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) Hon. Sidney H. Stein Hon. Ona T. Wang

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

TABLE OF CONTENTS

Page

INTRODUCTION 1

STATEMENT 4

A. OpenAI Retained 50 Billion User Chats and Made Them Available to Plaintiffs for Sampling 4

1. OpenAI Maintains Robust Data Retention Policies and Procedures 5

2. OpenAI Informs Plaintiffs and the Court of Its Preservation Policies Early in the Litigation 7

3. The Court Enters a Preservation Order But Does Not Require "Wholesale" Preservation 9

4. A Court-Ordered Sampling Procedure Confirms that OpenAI's Retained Chats are Representative of How Users Utilize Its Products 12

5. OpenAI Creates and Produces the Merits Sample, A Representative Sample of 20 Million Output Logs 13

B. OpenAI Makes its Training Data Available for Inspection at OpenAI's Expense 15

C. OpenAI Has Not Misrepresented Its Discovery Capabilities to the Court 16

1. Project Giraffe and Anti-Regurgitation Efforts 17

2. Project Giraffe is Consistent with OpenAI's Representations to the Court 18

3. OpenAI has Produced Relevant Documents and Datasets Pertaining to Its Anti-Regurgitation Efforts 21

ARGUMENT 22

I. Plaintiffs Fail to Establish a Basis for Preclusion of the Merits Sample Under Rule 37 24

A. Rule 37(b) Does Not Support Preclusion 24

B. Rule 37(e)(1) Does Not Support Preclusion 31

II. This Court Should Not Enter Established Fact Sanctions Under Rule 37 33

III. Plaintiffs' Rule 37(e)(2) Claim Fails Because OpenAI Did Not Act with Intent to Deprive 38

IV. Plaintiffs Are Not Entitled to Attorneys' Fees 40

V. The Court Should Not Exercise Its Inherent Authority to Impose Sanctions 43

CONCLUSION 45

TABLE OF AUTHORITIES

Page

CASES

Agiwal v. Mid Island Mortg. Corp., 555 F.3d 298 (2d Cir. 2009) 24

Chambers v. NASCO, Inc., 501 U.S. 32 (1991) 42

Cine Forty-Second St. Theatre Corp. v. Allied Artists Pictures Corp., 602 F.2d 1062 (2d Cir. 1979) 36

Eid v. Koninklijke Luchtvaart Maatschappij N.V., 310 F.R.D. 226 (S.D.N.Y. 2015) 43

Evans v. Griffin, 932 F.3d 1043 (7th Cir. 2019) 34

Health & Happiness (H&H) US LLC v. Nutramax Lab'ys, Inc., 2026 WL 836132 (S.D.N.Y. Mar. 26, 2026) 32

Hoffer v. Tellone, 128 F.4th 433 (2d Cir. 2025) 2, 23, 36, 39

Joint Stock Co. Channel One Russ. Worldwide v. Infomir LLC, 2018 WL 4760345 (S.D.N.Y. Sept. 28, 2018) 34

Klein v. Torrey Point Grp., LLC, 979 F. Supp. 2d 417 (S.D.N.Y. 2013) 41

Kortright Cap. Partners v. Investcorp Inv. Advisers Ltd., 330 F.R.D. 134 (S.D.N.Y. 2019) 37, 38

Outley v. City of New York, 837 F.2d 587 (2d Cir. 1988) 29

Rakower Law PLLC v. Berner, 2022 WL 17820861 (S.D.N.Y. Dec. 14, 2022) 41

Residential Funding Corp. v. DeGeorge Fin. Corp., 306 F.3d 99 (2d Cir. 2002) 23, 35, 36, 37

Softel, Inc. v. Dragon Med. & Sci. Commc'ns, Inc., 118 F.3d 955 (2d Cir. 1997) 30

Update Art, Inc. v. Modiin Publ'g, Ltd., 843 F.2d 67 (2d Cir. 1988) 29

Wisser v. Vox Media, Inc., 2020 WL 1547381 (S.D.N.Y. Apr. 1, 2020) 35

Wolters Kluwer Fin. Servs., Inc. v. Scivantage, 564 F.3d 110 (2d Cir. 2009) 23, 43, 44

Zubulake v. UBS Warburg LLC, 220 F.R.D. 212 (S.D.N.Y. 2003) 32

Page

RULES

Fed. R. Civ. P.
26.....34, 35
33.....34
37.....*passim*
72.....15

OTHER AUTHORITIES

API Platform, OpenAI, https://openai.com/api.....7

Data Controls in the OpenAI Platform, OpenAI,
https://developers.openai.com/api/docs/guides/your-data.....7

8B Wright & Miller, Federal Practice & Procedure (3d ed. 2026)
§ 2284.2.....34
§ 2289.....34
§ 2291.....35

INTRODUCTION

News Plaintiffs' motion for sanctions attempts to rewrite history, ignores OpenAI's good-faith compliance with this Court's discovery orders, and seeks punitive sanctions that are unsupported by the law or the facts. Here are the facts: OpenAI has spent hundreds of hours and millions of dollars working with the Court and opposing counsel to make one of the largest, most complex, and unprecedented productions in American litigation history. OpenAI's efforts included producing a representative sample of 20 million ChatGPT conversation logs that News Plaintiffs themselves selected (the "Merits Sample"). OpenAI made its training data available for inspection at its own expense and even supplied on-call engineers to fix technical problems of Plaintiffs' own making. And OpenAI produced troves of information related to Project Giraffe, an anti-regurgitation project Plaintiffs have known about for years.

These novel, complex discovery feats required managing tens of billions of conversation records and hundreds of terabytes of training data, all while implicating the privacy interests of hundreds of millions of users in over 160 countries. At every turn, OpenAI pioneered solutions, coordinated with News Plaintiffs, and—most importantly—followed this Court's orders.

Despite all this, Plaintiffs now seek sweeping discovery sanctions. Plaintiffs' motion asks the Court to preclude OpenAI from relying on the very Merits Sample that both sides and the Court spent eight months collaboratively designing; to enter sweeping factual findings; and to award attorneys' fees for ordinary litigation activity. There is no mystery about why Plaintiffs seek those extreme sanctions now: fact discovery has revealed *no evidence* whatsoever to support Plaintiffs' allegations that ChatGPT systematically reproduced Plaintiffs' work. Rather than accepting the stark reality that their claims fail on the merits, Plaintiffs instead seek to use discovery sanctions as an end run around proving their case. Each request fails on its own terms.

Plaintiffs' initial preclusion argument under Rule 37(b) fails because OpenAI did not willfully violate the Court's Preservation Order. Immediately after the May 13, 2025 Order, OpenAI took the unprecedented step of provisionally retaining, in a segregated legal-hold database, *all* consumer ChatGPT logs generated between April 14, 2025 (*i.e.*, 30 days prior to the Order) and September 26, 2025, when the Order terminated. When OpenAI investigated fully disabling its deletion pipelines to prevent users from removing conversations generated *before* the Court's Preservation Order (and already saved in a compressed data warehouse), OpenAI determined that doing so was technically infeasible and potentially disruptive. OpenAI promptly informed the Court of these technical limitations and the Court expressly clarified on May 27, 2025, that its Order did not require "wholesale" preservation or require OpenAI to stop all "deletion processes." OpenAI openly and repeatedly explained the steps it took to comply with the Order at the May 27 hearing, on the public docket, and during off-the-record settlement conferences that the Court convened in June, August, and September 2025.

Plaintiffs' demand for court-ordered factual findings fares no better. Rule 37 authorizes such sanctions only in specifically enumerated circumstances—violation of a court order under Rule 37(b)(2)(A), failure to provide information under Rule 37(c)(1), or a wholesale failure to respond to discovery under Rule 37(d)(1)(A)(ii). None of those predicates applies here. Plaintiffs instead implausibly dress up ordinary disputes about the sufficiency and speed of discovery responses as the kind of severe disobedience Rule 37 addresses.

Nor can Plaintiffs satisfy Rule 37(e)(2)'s exacting intent standard for sanctions. This provision requires proof of "*intentional* loss or destruction of evidence"; mere negligence or knowing conduct "will not suffice." *Hoffer v. Tellone*, 128 F.4th 433, 438 (2d Cir. 2025). OpenAI's retention policies were established before these suits were filed, communicated to the

Court and to Plaintiffs at the beginning of this suit, and maintained consistently throughout. Plaintiffs raised no issues with OpenAI's retention policies for nearly a year. OpenAI deleted the conversation logs at issue due to user requests and pre-existing product features that OpenAI disclosed to Plaintiffs and this Court from the very outset of this litigation. Transparent adherence to pre-existing policies is not intentional destruction.

Plaintiffs' overarching claim that OpenAI "misrepresented" its technical capabilities collapses under scrutiny. OpenAI's Project Giraffe work began before this litigation and included analyses that relied on limited training data and unrepresentative samples of chat records drawn from datasets with incomplete metadata. What OpenAI told the Court was and remains true: the time and resources required to search compressed chat records in a data warehouse of 50-billion-plus conversations with complete metadata differ fundamentally from OpenAI's pre-litigation limited searches. OpenAI also disclosed to the Court that it had *some* existing capacity to search its training data and even offered to run Plaintiffs' searches for them. That OpenAI's existing tools could not execute the hundreds of systematic searches Plaintiffs demanded across all training datasets and all output logs is not a misrepresentation—it is a technical reality.

Plaintiffs also fail to establish prejudice. A court-ordered sampling exercise has confirmed that the exclusion of user-deleted and temporary chats from the Merits Sample did not bias that sample. Plaintiffs thus possess everything to which they are entitled. They have the Merits Sample. They have 87 million additional conversation records. They have Project Giraffe source code, evaluation data, and related communications. And they have had unprecedented access to OpenAI's training data. This Court has already resolved many of the issues Plaintiffs now raise through targeted discovery orders and court-ordered settlement conferences. The Court should refuse Plaintiffs' invitation to destroy all those efforts, and it should deny the motion.

STATEMENT

Throughout their omnibus sanctions motion, News Plaintiffs conflate different discovery disputes in an effort to piece together a false narrative about OpenAI's purported failure to satisfy its discovery obligations. In reality, however, News Plaintiffs' allegations concern three separate issues: (1) OpenAI's preservation of chat logs generated by user conversations with OpenAI's artificial-intelligence products; (2) OpenAI's production of training datasets used to train its products; and (3) OpenAI's representations to the Court about how technically difficult it is to search its training datasets and chat logs. OpenAI takes each in turn.

A. OpenAI Retained 50 Billion User Chats and Made Them Available to Plaintiffs for Sampling

OpenAI has persistently complied with its preservation obligations. Throughout this litigation, OpenAI has repeatedly informed the Court and Plaintiffs about its pre-existing practices governing when and how it retains logs recording user conversations with its artificial-intelligence products, including ChatGPT. *See, e.g.*, MDL ECF 91-1, at 4-7. As Magistrate Judge Wang has recognized, those practices explain why retaining all conversation logs "wholesale" and stopping "all … deletion processes" is technologically infeasible and would require much more than a "flip of the switch." Pls.' Ex. 31 (Hr'g Tr. 34:17-20 (May 27, 2025)). OpenAI did not violate the Court's Preservation Order or any other discovery rule by maintaining its pre-existing policies. And notwithstanding those policies, OpenAI has worked tirelessly with Plaintiffs and the Court to provide Plaintiffs with the discovery they sought. OpenAI retained a warehouse of more than *50 billion* chat logs and made a representative sample of *20 million* logs available to Plaintiffs. OpenAI also conducted a court-ordered sampling exercise definitively proving that the exclusion of user-deleted and temporary chats from the Merits Sample did not bias the sample in any statistically significant way.

1. OpenAI Maintains Robust Data Retention Policies and Procedures

Hundreds of millions of people globally use OpenAI's products every day for countless tasks, many of which involve sharing deeply private information. This heavy use necessitates vast amounts of storage: As of January 2025, OpenAI's data warehouse contained "more than 50 billion ChatGPT user conversations," which, even compressed, took up [redacted] bytes of storage. Ex. 4 (NYT ECF 425, Trinh Decl. ¶ 4). Despite the volume, OpenAI goes to great lengths to protect its users' data and privacy. OpenAI's "data protection policies and procedures … are absolutely central to the trust [its] customers and users place in [its] services." Ex. 1 (MDL ECF 67, Kechida Decl. ¶ 3). OpenAI implemented data-retention policies before this litigation started, applying slightly different measures to its consumer-facing ChatGPT products than to its business-facing API.

ChatGPT Consumer. By default, OpenAI "indefinitely" stores text "prompts and outputs" generated by users' conversations with its consumer products (ChatGPT Free, Plus, and Pro). Ex. 2 (MDL ECF 64-4, Monaco Decl. ¶¶ 2-3). But users may "opt out" of data retention in one of two ways. *First*, users can manually delete a specific conversation, their entire chat history, or their account. *Id.* ¶ 3. Once a user has marked a conversation for deletion, OpenAI will delete the users' data "within 30 days." *Id.* *Second*, a user can use the "Temporary Chat" function. *Id.* Temporary chats do not appear in a user's chat history and, on the back-end, OpenAI does not create a permanent log of the chat. *Id.* Users delete about [redacted] and use the temporary chat function for [redacted]. Ex. 3 (MDL ECF 598, Cohen Decl. ¶ 6).



Id.

Even where a user has opted out of data retention by using the temporary chat function, OpenAI temporarily stores partial copies of all conversation data "in [an] uncompressed format" for up to 30 days, enabling OpenAI to review the data for safety reasons. Ex. 2 (MDL ECF 64-4, Monaco Decl. ¶¶ 3-4). After 30 days, this data is no longer stored. *Id.* ¶ 3. OpenAI also creates a separate copy of conversation data for long-term storage in a data warehouse. These copies are much more detailed than the temporary 30-day records and contain, "at a minimum, the text prompt, output, and metadata related to the user's conversation." Pls.' Ex. 32 (MDL ECF 435-1, Monaco Decl. ¶ 2). Consequently, they are also much larger than the temporary 30-day records. *See id.*; Pls.' Ex. 17. Because storing so much information in an uncompressed format would be technically difficult and expensive, OpenAI compresses the maintained data in the data warehouse. Ex. 4 (NYT ECF 425, Trinh Decl. ¶ 4). When Plaintiffs (at 13) claim that OpenAI has "compressed tens of billions of consumer ChatGPT logs … render[ing] them unsearchable" on "24 separate occasions," they are referring to OpenAI's ordinary practice of compressing data for long-term storage.

The chat records in long-term storage "cannot be searched without first decompressing the data." Ex. 2 (MDL ECF 64-4, Monaco Decl. ¶ 4). OpenAI can, however, locate chat records based on "user and date" information—permitting OpenAI to find specific posts. Ex. 4 (NYT ECF 425, Trinh Decl. ¶ 9).

OpenAI's API. This business product allows developers to integrate OpenAI's models into their own software. *See* API Platform, OpenAI, https://openai.com/api. Because many of these use-cases implicate commercially sensitive information, OpenAI gives business customers control over the retention of their records—telling them that "[y]our data is your data." Data Controls in the OpenAI Platform, OpenAI, https://developers.openai.com/api/docs/guides/your-data. Unless a customer contracts for zero-day retention, OpenAI keeps API data for 30 days to allow for "abuse monitoring," after which OpenAI deletes the data. *Id.* Accordingly, OpenAI never creates a copy of API conversations in its compressed data warehouse.

Strict adherence to these published data-retention policies is essential to keep OpenAI in compliance with various data-retention laws in jurisdictions around the world. If OpenAI were to "preserve user data regardless of OpenAI's policies and agreements," that step would "implicate the laws and regulations of the 160+ countries where OpenAI's users reside." Ex. 1 (MDL ECF 67, Kechida Decl. ¶ 4).

2. OpenAI Informs Plaintiffs and the Court of Its Preservation Policies Early in the Litigation

In response to a February 23, 2024 New York Times preservation letter, OpenAI informed The Times that OpenAI "ha[d] taken and will continue to take reasonable measures to preserve all potentially relevant evidence." Pls.' Ex. 61 (NYT ECF 379-3, at 1). But OpenAI also explained that, because of "privacy and contractual reasons," it would not go further than the Federal Rules require and place all input and output logs on "indefinite retention." *Id.* at 1-2. In particular,

OpenAI emphasized that it would not retain data if "users have affirmatively turned off chat history, deleted their past conversations, or deleted their account." *Id.* at 1. Plaintiffs did not respond to that letter or otherwise express disagreement.[1]

Ten months later, however, News Plaintiffs for the first time took the position that the Federal Rules required OpenAI to "ensure the preservation of *all* ChatGPT user conversations, notwithstanding the retention period prescribed by OpenAI's data retention policy." Pls.' Ex. 54 (NYT ECF 379-14, at 1).

When OpenAI did not agree to that demand, Plaintiffs filed a letter motion asking the Court to "order OpenAI to preserve output log data on a going-forward basis." NYT ECF 379, at 1. OpenAI opposed, explaining that retaining *all* chat records that users had marked for deletion would "harm [its] reputation and customer relationships" and would be technologically impracticable, "requir[ing] a new custom-designed system" that would take "at least two calendar quarters of dedicated full-time engineer time to build." NYT ECF 423, at 1-2.

On January 24, 2025, the Court denied Plaintiffs' request to compel "wholesale preservation of user conversations" and directed the parties to meet and confer on the "privacy and technological considerations" OpenAI had raised.[2] NYT ECF 441, at 2.

[1] News Plaintiffs (at App'x 14) suggest they responded on March 5. But the email they cite replied to a *different* letter that OpenAI sent The Times demanding that *The Times* preserve evidence regarding its use of OpenAI's products. *See* Ex. 5.

[2] The third-party privacy concerns raised by OpenAI are not speculative or hypothetical. After the Preservation Order, several third parties wrote the Court to directly express their interests in privacy protection. *See, e.g.*, MDL ECF 49, 95 & 484.

3. The Court Enters a Preservation Order But Does Not Require "Wholesale" Preservation

Shortly after the cases were consolidated in this MDL on April 30, 2025, Plaintiffs filed a May 7 letter that purportedly identified several "spikes" in the deletion data OpenAI had voluntarily disclosed, and claimed that these "spikes" corresponded to intentional deletions of user data by OpenAI for litigation-related reasons. Ex. 6 (MDL ECF 21-1, Letter to Magistrate Judge Wang at 2). Those accusations were unequivocally false, as OpenAI established in a sworn (and unrebutted) witness declaration. Ex. 2 (MDL ECF 64-4, Monaco Decl. ¶¶ 10-14) (explaining that these "spikes" occurred due to infrastructure issues). But before OpenAI had a chance to provide that explanation, Magistrate Judge Wang issued a "Preservation Order," directing OpenAI to take steps to "preserve and segregate all output log data that would otherwise be deleted on a going forward basis until further order of the Court," although it did not (as the Court later clarified) require wholesale preservation of all chats. MDL ECF 33, at 2 (emphasis omitted).

OpenAI moved for reconsideration, emphasizing that (1) OpenAI had already informed counsel and the Court about its data-retention practices from the outset of the litigation; (2) "OpenAI ha[d] not deleted user conversation data for any reason other than [those] practices"; and (3) preserving additional user chats was unnecessary because OpenAI had already retained *billions* of user chats. MDL ECF 40, at 1. OpenAI also detailed that, although it was "taking the steps it can to comply" with the Preservation Order and "retain[] data … otherwise slated for deletion," taking additional steps would require "significant engineering work, infrastructure changes, and compute resources." *Id.* at 1-2.

On May 16, 2025, Magistrate Judge Wang denied OpenAI's motion "without prejudice to renewal." MDL ECF 42, at 1. Judge Wang surmised that if people use ChatGPT for infringing purposes—*e.g.*, to "get around [The Times's] pay wall"—they might be more likely to "delete all

[their] searches" to cover their tracks. *Id.* at 4. In other words, Judge Wang observed that user-initiated deletions might mean that the billions of retained chat records are not representative of *all* user chats and thus could not be used to generate a representative sample. Accordingly, Judge Wang requested "proposals" for "a means to determine whether the retained output logs marked for deletion differ in any material way from the retained logs." *Id.* at 5. Judge Wang also acknowledged the "technological challenges of preservation and segregation," encouraged the parties to continue meeting and conferring, and suggested that OpenAI could renew its objections to the Preservation Order ahead of a May 27 status conference. *Id*. at 4-5; *see* MDL ECF 64-1 (renewing objections).

At the May 27 conference, OpenAI clearly explained the steps it had taken to comply with the Order and that due to technical challenges the preserved tables are not "a full representation of everything that is stored in the long-term," nor are they "designed to be the equivalent of [the] data warehouse." Pls.' Ex. 31 (Hr'g Tr. 36:12-15). Judge Wang declined to order OpenAI to take further steps to comply with the Preservation Order, and expressly clarified that the Order did not require the "wholesale" preservation of user chats: "I'm not saying stop all your deletion processes. Stop in any way. I am saying … [h]ow do we get our samples? How do we test and compare? And you know what, if we do the testing and the comparing, there's no difference between what's marked for deletion or what typically would be deleted versus what's normally retained, then maybe we have an answer there, and then maybe we move on." *Id.* at 34:19-35:2. On that basis, Judge Wang kept the Preservation Order in place, but expressed a desire "not to keep [it] going longer" and recognized that OpenAI should "do what [it] can to preserve" data given the "burden." *Id.* at 51:14-24.

Based on Judge Wang's clarification, on May 29, 2025, OpenAI notified the Court that it would "leave in-place the mechanisms it has implemented to comply with the Preservation Order, which (OpenAI understands) the Court found to be sufficient" but would not take further steps "[u]nless otherwise directed." Ex. 14 (MDL ECF 77, at 2). Those existing mechanisms included "provisionally retaining in a segregated legal hold database Conversation Data from 30 days prior to the Preservation Order, regardless of whether that data is subject to a user-initiated deletion or was generated via ChatGPT's Temporary Chat feature," but did not include "disregard[ing] user requests to delete ChatGPT Consumer Conversation Data that is more than 30 days old." Ex. 2 (MDL ECF 64-4, Monaco Decl. ¶¶ 7-8).

On that express understanding of the Preservation Order (which neither Plaintiffs nor the Court questioned), OpenAI began retaining and segregating *all* conversations stored in OpenAI's 30-day tables—*i.e.*, the data that OpenAI would ordinarily retain for just 30 days before deleting.[3] Given the burden involved, OpenAI could not, however, change the data-retention rules governing its compressed data warehouse of retained chats. Users could thus still delete chats that were more than 30 days old as of the date of the Preservation Order, resulting in the removal of the underlying data from the warehouse. OpenAI kept the Court and Plaintiffs' counsel abreast of its preservation practices in real time during the confidential settlement conferences the Court convened in June, August, and September 2025. *See* MDL ECF 217, 444, & 592.[4]

[3] The retained data includes both consumer and API chats. So it is simply not true, as Plaintiffs (at 13) allege, that OpenAI "destroyed virtually all of its API logs." Instead, OpenAI initially did not retain those chats for more than 30 days. Then, after the Court entered its Preservation Order, OpenAI began to preserve newly generated API chats.

[4] Mindful of the Court's admonition to not disclose "protected settlement communications" that took place during these confidential and off-the-record "settlement conferences," OpenAI will not disclose further details regarding these conferences here or elsewhere in its brief. *See* MDL ECF 708. However, they are an important part of the factual record, including because they informed OpenAI's own understanding of what its obligations were and whether it was complying with

4. A Court-Ordered Sampling Procedure Confirms that OpenAI's Retained Chats are Representative of How Users Utilize Its Products

In addition to clarifying the scope of the Preservation Order at the May 27 hearing, Magistrate Judge Wang also approved a sampling mechanism OpenAI had proposed for determining whether user-initiated deletions and temporary chats render the population of retained chats unrepresentative. Under that proposal, OpenAI would use the preserved-data tables to create two random samples of ChatGPT conversation data, one containing conversations that OpenAI ordinarily retains, and another containing conversations that OpenAI ordinarily retains plus conversations that OpenAI would ordinarily delete under its policies. MDL ECF 64-1, at 4-5. Those two samples, OpenAI explained, would allow the parties and the Court to assess whether excluding user-deleted chats and temporary chats from the population of retained chats would bias an analysis of outputs based only on the retained chats. *Id.* at 4. Judge Wang adopted that proposal. *See* Pls.' Ex. 31 (Hr'g Tr. 45:12-16); MDL ECF 79 (written order).

OpenAI then spent months working with Plaintiffs and the Court on a data-sampling protocol. That process culminated in an August 19, 2025 Order adopting an agreed-upon "random sampling methodology." MDL ECF 463. OpenAI generated the two datasets it had initially proposed, along with two additional datasets that Plaintiffs had demanded: one consisting entirely of temporary chats and another consisting entirely of chats subject to user-initiated deletions. *Id.* at 1. Each dataset contained five million "rows of Conversation data." *Id.* The parties then analyzed the "hit count" of key terms across each of those samples. *Id.* at 2.

---

them, as well as what Plaintiffs and the Court understood about OpenAI's actions. Additionally, OpenAI wishes to contextualize Plaintiffs' own descriptions of these conferences. *See* Mot. App'x A at 27-29.

The sampling exercise revealed *essentially no statistically significant difference* between the data OpenAI retains and the data OpenAI would have retained had it paused all deletions. The "hit counts prove that the exclusion of Temporary and Deleted Chats … have no discernable effect" on the representativeness of a sample drawn from OpenAI's retained data. MDL ECF 597, at 1 (joint letter on sampling). If users disproportionately used temporary chats or deletions to hide infringing uses of ChatGPT, then "that bias would have shown up in the form of higher hit counts for the relevant terms" in the sample including *all* conversation data than in the sample including just retained data. *Id.* at 2. But "[t]he hit counts show the opposite." *Id.* Notably, Plaintiffs' expert did not disagree with this finding, instead noting a higher incidence of potentially relevant content in temporary chats—a difference which did not produce sampling bias. *See infra*, p.27.

Because all parties agreed that sampling the logs preserved under the Preservation Order gave the parties a way to measure whether the retained logs differed from user-deleted logs or temporary chats, the parties later stipulated to the termination of "OpenAI's ongoing obligations" under the Preservation Order, effective September 26, 2025. Ex. 15 (MDL ECF 629).

5. OpenAI Creates and Produces the Merits Sample, A Representative Sample of 20 Million Output Logs

Alongside OpenAI's efforts to help News Plaintiffs assess temporary chats and chats deleted by users, OpenAI also worked with Plaintiffs to produce the Merits Sample—a representative sample of 20 million retained conversations that the parties could use to draw statistical conclusions about how often Plaintiffs' works appear in the full universe of ChatGPT outputs during a specified timeframe.

The path to the Merits Sample began when News Plaintiffs initially requested production of *all* "chat logs related to [their Content]." Pls.' Ex. 14 (MDL ECF 864-1, at 2). But OpenAI could not (and cannot) readily search billions of compressed chat logs stored in OpenAI's data

warehouse. Ex. 2 (MDL ECF 64-4, Monaco Decl. ¶ 5). Compressed chat logs, OpenAI explained, cannot be searched for keywords—*e.g.*, user mentions of "The New York Times"—until they are decompressed.[5] *Id.* ¶ 4. And because it would be infeasible to decompress 50 billion or more chat logs, the best way to search the logs would be to create a representative sample. *See id.* ¶¶ 4-5; *see also* Ex. 7 (MDL ECF 864-11, at 2).

News Plaintiffs ultimately agreed and proposed a sampling methodology for more than 1.4 billion ChatGPT conversation logs. Pls.' Ex. 58 (MDL ECF 864-12, at 3-4). OpenAI offered 20 million conversation logs as a compromise and, on August 11, 2025, Plaintiffs furnished OpenAI with a list of nearly 20 million conversation IDs corresponding to ChatGPT chat logs to use as the sample for searching records related to News Plaintiffs' content. *See* Pls.' Ex. 63 (MDL ECF 719-2); MDL ECF 717 at 3. OpenAI then began the extensive and complex process of retrieving, decompressing, processing, and anonymizing the logs corresponding to those 20 million conversation IDs. *See* Pls.' Ex. 32 (MDL ECF 435-1, Monaco Decl. ¶ 3) (describing this process); *see also* MDL ECF 717, at 3-4 (providing timeline).

OpenAI was ultimately able to retrieve and decompress "95%" of the 20 million chat logs Plaintiffs had selected. Ex. 8 (MDL ECF 656-2, at 1). OpenAI, however, "need[ed] to pull replacement logs for approximately 5% of the logs on [Plaintiffs'] list," *id.*, because those logs "were the subject of manual deletion requests by ChatGPT users" in the several months between when OpenAI sent Plaintiffs an initial list of conversation IDs and when Plaintiffs selected conversation IDs for sampling in August 2025, *see* Ex. 9 (MDL ECF 683, Monaco Decl. ¶ 4). As OpenAI had previously told the Court, it did not understand the Court's Preservation Order to

[5] Even transiently decompressing each of the 50+ billion compressed output logs for search purposes would be an extreme undertaking. *Contra* Mot. 12.

require "disregard[ing] user requests to delete ChatGPT Consumer Conversation Data that is more than 30 days old." Ex. 2 (MDL ECF 64-4, Monaco Decl. ¶¶ 7-8); Ex. 14 (MDL ECF 77, at 1). News Plaintiffs elected to have OpenAI select replacement logs "at random" rather than undertake to select replacement logs themselves. *See* Ex. 10 (MDL ECF 681-1, at 2).

With the sampling process complete, the parties disputed whether OpenAI should produce the *entire* 20 million conversation sample or instead search the sample and produce just those conversations implicating News Plaintiffs' works, especially given the "privacy concerns" at issue. MDL ECF 717 at 4-5. On November 7, 2025, Magistrate Judge Wang ordered OpenAI to produce all 20 million "anonymized … output logs," MDL ECF 734, and Judge Stein overruled OpenAI's Rule 72 objections, MDL ECF 1021. In doing so, Judge Stein anticipated that OpenAI would use a "custom de-identification tool," to "remov[e] both personally identifiable information and other private information from the logs" before producing them. *Id.* at 1-2.

As ordered, OpenAI produced the Merits Sample on December 14, 2025, and de-identified the logs using an "automated system" that was "designed to look for information that may personally identify a user and remove it." MDL ECF 1126, at 2; Ex. 11 (MDL ECF 1128, Monaco Decl. ¶ 3). Despite OpenAI's efforts, News Plaintiffs filed a letter motion objecting to the extent of OpenAI's redactions. *See* MDL ECF 1061, at 1-2. After a hearing, Judge Wang ordered OpenAI to reduce the number of redactions "and re-produce the sample by January 22, 2026." MDL ECF 1170. OpenAI did so.

B. OpenAI Makes its Training Data Available for Inspection at OpenAI's Expense

In addition to inspecting OpenAI's output logs for evidence of infringement, News Plaintiffs also requested access to the data OpenAI uses to train its models. *See* Pls.' Ex. 4, at 8. To allow News Plaintiffs to search OpenAI's training data, the parties extensively negotiated an

inspection protocol that culminated in a controlled "sandbox" environment. *See, e.g.*, NYT ECF 250, at 6-7. The sandbox was essential because OpenAI could not feasibly produce its proprietary and commercially sensitive training data outright. The datasets are so large (hundreds of terabytes) that they would be extremely difficult for OpenAI to transfer and for Plaintiffs to manage. Nor did OpenAI have the "tools to search" its "extremely voluminous datasets" for all the information Plaintiffs identified as potentially relevant. Pls.' Ex. 1 (Hr'g Tr. 7:17-20 (Sept. 12, 2024)).

Under the protocol, OpenAI made its training data available for inspection on a secure computer located at its counsel's offices—hosting the data and supplying the compute at its own expense. NYT ECF 254 ¶¶ 3, 7. OpenAI went to extreme lengths to ensure that the protocol worked for Plaintiffs. When Plaintiffs ran "flawed code that overwhelmed and crashed the file system," OpenAI "invested significant resources, including days of engineer time as well as hardware and compute resources to get Plaintiffs back on track." NYT ECF 345, at 1-2. OpenAI even offered to run Plaintiffs' searches for them and "worked diligently with its engineering team" to do so. *Id.* at 2. Although Plaintiffs (at 8) now suggest that OpenAI "erased" their "search result data stored on one of the … machines," Plaintiffs omit the fact that the purported data loss happened due to a *configuration change that Plaintiffs themselves requested*. *Id.* Nor did the configuration change delete any files—at most it erased "organizational metadata." *Id.*

In the end, the inspection protocol was a success: it enabled Plaintiffs to search OpenAI's training data for their works. That is why Plaintiffs have not pushed for additional searches of OpenAI training data in months.

C. OpenAI Has Not Misrepresented Its Discovery Capabilities to the Court

Plaintiffs allege (at 5-6) that OpenAI misrepresented the difficulty of searching for Plaintiffs' works in OpenAI's training data and output logs. In making this argument, Plaintiffs mischaracterize Mr. Monaco's deposition testimony and conflate OpenAI's anti-regurgitation

analyses with the systematic searches that litigation discovery demanded. The data OpenAI used for its anti-regurgitation analyses are entirely different from the data Plaintiffs requested. To perform its anti-regurgitation analysis, OpenAI compared an unrepresentative sample of 78 million conversation logs against a single training set. By contrast, News Plaintiffs' discovery requests sought historic, detailed conversation logs and associated metadata from OpenAI's compressed data warehouse and all its training sets. These searches differ enormously in size and in the resources required. Nothing in Mr. Monaco's testimony contradicts OpenAI's representations that the methodical and careful analysis of training data and output logs that Plaintiffs requested in discovery is a technically difficult and expensive endeavor.

1. Project Giraffe and Anti-Regurgitation Efforts

OpenAI has engaged in anti-regurgitation efforts since 2022 to prevent its products from reproducing copyrighted works to users under the label "Project Giraffe." One of those efforts was a "Bloom filter." *See* Pls.' Ex. 25 (Monaco Dep. 46:4-24 (Jan. 27, 2026)). The Bloom filter works by taking continuous sequences of words (called "n-grams") found in certain copyrighted works and creating a "block list," which then prevents OpenAI's products from repeating those same sequences of words to users—unless the sequence also appears on an "allow list" of commonly used n-grams (*e.g.*, a quote from the Bible or an everyday idiom). In late 2023 and early 2024, OpenAI added certain content from Plaintiffs' websites to the Bloom filter to prevent its products from reproducing those words. *See* Pls.' Ex. 2 (Monaco Dep. 229:15-17 (Apr. 8, 2026)).

Before adding a new dataset of additional n-grams to the Bloom filter, OpenAI ran "evals" to ensure that it had "an accurate view … of how the filter's going to work in production." *See id.* 236:12-22. Those evaluations would determine how often n-grams included in a particular dataset "intersected" with a sample of ChatGPT output logs. *Id.* at 237:17-25. The result would be "a

file … containing the aggregate statistics as well as, typically, a sample of the n-grams" that appeared in the ChatGPT outputs. *Id.* In total, OpenAI has run "several dozen evals like this." *Id.* at 236:12-15.

In addition to Project Giraffe, OpenAI also conducted an evaluation of possible regurgitation of New York Times content. That analysis, codenamed "NYT Recall Pipeline," began in May 2023 and evaluated whether New York Times content in OpenAI's training data also appeared in output logs of ChatGPT user conversations. Pls.' Ex. 2 (Monaco Dep. 241:5-242:25). To locate New York Times training data, OpenAI searched one of its many training datasets—derived from "Common Crawl," a database of websites scraped from the internet—for New York Times URLs. *See id.* at 222:18-23. From the resulting set of New York Times content, OpenAI then generated a list of "n-grams." *Id.* at 232:23-25. OpenAI then searched for those n-grams across a rough-and-ready sample of 78 million conversations. Ex. 19 (Kellogg Decl. ¶ 3). Searching that limited sample differs fundamentally from searching the Merits Sample, which contains a representative sample set of conversation outputs—assembled from OpenAI's compressed data warehouse—with complete metadata. *Id.* ¶ 4.

2. Project Giraffe is Consistent with OpenAI's Representations to the Court

The above background on Project Giraffe and the NYT Recall Pipeline makes abundantly clear that OpenAI did not misrepresent its ability to search training data or output logs.

a. Output Logs. Plaintiffs (at 1, 9-12, 15, 17) repeatedly criticize OpenAI for telling the Court that it could not readily search its *full universe* of conversation data for specific search terms. That was true then, and it remains true now. As OpenAI explained, any searches of all historic, compressed conversation data would first require decompressing a representative sample of that data. *See, e.g.*, Pls.' Ex. 16 (Hr'g Tr. 284:8-15 (June 26, 2025)).

Nor is there any merit to Plaintiffs' suggestion (at 9) that OpenAI "forc[ed]" them to accept the Merits Sample by failing to disclose the existence of the 78 million output logs at the outset of litigation. For one, OpenAI never concealed the existence of the 78 million logs. *See infra*, pp.21-22. For another, the 78 million logs do not constitute a representative sample that would support statistically significant conclusions about the contents of general ChatGPT outputs. *See* Ex. 19 (Kellogg Decl. ¶ 4).

Two key differences between the 78 million set and the Merits Sample prove the point. *First*, the 78 million logs do not randomly sample outputs across a given period. Instead, OpenAI drew sets of "5 million" chat logs from various "weekly" sets of ChatGPT data, which were smaller files without complete metadata. Pls.' Ex. 2 (Monaco Dep. 246:22-25). The Merits Sample, by contrast, includes a cross-section of logs drawn from compressed logs stored in OpenAI's data warehouse (*i.e.*, long-term storage) spanning from December 2022 through November 2024, which (all agree) is the relevant time period. Ex. 19 (Kellogg Decl. ¶ 5). *Second*, the 78 million sample lacked complete metadata. While the Merits Sample included *all* metadata associated with each log—metadata required to, for example, analyze back-end processes like retrieval augmented generation used to generate the response—the logs in the sample used for the NYT Recall Pipeline had no metadata other than a date and model name. *See supra*, p.18. A single log from the Merits Sample has so many fields and data that each log would be more than 300 pages long if printed. Ex. 19 (Kellogg Decl. ¶ 4). By contrast, each record in the 78 million sample would barely fill one printed page. *Id.* In other words, the Merits Sample is *actually* 75 times larger by total data volume than the 78 million sample. *Id.*

Those differences made producing the Merits Sample for this litigation a significant technical challenge. The compressed logs "contain, on average, at least 128,000 words, much of

which consists of metadata," and each log needed to be retrieved, decompressed, processed, and stored. Pls.' Ex. 32 (MDL ECF 435-1, Monaco Decl. ¶¶ 2-3). Handling the logs for the Merits Sample thus required significant "computational, data engineering, system architecture, and software engineering resources," just as OpenAI represented. Pls.' Ex. 60 at 2.

b. Training Data. News Plaintiffs (at 1) fault OpenAI for telling the Court that it did "not currently have tools to search [all of its training datasets] efficiently" and that searching for Plaintiffs' works would take "a very long time." Pls.' Ex. 1 (Hr'g Tr. 7:17-18). That was and remains true.

When OpenAI made that representation, OpenAI was responding to a request from a different set of Plaintiffs—the Books Plaintiffs—for OpenAI to determine whether *any* content in *any* of 300-plus books appeared in *any* of OpenAI's training datasets. *See id.* at 14:17-19. As OpenAI explained, although it could search training data for "particular phrases," it could not search for entire books—especially "without some input from plaintiffs" about "what … [to] search for." *Id.* at 8:3-6.

The training-data searches OpenAI ran under the NYT Recall Pipeline or Project Giraffe are consistent with that representation. The NYT Recall Pipeline, for instance, involved searching just one training dataset (Common Crawl) for just one thing (New York Times URLs). Likewise, although OpenAI located certain word sequences in its training data to add to the Project Giraffe filter using "tools," "heuristics," and "techniques" that the company was "continuously developing," Pls.' Ex. 2 (Monaco Dep. 223:14-21), none of those tools enabled OpenAI to search all of its datasets for vast amounts of material.

In other words, just as OpenAI told the Court, OpenAI's existing ability to "search for particular phrases of things," did not allow it to perform all the searches Plaintiffs desired—

instead, "[P]laintiffs would have to develop their own tool to do the search." Pls.' Ex. 1 (Hr'g Tr. 7:21-24, 8:3-6). Regardless, OpenAI expressly disclosed to the Court that it had *some* existing capacity to search its training data, *id.* at 8:3-6, and even offered to run Plaintiffs' searches for them using that method, *see, e.g.*, NYT ECF 462, at 2. Underscoring that OpenAI did not hide the ball, OpenAI ultimately ran the *same* search of Common Crawl for News Plaintiffs that OpenAI had used for the NYT Recall Pipeline. *See id.*

3. OpenAI has Produced Relevant Documents and Datasets Pertaining to Its Anti-Regurgitation Efforts

News Plaintiffs (at 6, 20-21) also criticize OpenAI for purportedly obstructing discovery into its anti-regurgitation efforts. But what Plaintiffs characterize as obstruction is simply routine discovery practice in a heavily contested litigation. And Plaintiffs were not prejudiced by the supposed discovery issues in any case.

OpenAI produced the source code for its Project Giraffe evaluations in February 2025, disclosing to Plaintiffs the existence of several underlying datasets. *See* Pls.' Ex. 11. And when Plaintiffs first moved to compel production of the 78 million sample in June 2025, they did so in a wide-ranging motion requesting the production of "over 20 billion files containing over 320,000,000 gigabytes [of data] from an array of confidential projects," only one of which constituted the 78 million sample. *See* MDL ECF 201-13, at 3. Only when Magistrate Judge Wang ordered the parties to meet and confer over the motion did Plaintiffs explain that they wished to access the "Giraffe ChatGPT Dataset" because they believed it "contains a regurgitation analysis of actual ChatGPT user conversations." Pls.' Ex. 20 at 3. OpenAI *accurately* clarified that the dataset—which, again, comprises a sample of user chats—did not "contai[n] a regurgitation analysis." *Id.* At no point did OpenAI hide the fact that the dataset contained "user data"—a fact

that OpenAI clearly communicated to News Plaintiffs in the meet-and-confer process, and a fact that was obvious given it was a dataset used for evaluations. *See* MDL ECF 201-13, at 3.

News Plaintiffs then let the issue drop for months, only reviving their interest in Project Giraffe days before the close of fact discovery on January 30, 2026. *See* MDL ECF 1220. Plaintiffs renewed their request after Mr. Monaco's first deposition, suggesting (without foundation) that OpenAI had "withheld" relevant information about its anti-regurgitation efforts. *Id.* at 2. In reality, as OpenAI explained, it had already produced "all source code relating to Project Giraffe," all of "the metrics [OpenAI] maintains relating to regurgitation," "analyses and communications relating to Project Giraffe," and a spreadsheet with the results of the NYT Recall Pipeline. MDL ECF 1267, at 2; *see also* Pls.' Ex. 24 (Hr'g Tr. 11:19-22 (Feb. 11, 2026)) (explaining that "Project Giraffe" had been a discovery search term "since 2024" and that "Plaintiffs have known what Project Giraffe is").

Even so, the Court ordered OpenAI to produce the 78 million sample used for the NYT Recall Pipeline plus a further 9 million chats used for another Project Giraffe evaluation and granted Plaintiffs additional deposition time to re-examine Mr. Monaco on the subject. *See* MDL ECF 1419.[6] News Plaintiffs thus have had *extensive* discovery into OpenAI's anti-regurgitation efforts and even now fail to identify what additional information OpenAI should have disclosed.

ARGUMENT

Plaintiffs' omnibus motion for sanctions fails at every turn because it rests on a fundamental mischaracterization of both the facts and the law. Plaintiffs try to weave together discrete discovery disputes to support their allegations of wrongdoing. But there is no wrongdoing

[6] The subsequently produced 9-million log set is similarly unrepresentative because it was specifically constructed to overrepresent conversations featuring RAG data. Ex. 13 (Supp. Berg-Kirkpatrick Rep. ¶¶ 10, 16-20).

here. This Court has already resolved disputes over the preservation of user chats, the assembly of the Merits Sample, and the production of Project Giraffe files through tailored discovery orders. Discovery is long over, and Plaintiffs are entitled to nothing more.

When examined against the applicable legal standards, none of Plaintiffs' five categories of requested relief withstand scrutiny. OpenAI did not willfully violate the Court's Preservation Order, and preclusion of the Merits Sample is unwarranted under either Rule 37(b) or Rule 37(e)(1). Plaintiffs identify no textual predicate for established fact sanctions under Rule 37 and, even if such a predicate existed, they cannot satisfy the controlling framework from *Residential Funding Corp. v. DeGeorge Financial Corp.*, 306 F.3d 99, 107 (2d Cir. 2002), *superseded by rule in other part as recognized in Hoffer*, 128 F.4th 433. Rule 37(e)(2)'s exacting intent standard—which requires proof of "*intentional* loss or destruction of evidence," *Hoffer*, 128 F.4th at 438—is not remotely satisfied here. Plaintiffs' fee requests fail because OpenAI's conduct was substantially justified at every stage of discovery. And this Court's inherent authority provides no independent basis for sanctions where, as here, Plaintiffs cannot demonstrate bad faith or conduct "entirely without color." *Wolters Kluwer Fin. Servs., Inc. v. Scivantage*, 564 F.3d 110, 114 (2d Cir. 2009).

At bottom, what Plaintiffs mischaracterize as a pattern of obstruction is in fact the predictable complexity of managing unprecedented discovery into tens of billions of compressed conversation records and hundreds of terabytes of training data—implicating the privacy interests of hundreds of millions of users in over 160 countries. OpenAI has consistently engaged in good faith with Plaintiffs and this Court to navigate those challenges, ultimately giving Plaintiffs everything they could plausibly require and more.

I. Plaintiffs Fail to Establish a Basis for Preclusion of the Merits Sample Under Rule 37

Plaintiffs advance two bases for precluding OpenAI from relying on the Merits Sample: (1) Rule 37(b), based on alleged violations of the Court's orders; and (2) Rule 37(e)(1), based on OpenAI's alleged failure to preserve output logs. Neither ground has merit.

A. Rule 37(b) Does Not Support Preclusion

1. Rule 37(b)(2)(A) authorizes sanctions only when a party "fails to obey an order to provide or permit discovery." Here, Plaintiffs (at 27) argue that OpenAI (1) violated the Preservation Order, MDL ECF 33, by failing to stop users from deleting historic chat logs and by substituting 5% of the logs in the sample, and (2) violated a November 2025 order requiring OpenAI to produce the Merits Sample, MDL ECF 734, by over-redacting the remaining entries. OpenAI did not disobey those orders and each of the Second Circuit's Rule 37 factors cut against the imposition of sanctions. Those factors are: "(1) the willfulness of the non-compliant party or the reason for noncompliance; (2) the efficacy of lesser sanctions; (3) the duration of the period of noncompliance; and (4) whether the non-compliant party had been warned of the consequences of … noncompliance." *Agiwal v. Mid Island Mortg. Corp.*, 555 F.3d 298, 302 (2d Cir. 2009) (omission in original) (citation omitted).

Compliance and Willfulness. Plaintiffs (at 27) contend that OpenAI willfully failed to comply with the Court's orders because it "thought about complying" and decided not to. That argument is meritless.

OpenAI consistently complied with its obligations under the Preservation Order. Immediately after the Court issued its May 13 Order, OpenAI took affirmative preservation steps, including retaining conversation data from the 30 days prior to the Order in a segregated legal-hold database. MDL ECF 40, at 2. Then, at a May 27, 2025 hearing, Judge Wang expressly clarified that the Preservation Order did not require OpenAI to "stop all [its] deletion processes"

or require "wholesale" preservation of all chat logs. Pls.' Ex. 31 (Hr'g Tr. 34:18-35:2). Based on that clarification, OpenAI notified the Court that it would continue to preserve any temporary chats or chats marked for deletion by users where those chats were generated *after* the Order went into effect, but that it would not take additional steps—including disregarding user requests to delete chats generated *before* the Order's effective date. *See* Ex. 14 (MDL ECF 77, at 2); Pls.' Ex. 3 (MDL ECF 64-4, Monaco Decl. ¶¶ 7-8). OpenAI ultimately preserved every consumer and API conversation subject to at least a 30-day data retention generated from April 14, 2025, until OpenAI's obligations under the Order terminated on September 26, 2025.

News Plaintiffs ignore OpenAI's compliance with the Order, instead suggesting that OpenAI simply "decided not to" obey because doing so would be "too complex." Mot. 27 (citation omitted). Again, that is not what happened. At the May 27 hearing, the Court clarified that the Order *did not* require OpenAI to stop users from deleting historic chats precisely *because* OpenAI had already explained the extreme technical challenges involved. *See* Pls.' Ex. 3 (MDL ECF 64-4, Monaco Decl. ¶ 7); Pls.' Ex. 31 (Hr'g Tr. 33:12-35:2). And at the confidential settlement conferences in June, August, and September 2025, OpenAI was forthright with the Court and Plaintiffs' counsel regarding the extent of OpenAI's preservation practices following the Preservation Order. *See* MDL ECF 217, 444, & 592. These conferences reinforced OpenAI's understanding that it was acting in compliance with the Preservation Order.

Given this history, OpenAI did not "violate[] the Court's [Preservation Order] by deleting 5% of the output logs called for by the [Merits Sample] protocol." *Contra* Mot. 27. As explained above, *see supra* pp.14-15, these logs were manually deleted by ChatGPT users in the several months between when OpenAI sent Plaintiffs an initial list of conversation IDs and when Plaintiffs selected conversation IDs for sampling in August 2025. *See* Ex. 9 (MDL ECF 683, Monaco Decl.

¶ 4). And shortly after the Preservation Order, OpenAI informed the Court that it did not understand the Court's Preservation Order to require "disregard[ing] user requests to delete ChatGPT Consumer Conversation Data that is more than 30 days old." Pls.' Ex. 3 (MDL ECF 64-4, Monaco Decl. ¶¶ 7-8). Far from contradicting this understanding, the Court confirmed that the Preservation Order did not require OpenAI to "stop all [its] deletion processes" or require "wholesale" preservation of all chat logs. Pls.' Ex. 31 (Hr'g Tr. 34:18-35:2).

Plaintiffs (at 27) next suggest that OpenAI violated the November Order by over-redacting the sample. But the Court acknowledged that OpenAI would use a "custom de-identification tool" to "remov[e] both personally identifiable information and other private information from the logs" before producing the sample. MDL ECF 1021, at 1-2 (citation omitted). And OpenAI notified the Court and Plaintiffs prior to the initial production that it would "use best efforts" to address the privacy concerns associated with the wholesale production of user chat logs by "augmenting its de-identification process." MDL ECF 752, at 2 n.4. When the Court questioned the initial redaction volume that resulted from these additional privacy protections, OpenAI reproduced an updated version within the ordered timeframe.[7] Since then, News Plaintiffs have not identified any specific redaction that they think is overzealous.

Efficacy of Lesser Sanctions. Plaintiffs (at 27-28) argue that lesser sanctions would not suffice because OpenAI "effectively '[ran] out the discovery clock.'" This argument fails because it rests on a deeply flawed premise: that OpenAI's productions were the product of tactical delays rather than genuine complexity. Plaintiffs' claim (at 28) that OpenAI's productions are "unusable"

[7] To account for the modification of these privacy protections, the parties stipulated to a protocol allowing for additional redactions (as needed) for output logs produced outside of the secure environment (MDL ECF 1503, at 9), and OpenAI has recently utilized this protocol to protect user privacy.

is hard to comprehend. The Merits Sample has been produced. *See* Pls.' Ex. 32 (MDL ECF 435-1, Monaco Decl. ¶ 3) (describing process); *see also* MDL ECF 717, at 3 (providing timeline). The "reservoirs" of 78 million and 9 million outputs have been produced. *See supra*, p.22. Project Giraffe source code, evaluation data, and related communications have been produced. *See* MDL ECF 1267, at 2; *see also* Pls.' Ex. 24 (Hr'g Tr. 11:19-22). And Mr. Monaco sat for a supplemental deposition where he testified at length about all these subjects. *See* MDL ECF 1419. Given this extensive discovery, the sanctions Plaintiffs seek would only result in an unjustified windfall.

Nor can Plaintiffs (at 16-18, 28-29) seriously claim that greater sanctions are warranted based on the (false) notion that OpenAI "skewed" the Merits Sample. News Plaintiffs offer two theories for how OpenAI purportedly biased the sample; both fail. Plaintiffs first suggest that excluding user-deleted chats and temporary chats, including OpenAI's substitution of 5% of the chat logs Plaintiffs had initially selected, renders the sample unrepresentative. Not so. Even if ChatGPT users are slightly more likely to mention certain words (like "paywall") in temporary chats than in the retained chats writ large, as Plaintiffs (at 16) claim, excluding those chats from the sample would not cause bias because the population of temporary chats is vanishingly small ([redacted] of conversations). Ex. 3 (MDL ECF 598, Cohen Decl. ¶ 6). As such, only significant differences between temporary chats and retained chats would result in bias. *Id.* ¶¶ 6, 10-12. But that is not the case here, because there is almost no statistically significant difference when comparing samples of (1) all chats, including temporary chats and user-deleted chats; and (2) retained chats. *Id.* ¶¶ 21-32; *supra*, pp.12-13. Tellingly, Plaintiffs' expert did not attempt to rebut that conclusion, leaving it uncontested in the record. *See* Ex. 12 (MDL ECF 599, Goldstein Decl.).

News Plaintiffs' allegation (at 17-18, 33) that OpenAI skewed the sample by excluding chat logs that failed to decompress due to processing errors is likewise flawed. In Plaintiffs' view (at 18), those processing errors removed chats containing retrieval-augmented generation ("RAG") data—*i.e.*, conversations in which ChatGPT obtained content from the internet in real time to respond to user queries. That is pure speculation. As Mr. Monaco testified at his first deposition, "there were a few different ways that a row would be deemed … unprocessable"—including rows that "simply didn't decompress" or became "malformed for some reason." Pls.' Ex. 34 at 23:1-18. Nothing in Mr. Monaco's testimony suggests that the processing errors somehow disproportionately impacted RAG files—let alone files allegedly regurgitating Plaintiffs' asserted works. Regardless, the Court could redress any perceived prejudice by permitting Plaintiffs to introduce expert testimony contextualizing the Merits Sample.[8]

Because News Plaintiffs already have access to a representative sample of OpenAI's output logs, Plaintiffs' alleged claims of harm and prejudice (if any) have already been remedied through productions and testimony. At most, any sanction should be limited to remedying any demonstrated prejudice. There is no need to take the extraordinary step of precluding OpenAI from relying on the representative Merits Sample that both sides (and the Court) spent months developing.

Duration of Noncompliance. Plaintiffs (at 28) characterize the relevant period of noncompliance with the Preservation Order as spanning the full two years of discovery. That is

[8] Plaintiffs (at 18) complain that the population of RAG chats in the Merits Sample is "low" and "should have been higher." But the low percentage of RAG chats is a product of *Plaintiffs'* choice to take an equal number of chats (833,333) from each month in the sampling period, rather than taking a larger number of chats from later months when users interacted with ChatGPT much more frequently. *See* Pls.' Ex. 27, Goldstein Decl. ¶ 439. Plaintiffs' choice skewed the sample toward earlier periods, when RAG was either not available or available only to a small subset of users.

absurd, given that the Court's Preservation Order only ran from May 13, 2025, to September 26, 2025. *See* MDL ECF 33; Ex. 15 (MDL ECF 629). Regardless, the duration of a discovery dispute is not itself evidence of willfulness where, as here, the dispute concerns genuinely complex technical and privacy challenges. OpenAI's training datasets are hundreds of terabytes in size; its data warehouse contains tens of billions of compressed conversation records; and de-identification of large-scale user data to protect privacy is time consuming and technically complex. These are not manufactured obstacles; they are intrinsic features of the data systems at issue, and the two-year timeline reflects the scale of the litigation and the underlying data, not willful defiance.

Notice. Plaintiffs (at 28) argue that OpenAI is a sophisticated, well-resourced company and this Court "repeatedly alerted OpenAI that its conduct potentially implicated the imposition of sanctions." But sophistication and resources led OpenAI to seek clarification of the Order's scope, both at the May 27 hearing and in a letter filed two days later. *See* Pls.' Ex. 31 (Hr'g Tr. 34:19-35:2); Ex. 14 (MDL ECF 77, at 2). When OpenAI transparently informed the Court about its interpretation of the Order, the Court *confirmed* that the Preservation Order did *not* require OpenAI to "stop all [its] deletion processes" or require "wholesale" preservation of all chat logs. Pls.' Ex. 31 (Hr'g Tr. 34:19-35:2). That clarification supports OpenAI's good-faith understanding of its obligations. At the very least, it shows that OpenAI was not on notice that it could face sanctions for proceeding under its interpretation of the Order—an interpretation OpenAI reasonably understood this Court to have approved.

2. Preclusion of the Merits Sample is especially unwarranted. Preclusion of evidence is one of the "harshest sanctions available," and it "should be imposed only in rare situations." *Update Art, Inc. v. Modiin Publ'g, Ltd.*, 843 F.2d 67, 71 (2d Cir. 1988). "Before the extreme sanction of preclusion may be used," courts "must consider less drastic responses." *Outley v. City*

of New York, 837 F.2d 587, 591 (2d Cir. 1988). To assess the appropriateness of preclusion, courts must consider: "(1) the party's explanation for the failure to comply with the discovery order; (2) the importance of the precluded [evidence]; (3) the prejudice suffered by the opposing party as a result of having to prepare to meet the" evidence targeted for preclusion; and "(4) the possibility of a continuance." *Softel, Inc. v. Dragon Med. & Sci. Commc'ns, Inc.*, 118 F.3d 955, 961 (2d Cir. 1997). None of these factors favor preclusion.

First, as stated, OpenAI has fully complied with its discovery obligations. *See supra*, pp. 4-15. News Plaintiffs (at 28) complain that OpenAI "continu[ed] to delete output logs" under the Preservation Order, including 5% of the logs News Plaintiffs had selected for the Merits Sample. But (again) the Order *did not require OpenAI to suspend deletion of those logs* given the technical difficulties involved in stopping user-initiated deletions of historic chat logs. Pls.' Ex. 31 (Hr'g Tr. 34:19-35:2) ("I'm not saying stop all your deletion processes."). And OpenAI contemporaneously told the Court that was its understanding of the Order—at the very least giving OpenAI a reasonable explanation for its actions. *See Softel*, 118 F.3d at 961-62.

Second, the Merits Sample is critically important to both parties; indeed, it is the only representative sample of conversation logs available to examine the claims and defenses in this case. The parties and the Court spent countless hours cooperatively designing the Merits Sample—work that should not be wasted in service of the ill-conceived pursuit of severe, liability-altering sanctions.

Third, Plaintiffs will suffer no prejudice from having to grapple with a carefully curated dataset *they designed* by selecting exactly which chat logs they wished to include—especially because Court-ordered analysis of temporary chats and chats deleted by users revealed that the

exclusion of those chats from the Merits Sample caused no statistically significant bias. MDL ECF 597, at 1; *accord* MDL ECF 598, Cohen Decl.

And *fourth*, Plaintiffs cannot identify any prejudice when additional sampling was offered if Plaintiffs provided a reasonable basis for that request. Specifically, in August 2025, OpenAI offered to conduct "targeted supplemental sampling" if "News Plaintiffs … demonstrate[d] that their ability to prosecute their claims w[as] materially prejudiced absent another sample," MDL ECF 435, at 3—but Plaintiffs never took OpenAI up on that offer. In short, no basis exists for the extraordinary remedy of precluding the very evidence that the parties and the Court spent eight months designing, particularly where News Plaintiffs declined other remedies along the way.

B. Rule 37(e)(1) Does Not Support Preclusion

Plaintiffs (at 31-33) separately argue that preclusion is warranted under Rule 37(e)(1) because OpenAI allegedly failed to preserve output logs that "should have been preserved," resulting in prejudice because the deleted chats purportedly contained disproportionately high instances of copyright infringement. Rule 37(e)(1) permits sanctions only in narrow circumstances: if "a party failed to take reasonable steps to preserve" electronically stored information ("ESI") that "should have been preserved," the ESI "cannot be restored or replaced through additional discovery," and the "loss of information" causes "prejudice to another party." Fed. R. Civ. P. 37(e)(1). Even then, courts are only empowered to "order measures no greater than necessary to cure the prejudice." *Id.* Plaintiffs' preclusion argument under this rule fails for multiple independent reasons.

First, the deletions at issue do not reflect a failure to take "reasonable steps" to preserve ESI. *See id.* By requiring "reasonable steps," Rule 37(e)(1) "does not call for perfection." Fed. R. Civ. P. 37(e) advisory committee's note to 2015 amendment. And the "routine, good-faith operation of an electronic information system" is relevant evidence that preservation efforts were

reasonable. *Id.* Accordingly, parties must generally retain "accessible" data but not "inaccessible" data that is routinely deleted "on the schedule set forth in the company's policy." *Zubulake v. UBS Warburg LLC*, 220 F.R.D. 212, 218 (S.D.N.Y. 2003). Here, OpenAI maintained a "routine" data-retention system, retaining tens of billions of user logs in its data warehouse, while allowing users to maintain control of their data through user-initiated deletions and temporary chats, as set forth in the company's policy. OpenAI notified Plaintiffs about that system in February 2024—a system with which Plaintiffs took no issue for ten months. *See supra*, pp.7-8. And when Plaintiffs finally moved to compel OpenAI to preserve additional output logs in January 2025, Judge Wang denied that request, expressly recognizing the "privacy and technological considerations" OpenAI had raised. NYT ECF 441, at 2. OpenAI also investigated the technical feasibility of halting automated deletions and implemented a litigation hold on electronically stored corporate documents. Those efforts amply satisfied OpenAI's obligations.

This Court's Preservation Order confirms that OpenAI took reasonable steps. *See* Fed. R. Civ. P. 37(e)(1) advisory committee's note to 2015 amendment ("The duty to preserve may … be triggered or clarified by a court order in the case."). The Order addressed OpenAI's preservation obligations explicitly and, in Magistrate Judge Wang's words, did not require "wholesale" preservation or for OpenAI to "stop all … deletion processes." Pls.' Ex. 31 (Hr'g Tr. 34:19-35:2). Instead, the Order directed OpenAI to preserve in-scope chats during the period the Preservation Order remained effective. *See* Ex. 15 (MDL ECF 629) (terminating order). OpenAI complied with that directive to the letter.

Second, any allegedly lost data does not prejudice Plaintiffs. Under Rule 37(e)(1), the moving party must show that the loss of ESI caused "*actual* prejudice." *See Health & Happiness (H&H) US LLC v. Nutramax Lab'ys, Inc.*, 2026 WL 836132, at *12 (S.D.N.Y. Mar. 26, 2026).

Here, retention of additional output logs would have been entirely duplicative, dispelling any suggestion of prejudice. Plaintiffs have a representative sample of 20 million output logs they can use to assess whether ChatGPT reproduced copyrighted works to users. *See* MDL ECF 104 (proposed order describing methodology); MDL ECF 601, at 1 (joint letter on sampling).[9] Given their access to tens of millions of representative output logs, News Plaintiffs' inability to access still more logs does not cause *actual* prejudice to their case.

Third, preclusion of the Merits Sample would be "greater than necessary to cure" any asserted prejudice. Fed. R. Civ. P. 37(e)(1). Precluding that dataset would significantly hamper OpenAI's ability to mount a defense. It would also eliminate the *only* representative dataset available to either side on the central regurgitation issue in this case, injuring *Plaintiffs*' own ability to prove their claims. A narrower remedy, such as permitting Plaintiffs to note the 5% deletion of sampling materials for the factfinder, would fully address any possible prejudice without the disproportionate and self-defeating step of eliminating the sample altogether.

II. This Court Should Not Enter Established Fact Sanctions Under Rule 37

The Court should decline Plaintiffs' extreme request to find as established 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 user outputs," and to prohibit OpenAI from arguing the contrary. Mot. 33 (capitalization omitted). Plaintiffs' request ignores Rule 37's text and fails to satisfy the controlling standard, which Plaintiffs conspicuously omitted from their motion.

[9] OpenAI also *preserved* millions of API logs under the Preservation Order, but Plaintiffs dropped any request that OpenAI *produce* a sample of those logs.

1. Rule 37 authorizes established fact sanctions in only three circumstances, none of which is present here. *First*, Rule 37(b)(2)(A)(i) and (ii) authorize a court to "direct[] that … designated facts be taken as established for purposes of the action" and "prohibit[] the disobedient party from supporting or opposing designated claims or defenses." But those "severe" sanctions, *Joint Stock Co. Channel One Russ. Worldwide v. Infomir LLC*, 2018 WL 4760345, at *6 (S.D.N.Y. Sept. 28, 2018), are available "only when a litigant fails 'to comply with a court order,'" *Evans v. Griffin*, 932 F.3d 1043, 1046 (7th Cir. 2019) (quoting Fed. R. Civ. P. 37(b)); 8B Wright & Miller, Federal Practice & Procedure § 2289 (3d ed. 2026) ("Rule 37(b) usually has no application if there has not been a court order."). *Second*, Rule 37(c) permits an established fact sanction where "a party fails to provide information … as required by Rule 26(a) or (e)." Fed. R. Civ. P. 37(c)(1)(C). And *third*, Rule 37(d) permits such a sanction where a party "fails to serve its answers, objections, or written response" "after being properly served with interrogatories under Rule 33." Fed. R. Civ. P. 37(d)(1)(A)(ii).[10]

Plaintiffs' (at 33-37) justifications for established fact sanctions do not match up with any of those Rules. Plaintiffs (at 33-34) first argue that OpenAI "refused to produce" certain information in response to requests for production. In their view (at 34), OpenAI should have produced information about its Project Giraffe "intersection evaluations." But OpenAI *did*

[10] Plaintiffs (at 33) make the passing assertion that subdivision "(e)(1)" provides another ground for established fact sanctions. That is incorrect. Rule 37(e)(1) simply empowers a court to "order measures no greater than necessary to cure the prejudice" a party suffers due to the loss of ESI. Fed. R. Civ. P. 37(e)(1). But courts "must be careful to ensure that 'curative measures under subdivision (e)(1) do not have the effect of measures that are permitted under (e)(2) only on a finding of intent to deprive another party of the lost information's use in the litigation.'" 8B Wright & Miller, Federal Practice & Procedure § 2284.2 (3d ed. 2026) (quoting Fed. R. Civ. P. 37(e) advisory committee's note to 2015 amendment). Here, an order designating certain facts as *established* would be more severe than an order "*presum[ing]* that the lost information was unfavorable to the party"—a sanction available only upon a showing of "intent to deprive" under Rule 37(e)(2).

produce these evaluations (and their source code) with ample time for Plaintiffs to utilize them. *See* MDL ECF 1295, at 1-2. In any case, Plaintiffs' complaints about the sufficiency and speed of OpenAI's discovery responses are not a basis for established fact sanctions under Rule 37. Those complaints do not constitute "an order" under Rule 37(b)(2)(A), and because OpenAI produced the evaluations, OpenAI complied with its obligations under Rule 26(e).

Plaintiffs (at 35) next contend that an established fact sanction "is appropriate due to OpenAI's failure to produce the same information in response to multiple interrogatories and requests for inspection." That is mistaken. Rule 37(d)(1)(A)(ii) and (d)(3) authorize established fact sanctions "only [for] a complete failure to respond" to a discovery request. *Wisser v. Vox Media, Inc.*, 2020 WL 1547381, at *5 (S.D.N.Y. Apr. 1, 2020). The Rule "is inapplicable if the party has served answers to interrogatories but the answers are thought to be incomplete or evasive." 8B Wright & Miller, Federal Practice & Procedure § 2291 (3d ed. 2026). Because OpenAI responded and objected to Times Interrogatory No. 9 and Times Request for Inspection No. 6 on June 21, 2024, Plaintiffs cannot invoke Rule 37(d) as a basis for obtaining established fact sanctions under Rule 37(b).

Plaintiffs (at 35-36) also attempt to root their request in unspecified "meet-and-confer orders." But OpenAI did not "fail[] to obey" the Court's meet-and-confer orders: it exhaustively and routinely met-and-conferred with Plaintiffs' counsel over the course of discovery. Fed. R. Civ. P. 37(b)(2)(A). And in any event, OpenAI *did* produce a spreadsheet containing New York Times-related intersection evaluations, just as Plaintiffs demand in their motion, *see* Mot. 36.

2. Even had Plaintiffs identified a proper textual predicate for established fact sanctions, such sanctions would be unwarranted. Under the framework established in *Residential Funding Corp. v. DeGeorge Financial Corp.*, a party seeking established fact sanctions based on the non-

production of evidence must establish three elements: "(1) that the party having control over the evidence had an obligation to timely produce it; (2) that the party that failed to timely produce the evidence had 'a culpable state of mind'; and (3) that the missing evidence is 'relevant' to the party's claim or defense such that a reasonable trier of fact could find that it would support that claim or defense." 306 F.3d at 107.[11] Plaintiffs do not mention, much less satisfy, these elements.

First, even assuming OpenAI had an obligation to produce "'intersection' evaluations," "the over 80 million output logs[,] and other Project Giraffe-related discovery," Mot. 34, 37, OpenAI timely produced much of that information. It produced a spreadsheet containing OpenAI's intersection evaluation for New York Times material on January 26, 2026. Pls.' Ex. 50. It produced the "reservoirs" of 78 million and 9 million de-identified ChatGPT conversations on March 20, 2026, pursuant to the Court's March 9, 2026 Order. Ex. 20 at 1. And it produced source code for Project Giraffe in February 2025. Pls.' Ex. 11. Producing material in time for Plaintiffs to incorporate it into their expert reports and overall case strategy is a far cry from "fail[ing] to produce [material] in time for trial." *Residential Funding*, 306 F.3d at 101.

Plaintiffs' reliance (at 37) on *Cine Forty-Second St. Theatre Corp. v. Allied Artists Pictures Corp.*, 602 F.2d 1062, 1068 (2d Cir. 1979), for the proposition that "hopelessly belated compliance should not be accorded great weight," is inapposite. *Cine* involved a party who defied multiple explicit orders to respond to specific interrogatories for several years. *See id.* at 1064-65, 1068. But here, OpenAI's production timeline reflects the genuine scale and complexity of the underlying data—not strategic delay—and, critically, OpenAI produced the at-issue information

[11] The Second Circuit has abrogated *Residential Funding*'s holding that "sanctions under Rule 37(e)(2)" may be imposed upon a mere showing of "negligence." *Hoffer*, 128 F.4th at 435. But *Hoffer* expressly declined to "call into question the applicability of *Residential Funding* to cases that fall outside the province of Rule 37(e)(2)." *Id.* at 439. *Residential Funding*'s framework thus governs Plaintiffs' request for established fact sanctions under Rule 37(b)(2)(A).

pursuant to the Court's orders, giving Plaintiffs ample time to incorporate it into their expert reports and litigation strategy.

Second, OpenAI did not withhold "intersection evaluations" or other Project Giraffe materials with a "culpable state of mind." *Residential Funding*, 306 F.3d at 107. Culpability in this context means that evidence was withheld "knowingly" or "negligently." *Id.* at 108 (citation and emphasis omitted). But a showing of culpability must be tied to specific acts that caused the non-production. *Cf. Kortright Cap. Partners v. Investcorp Inv. Advisers Ltd.*, 330 F.R.D. 134, 139 (S.D.N.Y. 2019) (examining specific "conduct … to establish that [a party] acted with a negligent state of mind").

OpenAI did not culpably withhold intersection evaluations or Project Giraffe-related information. On the contrary, OpenAI produced Project Giraffe's source code, intersection evaluation pipeline code, and related ESI through agreed-upon search terms, and the underlying 78 million dataset used for the NYT Recall Pipeline was available for inspection for over a year before Mr. Monaco's second deposition. *See* MDL ECF 1267, at 2. Plaintiffs therefore had this material in their possession in more than enough time to make use of it. That they did not appreciate its significance until Mr. Monaco's second deposition reflects a failure of diligence on Plaintiffs' part, not culpable concealment by OpenAI.

Third, Plaintiffs cannot establish that the allegedly "withheld" information is "relevant" to their claims for purposes of obtaining established fact sanctions. *See Residential Funding*, 306 F.3d at 108-09. Relevance for purposes of obtaining adverse inference or established fact sanctions requires "something more than sufficiently probative to satisfy Rule 401 of the Federal Rules of Evidence." *Id.* When "an assessment of relevance is possible" because, as here, the evidence has been "produced," the nonmoving party "should sustain liability for breaching its

discovery obligations where such breach causes injury, but the moving party should not obtain a windfall for uncovering evidence that would have made little difference in the underlying case." *Kortright*, 330 F.R.D. at 139 (citation omitted).

While the "intersection evaluations," "over 80 million output logs," and "other Project Giraffe-related discovery," Mot. 37, may be "arguably relevant to" Plaintiffs' infringement claims, "[t]he inference that [Plaintiffs] seek"—that ChatGPT output logs include substantial and systematic grounding on and regurgitation of News Plaintiffs' copyrighted material into user outputs—would hand Plaintiffs an unwarranted windfall. *Kortright*, 330 F.R.D. at 139. Plaintiffs cannot show that the timing of the productions prevented them from litigating their claims. Indeed, Plaintiffs' real gripe appears to be that these materials show their claims premised on regurgitation are baseless, because they prove OpenAI's products *do not* systematically regurgitate Plaintiffs' asserted works.

In addition, Plaintiffs *have access* to the material they seek in relation to their established fact sanction argument (*i.e.*, intersection evaluations, the 78 million de-identified outputs, and Project Giraffe-related information). The only purpose for established fact sanctions in this circumstance would be to "essentially absolve[] [Plaintiffs] of the duty to prove liability." *Id.* Given "the comparative windfall an [established fact sanction] would yield, the amorphous prejudice suffered by [Plaintiffs], and the fact that [OpenAI] has come forward with the evidence, the [sanction] is unwarranted." *Id.*

III. Plaintiffs' Rule 37(e)(2) Claim Fails Because OpenAI Did Not Act with Intent to Deprive

OpenAI did not act with the intent to deprive Plaintiffs of the use of output log data in this litigation. As relevant here, Rule 37(e)(2) permits courts to: (A) "presume that" lost ESI "was unfavorable" to a sanctioned party or (B) "instruct the jury that it may or must presume the

information was unfavorable to the party." Fed. R. Civ. P. 37(e)(2). But these severe sanctions are available only when "a party failed to take reasonable steps to preserve" the lost ESI and that "party acted with the intent to deprive another party of the information's use in the litigation." *Id.* Here, neither condition is met. OpenAI took "reasonable steps" to preserve its output logs for this litigation. *See supra*, pp.31-33. And OpenAI did not "inten[d] to deprive" Plaintiffs of the ability to use those output logs.

The "specific intent required to impose sanctions under Rule 37(e)(2) sets a … high bar" that Plaintiffs cannot clear. *Hoffer*, 128 F.4th at 440. Indeed, sanctions under Rule 37(e)(2) require proof of "*intentional* loss or destruction of evidence"; "negligent[] or knowing[]" misconduct "will not suffice." *Id.* at 438. This "standard is 'both stringent and specific,' and contemplates 'not merely the intent to perform an act that destroys ESI but rather the intent to actually deprive another party of evidence.'" *Id.* at 440 (citation omitted). "Proving such specific intent is not easy," and Plaintiffs come nowhere close. *Id.*

Plaintiffs' complaint (at 38) about the deletion of "billions of output logs" presumably refers to two categories of conversation logs: (1) conversations deleted at the request of individual ChatGPT users who exercised their contractual and privacy-law rights, and (2) conversations generated through the "Temporary Chat" feature, which by design are not retained beyond 30 days. In Plaintiffs' view, OpenAI's deletion of these logs cannot be "credibly explained other than bad faith." Mot. 38-39 (citation omitted). Plaintiffs are wrong. Deletion of these logs reflects user choices and pre-existing product design, not OpenAI's "intent to actually deprive [Plaintiffs] of evidence." *Hoffer*, 128 F.4th at 440 (citation omitted). OpenAI's retention policies were established before this litigation, disclosed to Plaintiffs and the Court, and maintained consistently throughout the litigation.

OpenAI's response to the Preservation Order demonstrates the opposite of an intent to deprive Plaintiffs of the use of conversation data in this litigation. As discussed, as soon as the Court entered the Order, OpenAI investigated the steps it could take, began retaining and segregating all covered, in-scope conversation data generated after the Order went into effect, and clarified with the Court that additional steps were not required. *See supra*, p.11. OpenAI thus took "reasonable" steps given the circumstances. *See* Fed. R. Civ. P. 37(e).

Plaintiffs (at 38) also argue that OpenAI misrepresented its "existing technical capacity" to search output logs, suggesting that "OpenAI's [Project] Giraffe evaluations, its search of API output logs, and the spoliation sampling exercise" reveal that OpenAI had greater search capabilities than it had initially claimed. But Plaintiffs make no effort to grapple with the key technical distinctions between the searches it lists (each of which required OpenAI to search a limited set of *uncompressed* output logs) and the searches that OpenAI described as technically difficult (which required searching all 50 billion *compressed* chat logs stored in OpenAI's data warehouse). *See supra*, pp.18-20 (discussing differences). To this day, OpenAI *still lacks the capacity* to search its compressed data warehouse without first decompressing the data, and Plaintiffs fail to identify any representation to the contrary. Thus, OpenAI did not misrepresent its "existing technical capacity," Mot. 38, let alone make misrepresentations that would impute an intent to deprive Plaintiffs from using conversation data.

IV. Plaintiffs Are Not Entitled to Attorneys' Fees

OpenAI's conduct throughout discovery has been substantially justified and marked by good-faith efforts to resolve genuinely complex issues of first impression. Although Plaintiffs invoke four legal grounds for attorneys' fees, each is meritless.

First, Plaintiffs (at 39) claim that Rule 37(b)(2)(C) entitles them to fees. But that Rule mandates fees only against a "disobedient party." Fed. R. Civ. P. 37(b)(2)(C). And even then,

fees are improper if "the failure" to follow a court order "was substantially justified or other circumstances make an award of expenses unjust." *Id.* "Substantial justification for refusing discovery is determined according to 'an objective standard of reasonableness and does not require that the party have acted in good faith.'" *Klein v. Torrey Point Grp., LLC*, 979 F. Supp. 2d 417, 442 (S.D.N.Y. 2013) (citation omitted). Under that "objective standard," "[c]onduct is substantially justified if there was a genuine dispute or if reasonable people could differ as to the appropriateness of the contested action." *Id.* (citations omitted).

Genuine disputes pervade every aspect of the discovery at issue. For instance, reasonable parties differed—and this Court carefully considered both sides—on whether OpenAI's data warehouse could feasibly be searched at scale; on the scope and mechanics of the Preservation Order; on the appropriate sample size for output-log analysis; on the privacy implications of producing de-identified user conversations to opposing counsel; and on the relevance and reliability of incomplete, non-representative datasets like the 78-million-record collection. As this Court has recognized, where "there are genuine disputes over the appropriateness of the [conduct]," fees and "sanctions are not appropriate." *Rakower Law PLLC v. Berner*, 2022 WL 17820861, at *2 (S.D.N.Y. Dec. 14, 2022) (Wang, M.J.).

Plaintiffs (at 40-41) catalog the time and expense they devoted to the sandbox, Mr. Monaco's supplemental deposition, motion practice, and the sampling process, and then attribute it all to OpenAI's alleged misconduct. According to Plaintiffs, none of this would have been necessary had OpenAI just "produced its training dataset searches" and admitted it "had the ability to search its output logs all along." Mot. 40-41. But Rule 37(b)(2)(C) only authorizes fees against a party that "fail[ed] to obey an order"—and there was never an order to produce "training dataset searches" or search all "output logs." For good reason: OpenAI never conducted searches for

Plaintiffs' works across all its training datasets; nor has OpenAI had the ability to search for those works across all its output logs. *See supra*, pp.18-20. These genuine limitations, not any disobedience by OpenAI, gave rise to the circumstances Plaintiffs now condemn.

Second, Plaintiffs (at 39) suggest that Rule 37(d)(3) entitles them to fees. But that Rule permits fees only against a party who failed "to appear" for a deposition or entirely "fail[ed] to serve its answers, objections, or written responses" to interrogatories or a request for inspection. Fed. R. Civ. P. 37(d)(1)(A)(ii) & (d)(3). Plaintiffs have not shown that OpenAI made such a wholesale failure, *see supra*, pp.35, 43, so Rule 37(d)(3) cannot serve as a basis for fees.

Third, Plaintiffs (at 39-40) argue that this Court should exercise its "inherent authority to award attorneys' fees as a sanction for litigation misconduct." But a court "may assess attorney's fees" under that circumstance only "when a party has acted in bad faith, vexatiously, wantonly, or for oppressive reasons." *Chambers v. NASCO, Inc.*, 501 U.S. 32, 45-46 (1991) (cleaned up). That stringent standard is met only in extreme circumstances, like when "fraud has been practiced upon" the court, "the very temple of justice has been defiled," or "a party shows bad faith by delaying or disrupting the litigation or by hampering enforcement of a court order." *Id.* at 46 (citations omitted). None of those circumstances are present here.

Plaintiffs (at 39-40) resist this conclusion and argue that "misrepresentations about [OpenAI's] technical capabilities, the destruction of billions of user conversations, and the resulting need for extensive motion practice" warrant fees based on this Court's inherent authority. But once again, Plaintiffs get the premise wrong: OpenAI has not misrepresented its technical capabilities, user conversations were deleted pursuant to user requests and Temporary Chat settings that OpenAI disclosed to Plaintiffs and the Court, and the motion practice resulted from genuine disputes, not fraud or bad faith. *See supra*, pp.14-15.

Fourth, Plaintiffs (at 40-41) claim they are entitled to fees relating to Mr. Monaco's Rule 30(b)(6) deposition. But "to impose sanctions" and fees for an allegedly unprepared corporate deponent, "the inadequacies in [the] deponent's testimony must be egregious and not merely lacking in desired specificity in discrete areas." *Eid v. Koninklijke Luchtvaart Maatschappij N.V.*, 310 F.R.D. 226, 228 (S.D.N.Y. 2015) (citation omitted). Here though, Plaintiffs just raise specificity concerns. *See* Mot. 40-41 (objecting to responses where Mr. Monaco provided certain details but could not "remember" other details that happened "over three years ago" (citation omitted)). And Mr. Monaco was never "expected to know and testify about every detail on [every] topic." MDL ECF 848 at 4 (citation omitted). In any event, this Court ordered a supplemental deposition of Mr. Monaco, after which Plaintiffs raised no objection to his preparation.

V. The Court Should Not Exercise Its Inherent Authority to Impose Sanctions

Plaintiffs have not shown, and cannot show, their entitlement to sanctions as an exercise of this Court's inherent authority. Plaintiffs (at 42) start by badly misconstruing the relevant legal standard, arguing that claimed violations of Rule 37 "should be sufficient in itself to bring the Court's inherent authority into play." But the standard for inherent-authority sanctions is much higher than sanctions under Rule 37.

"Imposition of sanctions under a court's inherent powers requires a specific finding that [a party] acted in bad faith." *Wolters Kluwer*, 564 F.3d at 114. "Moreover, inherent-power sanctions are appropriate only if there is clear evidence that the conduct at issue is (1) entirely without color and (2) motivated by improper purposes." *Id.* "Conduct is entirely without color when it lacks any legal or factual basis; it is colorable when it has some legal and factual support, considered in light of the reasonable beliefs of the [party] whose conduct is at issue." *Id.* "A finding of bad faith, and a finding that conduct is without color or for an improper purpose, must be supported

with a high degree of specificity in the factual findings." *Id.* Plaintiffs do not mention this standard, let alone satisfy it.

First, OpenAI did not act with bad faith, nor was it motivated by "improper purposes." *Id.* Plaintiffs' central accusation (at 42) is that OpenAI should have sooner produced "intersection evaluations" that OpenAI "created by searching its training data and over 80 million output logs." But OpenAI did not immediately produce these evaluations because Plaintiffs asked OpenAI to identify instances of Plaintiffs' content in *all* of the company's training and output data, a different and broader request. *See, e.g.*, Ex. 17 (NYT RFP Nos. 22, 45, and 93). OpenAI worked in good faith with Plaintiffs and the Court to accommodate those requests, culminating in the sandbox (which gave Plaintiffs access to *all* relevant OpenAI training data, at OpenAI's expense) and the Merits Sample (which gave Plaintiffs a representative sample of *all* OpenAI conversation outputs). The intersection evaluations, which OpenAI did produce, compare a single training dataset to incomplete conversation outputs—making the evaluations useless for generalizing ChatGPT's behavior. *See* Ex. 19 (Kellogg Decl. ¶ 3). This record does not establish bad faith—tellingly, even Plaintiffs do not argue as much.

Second, OpenAI's conduct is not "entirely without color." *Wolters Kluwer*, 564 F.3d at 114. OpenAI's position that Plaintiffs wanted *all* of the company's training and output data—not useless comparisons between a single dataset and a sliver of incomplete outputs—had "some … factual support," *id.*; namely, Plaintiffs' own discovery requests, *see* Ex. 17 (NYT RFP Nos. 22, 45, and 93).

OpenAI has gone to herculean efforts to give Plaintiffs the discovery they want. To summarize: OpenAI gave Plaintiffs extensive access to all of OpenAI's training data; it assembled and produced a comprehensive Merits Sample of representative conversation outputs; and when

Plaintiffs learned of Project Giraffe and OpenAI's intersection evaluations (based on material OpenAI voluntarily produced), OpenAI produced an additional 87 million incomplete conversation outputs. *See supra*, pp.22. As a result of this massive exchange of information, Plaintiffs are going into trial preparation possessing everything available they could possibly require to build their case. OpenAI's facilitation of that exchange of information does not constitute clear and convincing evidence of bad faith or sanctionable conduct.

CONCLUSION

For the foregoing reasons, OpenAI respectfully requests that the Court deny News Plaintiffs' motion for sanctions in its entirety.

Respectfully submitted,

Dated: August 13, 2026

By: */s/ Lisa S. Blatt*

LISA S. BLATT
WILLIAMS & CONNOLLY LLP
680 Maine Avenue SW
Washington, DC 20024
(202) 434-5000
lblatt@wc.com

By: */s/ R. James Slaughter*

R. JAMES SLAUGHTER
KEKER, VAN NEST & PETERS LLP
633 Battery Street
San Francisco, CA 94111
(415) 391-5400
rslaughter@keker.com

By: */s/ Margaret Graham*

MARGARET GRAHAM
LATHAM & WATKINS LLP
1271 Avenue of the Americas
New York, NY 10020
(212) 906-1200
margaret.graham@lw.com

By: */s/ Joseph C. Gratz*

JOSEPH C. GRATZ
MORRISON & FOERSTER LLP
425 Market Street
San Francisco, CA 94105
(415) 268-7000
jgratz@mofo.com

Attorneys for Defendants
OPENAI, INC., OPENAI GP, LLC, OPENAI, LLC, OPENAI OPCO LLC, OPENAI GLOBAL LLC, OAI CORPORATION, LLC, OPENAI HOLDINGS, LLC

* All parties whose electronic signatures are included herein have consented to the filing of this document, as contemplated by Rule 8.5(b) of the Court's ECF Rules and Instructions.