when it engaged in numerous meet and confers without ever revealing the conversations it had already compiled and evaluated or its existing preservation abilities.

B. Sanctions are warranted under Rule 37(e) for OpenAI's deletion of ChatGPT conversations, which has prejudiced News Plaintiffs.

Sanctions are likewise appropriate under Rule 37(e)(1). The ChatGPT conversations OpenAI deleted after the Preservation Order "should have been preserved," but OpenAI "fail[ed] to take reasonable steps to preserve [them], and [they] cannot be restored or replaced through additional discovery." Fed. R. Civ. P. 37(e).

OpenAI argues lack of prejudice, stating that News Plaintiffs have access to other ChatGPT conversations and the "inability to access still more logs does not cause *actual* prejudice to their case." Opp. at 33 (emphasis in original). In a carefully qualified statement, OpenAI rests its argument on its assertion that the sampling exercise revealed "*essentially no statistically significant difference* between the data that OpenAI retains and the data OpenAI would have retained had it paused all deletions." Opp. at 13 (emphasis in original). OpenAI is incorrect. OpenAI's corruption of the 20 million sample deprived News Plaintiffs of discovery into at least two categories of conversations that were disproportionately likely to include evidence of that infringement: temporary conversations and RAG conversations.

Temporary Conversations: The spoliation sampling exercise demonstrated that temporary conversations [redacted] Dkt. 599 at 2.

OpenAI concedes that [redacted]. Opp. at 5. Applied to the 20 million ChatGPT population, this percentage translates to News

Plaintiffs having been deprived of approximately [redacted] conversations that were disproportionately more likely to contain favorable evidence to support their infringement claims. Extrapolated to the billions of ChatGPT conversations, these percentages translate to at least [redacted] temporary chats that OpenAI deleted. Ex. 23 at 10–11.

RAG Conversations: The record also shows that RAG conversations are disproportionately more likely to contain evidence favorable to News Plaintiffs. *See, e.g.*, Dkt. 1728 at 6-7, 16-21 (describing "hundreds of Custom GPTs designed to access, store, provide, or view news content" and providing exemplary RAG conversations outputting hundreds of words from Plaintiffs' articles); Dkt. 1742 ¶¶ 1020-21 ([redacted] of RAG conversations in the 20 million sample relied on a news website). News Plaintiffs' expert has identified three concrete indicators that such conversations are disproportionately absent from the 20 million ChatGPT conversations. *First*, the "processing errors" that OpenAI invoked to justify "substitutions" of 5% of the 20 million sample occurred in ChatGPT conversations with larger file sizes, which are consistent with the larger file sizes of conversations that use RAG. Ex. 27 (T. Goldstein Opening Report) at 164. *Second*, only about [redacted] of the 20 million ChatGPT conversations contain any RAG data, which is far below rates suggested by OpenAI's own documents. Ex. 27 at 164. *Third*, the Giraffe ChatGPT conversations that OpenAI deemed suitable for its own regurgitation evaluations contain over ***ten times*** the amount of RAG conversations as in the 20 million sample. *See* Declaration of Jennifer Maisel in Support of News Plaintiffs' Reply to OpenAI Defendants' Opposition to News Plaintiffs' Motion for Sanctions ("Maisel Decl."), Ex. 69 (Supplemental Expert Report of Dr. Tom Goldstein), ¶ 80.

OpenAI's reliance (Opp. at 32) on *Health & Happiness (H&H) US LLC v. Nutramax Labs., Inc.*, No. 23-cv-10849, 2026 WL 836132 (S.D.N.Y. Mar. 26, 2026), is misplaced. There, the

Court's conclusion that the defendant could not demonstrate "*actual* prejudice" was made in the context of rejecting the defendant's intrinsic prejudice argument. *Id.* at *12 (emphasis in original). This decision did not alter this District's standard that a party moving under Rule 37(e)(1) must show only that the spoliated evidence "would be helpful to them in connection with establishing several elements of their . . . claims." *See Doubleline Capital LP v. Odebrecht Fin., Ltd.*, No. 17-cv-4576, 2021 WL 1191527, at *7 (S.D.N.Y. Mar. 30, 2021); *see also Karsch v. Blink Health Ltd.*, 17-cv-3880, 2019 WL 2708125, at *21 (S.D.N.Y. June 20, 2019) ("It is sufficient if the existing evidence plausibly 'suggests' that the spoliated ESI could support the moving party's case."). News Plaintiffs have done so here.

C. Sanctions are appropriate under Rules 37(c) and 37(d) for OpenAI's repeated concealment of critical data repositories and evidence responsive to Plaintiffs' requests for discovery.

In May 2023, OpenAI had a "rough-and-ready sample of 78 million conversations" that it was already using to search for The Times's content. Opp. at 18. But OpenAI did not produce the Giraffe ChatGPT conversations until April 2026, after the Court ordered production, and two months after fact discovery closed. This timeline alone establishes that OpenAI failed to properly respond to discovery, and sanctions are appropriate under Rules 37(c) and 37(d).

1. OpenAI failed to respond to discovery requests.

OpenAI argues that News Plaintiffs requested something "entirely different" from the material OpenAI withheld. Opp. at 17. This is false. The Project Giraffe materials were responsive to multiple discovery requests, including:

- "Documents concerning the study, testing, or evaluation . . . for hallucinations, memorization, and any problematic behaviors," Ex. 13 at 12 (Times RFP No. 39);

- "Documents and communications concerning Defendants' development of filters and/or classifiers related to anti-regurgitation or copyright efforts, including the Giraffe filter," Ex. 38 at 11 (No. 121);
- "Documents and communications concerning Defendants' anti-regurgitation efforts," *id.* (No. 122);
- "[D]ocuments reflecting user sessions," "including *any* logs," Ex. 67 at 6 (Times Interrogatory No. 9) (emphasis added); and
- "Query, session, and chat logs related to Times Content reflecting user sessions on Your Generative AI Products and Services ('Session Data'), including, *for each session*, user queries paired with responses to those queries." Ex. 15 at 2 (emphasis added).

News Plaintiffs' requests were not expressly or by implication limited to "historic, detailed conversation logs" (Opp. at 17)—a phrase that OpenAI appears to have conceived for purposes of its opposition. What is more, while resisting production, OpenAI misrepresented the nature of the Giraffe ChatGPT conversations. When Plaintiffs asked directly whether the conversations were "analyzed to identify instances of infringement" or were "part of a regurgitation-related pipeline," OpenAI replied: "That belief was mistaken," Ex. 20 at 4. As it turns out, it was not. Mr. Monaco admitted in his second deposition, after the end of fact discovery, that by mid-2024 OpenAI had already conducted multiple searches of the Giraffe ChatGPT datasets for News Plaintiffs' content. *See, e.g.*, Ex. 2 at 236:2-9.

OpenAI also argues that the Giraffe ChatGPT conversations are not "representative." Opp. at 19. That is not OpenAI's call; the conversations were responsive and relevant, OpenAI was required to produce them. Nor does the purported absence of "complete metadata" excuse

concealment of the dataset. The Giraffe ChatGPT conversations contain the user queries paired with responses to those queries, which News Plaintiffs searched to identify their content. *See, e.g.*, Ex. 69 at ¶¶94-103; Dkt. 1673-21 (Kellogg Decl.) ¶ 3. Once News Plaintiffs showed relevance, as they did here, OpenAI bore the burden of justifying any limitation. *See Fireman's Fund Ins. Co. v. Great Am. Ins. Co. of N.Y.*, 284 F.R.D. 132, 135 (S.D.N.Y. 2012).

OpenAI now argues that it complied with Rule 26(e) because it "produced the evaluations" in 2025. Opp. at 35. But OpenAI made different representations to this Court. When News Plaintiffs moved to compel production of the Giraffe ChatGPT conversations in February 2026, OpenAI opposed their production because that "would threaten the privacy of the tens of millions of ChatGPT users who generated the conversation data therein." Dkt. 1295 at 2. OpenAI never argued that the conversations had already been produced. If, as OpenAI now asserts, the conversations had already been produced, why did OpenAI oppose the motion to compel, rather than explain that they had already been produced? The most likely answer is: they hadn't been (which is why OpenAI refused to respond to News Plaintiffs' requests for proof). *See* Ex. 37. The alternative explanation, that these conversations were so well-hidden in OpenAI's cumbersome training data sandbox that neither News Plaintiffs nor OpenAI knew about them, does not help OpenAI. An inadvertent, unnoticed, and unexplained disclosure cannot cure the violation. *See N. Am. Photon Infotech, Ltd. v. ZoomInfo LLC*, No. 20-cv-2180, 2022 WL 4132941, at *3–4 (S.D.N.Y. Sept. 7, 2022) (a "passing reference" is not sufficient disclosure).

2. Sanctions are warranted under Rules 37(c) and 37(d).

As noted above, the Giraffe ChatGPT conversations and evaluations were responsive to News Plaintiffs' requests for production, inspection, and interrogatories. OpenAI effectively denied the existence of the conversations, then produced them pursuant to Court order only after

fact discovery had closed. Whether viewed as a failure to respond to discovery under Rule 37(d) or a failure to supplement a response under Rule 37(c)(1), the rules allow the same set of sanctions.

OpenAI largely ignores Rule 37(c), but its conduct qualifies as a "fail[ure] to provide information . . . required by Rule 26(a) or (e)." Simply stated, Rule 26(e) requires supplementation of a discovery response that a party knows is "incomplete or incorrect," and OpenAI's failure to produce the 78 million output logs and related materials qualifies.

OpenAI argues that Rule 37(d) requires a "complete failure to respond," and that it served objections to the requests. Opp. at 35. But OpenAI never objected to the production of the evidence it had because it never disclosed that the evidence existed. It also repeatedly misrepresented to the Court that it did not have the capacity to provide responsive material. This is not a case in which a party objected to the production of evidence based on relevance or privilege, or in which an interrogatory response was merely "incomplete." *See Wisser v. Vox Media, Inc.*, No. 19-cv-1445, 2020 WL 1547381, at *5 (S.D.N.Y. Apr. 1, 2020) (Opp. at 35). Under any reasonable construction of the term, this constituted a "complete failure to respond."

D. Sanctions are appropriate under Rule 37(e)(2) for OpenAI's concealment and subsequent deletion of critical ChatGPT conversations.

What OpenAI now admits stands in stark contrast to what OpenAI said during discovery. OpenAI now admits that in May 2023, OpenAI drew "5 million" conversations from various "weekly" sets of ChatGPT conversations. Opp. at 19. In May 2023, OpenAI searched for "New York Times content in OpenAI's training data." *Id.* at 18. OpenAI then ran evaluations over the large datasets it created and "conducted an evaluation of possible regurgitation of New York Times Content." *Id.* at 18; Ex. 46 [redacted]). But for two years, OpenAI staunchly maintained that: (1) it did "not currently have tools to search [the training datasets] efficiently" (Ex. 1 at 7) and, (2) its "infrastructure . . . does not include *any* mechanism

that could search a large dataset of output logs like decompressed Conversation Data for specific content," such as News Plaintiffs' (Ex. 3) (cleaned up). These were not isolated statements. Repeatedly, OpenAI claimed it did not have the ability to search for News Plaintiffs' content. *See, e.g.*, *Times* Dkt. 423 at 2.

These representations were false.

1. OpenAI repeatedly misrepresented to this Court and News Plaintiffs its preservation and search capabilities.

OpenAI now recasts its statements as making a distinction between searching its "*full universe*" of conversation data and searching a smaller subset of those conversations. Opp. at 18 (emphasis added). But OpenAI told this Court it lacked *any* mechanism to search output logs for specific content, when it had already conducted several searches of that kind. In a sworn declaration to this Court, OpenAI said it could not "search a large dataset" for specific content (Ex. 3 at 3), a fact that OpenAI and its witnesses now acknowledge was untrue. OpenAI never disclosed in any discovery response, filing, hearing, settlement conference, or meet-and-confer letter, email, or call that it could, and had, preserved and searched *vast amounts* of ChatGPT conversations for News Plaintiffs' content. OpenAI does not try to excuse these misrepresentations, hiding instead behind irrelevancies, like "incomplete metadata." *See* Opp. at 3, 6, 16, 17, 18, 19, 20.

Because OpenAI hid its technological capabilities and existing infrastructure, vast quantities of relevant information were deleted. Had OpenAI not withheld this information, News Plaintiffs could have requested that OpenAI preserve ChatGPT conversations on a weekly basis and searched those conversations for News Plaintiffs' content, just as OpenAI had done previously. In fact, OpenAI makes no claim that repeating what it had already done—amassing a large dataset and searching it—would have been especially burdensome. Instead, OpenAI forced the Court and News Plaintiffs to design a sampling process based on a misrepresentation of what it could do.

Then, following months of negotiations and agreement on the collection of a 20 million ChatGPT conversation sample—which Mr. Monaco said would take [redacted] to prepare (Ex. 32 at 2)—OpenAI doctored the sample by substituting and deleting 10% of the requested 20 million conversations, despite the Preservation Order.

2. Sanctions are appropriate under Rule 37(e)(2).

OpenAI's response is that its users get to decide what it deletes, and it had no obligation to change that system. Opp. at 38-40. Specifically, OpenAI argues its "retention policies were established before this litigation . . . and maintained consistently throughout." Opp. at 39. This misses the point: the litigation imposed new preservation obligations, especially (but not exclusively) after the Court entered the Preservation Order. The record shows in other ways that OpenAI intended to deprive News Plaintiffs of relevant information. Fed. R. Civ. P. 37(e)(2). News Plaintiffs' requests were straightforward, yet OpenAI continued to conceal and delete relevant data. While there may be inadvertent deletions or failures to produce, this was not inadvertent. On the one hand, OpenAI asked Magistrate Judge Wang to reconsider the Preservation Order—and appealed the Order to Judge Stein—all to no avail. On the other hand, when asked about steps OpenAI took to preserve data that was routinely deleted, OpenAI's Rule 30(b)(6) witness confirmed that [redacted]. *See* Ex. 2 at 276:20-277:4.

OpenAI's intent is further evidenced by the statements of its CEO, Sam Altman, who publicly denounced the Preservation Order on June 5, 2025. Five days later, [redacted] Ex. 25 at 95:3-6; Ex. 27 at ¶ 436. OpenAI provides no justification for [redacted].

Finally, OpenAI also selectively preserved evidence that served its interests: [redacted], while simultaneously protesting it could not search those same logs for News Plaintiffs' content, Mot. at 15. This selective preservation and deletion pattern demonstrates that OpenAI's conduct was not the product of technical limitations or routine data management, but rather a deliberate strategy to manipulate the evidentiary record.

E. The remedies News Plaintiffs request are appropriate and proportional to the harm suffered.

1. OpenAI should be precluded from relying on the 20 million output logs for any purpose.

Preclusion of the 20 million sample is warranted under Rule 37(b) and (e)(1) given OpenAI's violation of this Court's Preservation Order, which required OpenAI to preserve and segregate output logs, Ex. 23 at 15, and for OpenAI's intentional deletion of relevant output logs. Mot. at 23-33. OpenAI's conduct was willful and persisted throughout discovery, and OpenAI was on notice of the consequences. *See Agiwal v. Mid Island Mortg. Corp.*, 555 F.3d 298, 302 (2d Cir. 2009).

Lesser sanctions cannot cure the prejudice to News Plaintiffs. OpenAI's deletions were permanent. Similarly, there is no new pool of data from which to constitute a new sample that has not been tainted by OpenAI's deletions. OpenAI argues that the 20 million sample and the Giraffe ChatGPT datasets have been produced, and Mr. Monaco sat for a supplemental deposition. Opp. at 25-26. But, beyond the fact that it was negotiated without Plaintiffs or the Court knowing the truth about OpenAI's search and retention capabilities, the 20 million sample is different from the sample the parties negotiated and agreed to; OpenAI deleted logs it had agreed to produce (and unilaterally substituted others). Moreover, production of the Giraffe ChatGPT conversations and the supplemental deposition occurred at the end of, or after the end of, fact discovery, precluding

Plaintiffs from fully exploring them with witnesses and through follow-up document requests.

Allowing OpenAI to use the 20 million sample would reward the very misconduct this motion addresses. OpenAI relies extensively on the 20 million sample to support its argument on summary judgment that "OpenAI's LLMs are not designed to 'regurgitate' their training data." Dkt. 1749 at 4. But as described above, the sample is not representative. It excludes temporary chats predating the Preservation Order, which News Plaintiffs' experts [redacted] Dkt. 599 at 2. The sample also underrepresents RAG conversations, in which OpenAI retrieved News Plaintiffs' content in real time to respond to user queries and which are central to News Plaintiffs' output-based claims. And even if OpenAI disputes these points, the fact remains that this is an issue that should have been explored in discovery—but OpenAI's misrepresentations and delays in production made that impossible.

OpenAI does not dispute that the 20-million log sample was not the sample to which the parties agreed, and that the differences are the result of OpenAI's deletion and substitution of logs that should have been preserved and produced. In these circumstances, all of which are the result of OpenAI's misconduct, it should not be allowed to rely on the sample.

2. Fact sanctions are warranted.

In addition to preclusion, News Plaintiffs ask that the Court grant Plaintiffs' request for the following fact sanctions:

i. This Court should find that ChatGPT's output logs include (or would have shown if properly produced in discovery) substantial and systematic grounding on and regurgitation of News Plaintiffs' copyrighted material in outputs.

ii. OpenAI should be prohibited from arguing that ChatGPT's output logs do not include (or would not have shown if properly produced in discovery) substantial and systematic grounding on and regurgitation of News Plaintiffs' copyrighted material into user

outputs.

iii. This Court should instruct the jury on these findings and their binding effect.

Mot. at 4-5.

OpenAI argues that the imposition of established-fact sanctions would run afoul of *Residential Funding Corp. v. DeGeorge Fin. Corp.*, 306 F.3d 99 (2d Cir. 1997). The uncontroverted record shows otherwise.

First, OpenAI had control over the concealed information and "had an obligation to timely produce it," *id.* at 107. OpenAI contends that it "timely produced *much* of that information" (Opp. at 36) (emphasis added) on January 26, 2026, and on March 20, 2026. These dates are, respectively, just a few weeks before and a few weeks *after* the close of fact discovery. OpenAI's suggestion that these productions of 97 million conversations allowed News Plaintiffs "to incorporate [the material] into their expert reports and overall case strategy" (Opp. at 36) is absurd.

Second, as described above, OpenAI had a "culpable state of mind." It made repeated misrepresentations to News Plaintiffs and this Court. OpenAI does not dispute that it and its lawyers knew about the concealed materials. Moreover, once OpenAI's deceit was uncovered, the technical obstacles to prompt production disappeared. OpenAI produced the 78 million sample ten days after being ordered to do so, and the 9 million sample one week later.

Third, OpenAI argues that the material is not sufficiently relevant. Opp. at 37. Really? For OpenAI to argue now that the ChatGPT conversations "would have made little difference in the underlying case" (Opp. at 38) (cleaned up) beggars belief in light of OpenAI's recently filed summary judgment motion, Dkt. 1749 at 4. And given the timing of the eventual productions at the end of and even after the end of fact discovery, OpenAI's argument that the News Plaintiffs

now have access to the material, and were not prevented from litigating their claims (Opp. at 38), fares no better.

3. Attorneys' fees are mandatory.

An award of the attorneys' fees and costs incurred as a direct result of OpenAI's discovery misconduct is mandatory. Under Rule 37(b)(2)(C), when a court grants a motion for sanctions, the court must, after giving an opportunity to be heard, require "the disobedient party, the attorney advising that party, or both to pay the reasonable expenses, including attorney's fees, caused by the failure, unless the failure was substantially justified or other circumstances make an award of expenses unjust." Similarly, under Rule 37(d)(3), sanctions "must" include reasonable expenses and attorneys' fees unless the failure was "substantially justified." *Novak v. Wolpoff & Abramson LLP*, 536 F.3d 175, 178 (2d Cir. 2008).

OpenAI points to "genuine disputes" as to the reasonableness of their conduct. Opp. at 40. The disputes OpenAI cites are not "genuine"—they are a rehash of OpenAI's previous misrepresentations. For example, the "dispute" over whether OpenAI could search its data warehouse at scale was not a legitimate discovery disagreement; it was the product of false declarations from Mr. Monaco, who swore under oath that building a search tool would take "at least [redacted] of dedicated full-time engineer time" and cost [redacted] when OpenAI had already built and deployed comparable tools for its own purposes. Ex. 3. The "dispute" over the appropriate sample size was similarly artificial: OpenAI argued for a smaller sample based on "technical costs" and engineering burden that turned out not to exist. OpenAI had already decompressed, anonymized, processed, and stored over 80 million output logs. Ex. 2 at 234:5-21.

OpenAI cannot claim that "other circumstances make an award of expenses unjust." Fed. R. Civ. P. 37(a)(5)(A)(iii). OpenAI has huge financial resources, including enough money to cash

out its executives for billions of dollars, and it has the most sophisticated legal and technical resources available to it. Its misconduct was not the product of confusion or limited resources, but a deliberate litigation strategy. The fees and costs News Plaintiffs incurred, including cross-country trips to the "sandbox," hundreds of hours reviewing over-redacted productions, multiple motions to compel, and the re-deposition of Mr. Monaco, were the direct and foreseeable consequence of OpenAI's decision to conceal its capabilities and obstruct discovery.

III. CONCLUSION

For the foregoing reasons, the Court should grant News Plaintiffs' motion for sanctions.

Dated: September 10, 2026

/s/ Steven Lieberman

Steven Lieberman (SL8687)
Jennifer B. Maisel (5096995)
Robert Parker *(admitted pro hac vice)*
Jenny L. Colgate *(admitted pro hac vice)*
Alexandra Hughes *(admitted pro hac vice)*
Mark Rawls *(admitted pro hac vice)*
Michael H. Jones *(admitted pro hac vice)*
Kristen J. Logan *(admitted pro hac vice)*
Bryan B. Thompson (6004147)
Mary L. Mullins *(pending pro hac vice)*
ROTHWELL, FIGG, ERNST & MANBECK, P.C.
901 New York Avenue, N.W., Suite 900 East
Washington, DC 20001
Telephone: (202) 783-6040
Facsimile: (202) 783-6031
slieberman@rothwellfigg.com
jmaisel@rothwellfigg.com
rparker@rothwellfigg.com
jcolgate@rothwellfigg.com
ahughes@rothwellfigg.com
mrawls@rothwellfigg.com
mjones@rothwellfigg.com
klogan@rothwellfigg.com
bthompson@rothwellfigg.com
mmullins@rothwellfigg.com

Attorneys for Daily News Plaintiffs

/s/ Ian Crosby

Ian Crosby *(pro hac vice)*
Katherine M. Peaslee *(pro hac vice)*
Susman Godfrey L.L.P.
401 Union Street, Suite 3000
Seattle, WA 98101
Telephone: (206) 516-3880
Facsimile: (206) 516-3883
icrosby@susmangodfrey.com
kpeaslee@susmangodfrey.com

Davida Brook *(pro hac vice)*
Emily K. Cronin (*pro hac vice*)
Adnan Muttalib *(pro hac vice)*
Susman Godfrey L.L.P.
1900 Ave of the Stars, Suite 1400
Los Angeles, CA 90067
Telephone: (310) 789-3100
Facsimile: (310) 789-3150
dbrook@susmangodfrey.com
ecronin@susmangodfrey.com
amuttalib@susmangodfrey.com

Elisha Barron (5036850)
Zachary B. Savage (ZS2668)
Alexander Frawley (5564539)
Eudokia Spanos (5021381)
Susman Godfrey L.L.P.
One Manhattan West
New York, NY 10001
Telephone: (212) 336-8330
Facsimile: (212) 336-8340
ebarron@susmangodfrey.com
zsavage@susmangodfrey.com
afrawley@susmangodfrey.com
espanos@susmangodfrey.com

Steven Lieberman (SL8687)
Jennifer B. Maisel (5096995)
Kristen J. Logan *(pro hac vice)*
Rothwell, Figg, Ernst & Manbeck, P.C.
901 New York Avenue, N.W., Suite 900 East
Washington, DC 20001

Telephone: (202) 783-6040
Facsimile: (202) 783 6031
slieberman@rothwellfigg.com
jmaisel@rothwellfigg.com
klogan@rothwellfigg.com

Attorneys for Plaintiff
The New York Times Company

/s/ Matthew Topic

Jon Loevy *(pro hac vice)*
Michael Kanovitz *(pro hac vice)*
Matthew Topic *(pro hac vice)*
Thomas Kayes *(pro hac vice)*
Steven Art *(pro hac vice)*
Kyle Wallenberg *(pro hac vice)*
Shelley Geiszler *(pro hac vice)*
LOEVY & LOEVY
311 North Aberdeen, 3rd Floor
Chicago, IL 60607
312-243-5900 (p)
312-243-5902 (f)
jon@loevy.com
mike@loevy.com
matt@loevy.com
kayes@loevy.com
steve@loevy.com
wallenberg@loevy.com
geiszler@loevy.com

Stephen Stich Match (No. 5567854)
Lauren Sonnenberg *(pro hac vice)*
LOEVY & LOEVY
222 Sutter Street, Suite 600A
San Francisco, CA 94108
(312) 243-5900
match@loevy.com
sonnenberg@loevy.com

Attorneys for Plaintiffs The Center for Investigative Reporting, Inc. and The Intercept Media, Inc.

/s/ Lacy H. Koonce, III
Lacy H. ("Lance") Koonce, III
Matthew A. Leish
Brendan Kehoe
Lucy Hoffman
Maya Katalan
Klaris Law PLLC
161 Water Street
New York, NY 10038
Telephone: (646) 779-4882
lance.koonce@klaris.com
matthew.leish@klarislaw.com
brendan.kehoe@klarislaw.com
lucy.hoffman@klarislaw.com
maya.katalan@klarislaw.com

Guy Ruttenberg *(pro hac vice)*
Ruttenberg IP Law, A Professional Corporation
1801 Century Park East
Suite 1920
Los Angeles, CA 90067
Telephone: (310) 627-2270
ZD-AICopyright@ruttenbergiplaw.com

Attorneys for Plaintiffs Ziff Davis, Inc., Ziff Davis, LLC, IGN Entertainment, Inc., Everyday Health Media, LLC, Mashable, Inc., and CNET Media, Inc.