UNITED STATES DISTRICT COURT
SOUTHERN DISTRICT OF NEW YORK

IN RE: OPENAI, INC., COPYRIGHT INFRINGEMENT LITIGATION This Document Relates To: 1:23-cv-11195-SHS-OTW 1:24-cv-04872-SHS-OTW 1:24-cv-03285-SHS-OTW 1:24-cv-01515-SHS-OTW 1:25-cv-04315-SHS-OTW	1:25-md-03143 (SHS) (OTW) **DECLARATION OF DEFENDANTS' EXPERT AVI GOLDFARB IN SUPPORT OF MOTION FOR SUMMARY JUDGMENT** REDACTED — PUBLIC VERSION

I, Avi Goldfarb, Ph.D., declare under penalty of perjury, pursuant to 28 U.S.C. § 1746, as follows:

1. I submit this Declaration in support of OpenAI Defendants' (collectively, "OpenAI") Motion for Summary Judgment.

A. Professional Background and Qualifications

2. I am the Rotman Chair of Artificial Intelligence and Healthcare and a Professor of Marketing at the Rotman School of Management, University of Toronto. I have held those positions since 2018 and 2012, respectively, and have served as the Rotman School of Management's Chief Data Scientist at the Creative Destruction Lab and as a Research Associate at the National Bureau of Economic Research. I received my Ph.D. in Economics from Northwestern University in 2002.

3. I have authored or co-authored over 120 peer-reviewed journal articles in Industrial Organization, Marketing, Digital Economics, and Economics of Artificial Intelligence. I have also authored or co-authored 28 articles and co-authored or edited four books on Artificial Intelligence. Many of my publications concern digital markets and use econometric analysis (statistical analysis used to study economic relationships).

4. My research has appeared in peer-reviewed journals, including the American Economic Review, Management Science, Marketing Science, and the Journal of Marketing Research. I have taught courses focused on econometric and data analysis, including statistical methods for studying economic data, and have testified before the United States Senate Judiciary Committee on competition and privacy in digital advertising.

B. Assignment and Materials Reviewed

5. I have been retained by counsel representing OpenAI to provide economic analysis of the impact, if any, of OpenAI's alleged use of the works Plaintiffs have identified in this case (the "Asserted Works") in training OpenAI's large language models ("LLMs"), in product outputs generated by those LLMs, and in retrieval-augmented generation ("RAG") and search functionality (features that retrieve information and use it to generate an answer) on the market for or value of those works and the Plaintiffs' journalism generally. In this declaration, "training data" means data used for pre-training. My work addressed economic substitution, traffic, revenue, subscriptions, and user engagement.

6. I submitted five opening reports on April 24, 2026, addressing The New York Times Company ("The Times" or "NYT")[1]; Daily News, LP; Chicago Tribune Company, LLC; Orlando Sentinel Communications Company, LLC; Sun-Sentinel Company, LLC; San Jose Mercury-News, LLC; DP Media Network, LLC; ORB Publishing, LLC; and Northwest Publications, LLC (together, the "Daily News Plaintiffs" or "DNP")[2]; the Center for Investigative Reporting, Inc. ("CIR")[3], including Mother Jones and Reveal; The Intercept Media, Inc. ("The Intercept")[4]; and Ziff Davis, Inc. ("Ziff Davis" or "ZD")[5] (collectively, the "News Plaintiffs"). I submitted a combined rebuttal report for The Times and Daily News Plaintiffs[6] and a separate rebuttal report for Ziff Davis[7] on May 29, 2026. I submitted a combined The Times-Daily News Plaintiffs Reply Report[8], a combined CIR-The Intercept Reply Report[9], and a Ziff Davis Reply Report[10] on June 26, 2026. On July 29, 2026, I submitted a supplemental report[11] addressing The Times's revised Asserted Works list.

7. In preparing this declaration, I reviewed those reports and the underlying data, documents, testimony, and economic publications cited in them. I also reviewed the applicable errata to the April 24, 2026, opening reports and incorporated the corrections relevant to the

[1] "Goldfarb NYT Opening Report"
[2] "Goldfarb DNP Opening Report"
[3] "Goldfarb CIR Opening Report"
[4] "Goldfarb Intercept Opening Report"
[5] "Goldfarb ZD Opening Report"
[6] "Goldfarb NYT-DNP Rebuttal Report"
[7] "Goldfarb ZD Rebuttal Report"
[8] "Goldfarb NYT-DNP Reply Report"
[9] "Goldfarb CIR-INT Reply Report"
[10] "Goldfarb ZD Reply Report"
[11] "Goldfarb NYT Supplemental Report"

opinions summarized below. This declaration incorporates my opinions about market harm and substitution in those reports that are relevant to the factor-four analysis (the fourth fair-use factor).

8. I was deposed in this matter on July 30, 2026.

C. Summary of Opinions

9. Based on my reports and the analyses summarized below, I found no empirical evidence that OpenAI's alleged use served as an economic substitute for the Asserted Works or News Plaintiffs' journalism, or caused a measurable market harm across traffic, revenue, subscriptions, or user behavior. Several widely used AI chatbots already existed during the periods covered by my data. If material substitution had occurred, I would have expected it to appear in the observed traffic, revenue, subscription, or individual-user activity measures. The data showed no such substitution. I also concluded that the same evidence did not support a substantially adverse impact from widespread use of the same type.

D. Economic Substitution and the Relevant Inquiry

10. In my reports, I addressed the economic question of whether OpenAI's alleged use of the News Plaintiffs' Asserted Works created an economic substitute for the Asserted Works specifically or for News Plaintiffs' journalism generally. My analysis covered alleged training use, product outputs, and RAG functionality and focused on whether those activities affected demand for the Asserted Works or News Plaintiffs' journalism more generally.

11. Economic substitution asks whether consumers consider the use of two products as reasonably interchangeable, and in fact use one in place of the other. The empirical test for economic substitution is not whether ChatGPT could answer a question; it is whether users who otherwise would have consumed News Plaintiffs' content instead used ChatGPT. I found no evidence of that kind of economic substitution in the available traffic, revenue, subscription, or user-level data for News Plaintiffs.

12. Technical capability, isolated prompts, or general statements that AI *can* generate news do not measure economic substitution. The relevant analysis is what consumers actually did and what effects followed in the market. I analyzed those outcomes to reach my opinions.

13. To assess whether News Plaintiffs experienced market harm as a result of OpenAI's alleged use of News Plaintiffs' Asserted Works, I used data produced by News Plaintiffs on traffic, revenue, subscriptions, and other business performance metrics. I also used third-party Comscore traffic data to analyze individual user behavior for the publishers which were large enough to be analyzed using this data source: The Times, DNP, and Ziff Davis. Plaintiffs' available revenue data could not be tied to individual Asserted Works. Therefore, I treated traffic and engagement as proxies for demand rather than as direct measures of revenue, and I separately evaluated overall business metrics such as revenue by type and subscribers.

14. The same conclusion applies to claims based on regurgitated or summarized outputs and to OpenAI's RAG or search functionality. Based on my review of News Plaintiffs' experts' reports, I did not identify any analysis showing that real users used such outputs instead

of News Plaintiffs' works, that any identified output displaced a visit to a particular News Plaintiff's website, or that such outputs caused a quantifiable loss of traffic, subscribers, revenue, or other business metrics.

E. Plaintiff-Specific Data Analyses

15. In the sections below, I present my Plaintiff-specific analyses of Asserted Works composition, financial performance, web traffic, and user-level behavior for each News Plaintiff. As explained below, these analyses are consistent with the conclusions summarized above: the data I reviewed for each News Plaintiff showed no measurable market harm attributable to OpenAI's alleged use and instead reflected a combination of long-term industry trends and Plaintiff-specific factors unrelated to OpenAI.

i. The New York Times

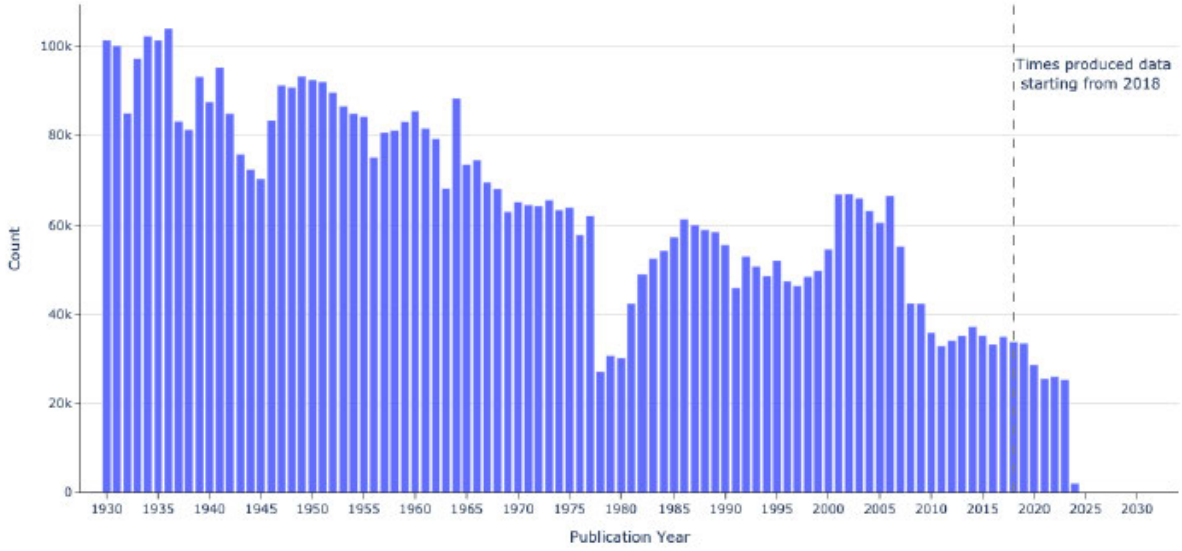
(1) The Times's Asserted Works

16. The Times revised its Asserted Works list on July 17, 2026. The Times's July 17, 2026, list contained 6,030,928 articles, added 1,657 articles, excluded 1,803 articles from the April 23, 2026, list, and was 146 articles smaller overall. The revised list covered publication dates from July 8, 1928, through April 6, 2024. The latest two entries with publication dates in February and April 2024 were Games pages.[12]

[12] Goldfarb NYT Supplemental Report ¶¶ 7-12.

17. As detailed below, the majority of The Times's Asserted Works pre-date OpenAI's LLMs by years or decades. Only 2.9% of the articles were published in or after 2018. Approximately 82.73% were published before 2000, 95.86% before OpenAI was founded in November 2015, 98.38% before GPT-3 was launched in May 2020, and 99.47% before ChatGPT was released in November 2022. Of the list, 3,993,611 articles, or 66.2%, were published before 1981, and 99.9% of those were archived. These figures show that the list largely consists of historical content rather than current reporting.[13]

Figure 1: The Times's Asserted Works by publication year



18. The Times's produced URL-level traffic data permitted only the direct analysis of article-level pageviews for articles published from January 1, 2018, through February 20, 2025, and only for up to 14 days after publication. Because only 2.9% of the revised Asserted Works

[13] Goldfarb NYT Supplemental Report ¶¶ 7-12.

were published in or after 2018, direct analysis could cover no more than that small fraction of The Times's Asserted Works and only within a short post-publication window.[14]

19. This data shows that reader demand for The Times's content was [redacted]. Within the URL-level data, [redacted] of pageviews came from the core news product. For [redacted] of pageviews recorded during the [redacted] after publication occurred within the [redacted]; the comparable figures were [redacted] for [redacted] within two days and [redacted] for [redacted] within [redacted]. Thus, according to the produced data, most reader demand for The Times's content is realized within [redacted] of publication.[15]

Figure 2: Cumulative share of pageviews per day after publication, by product



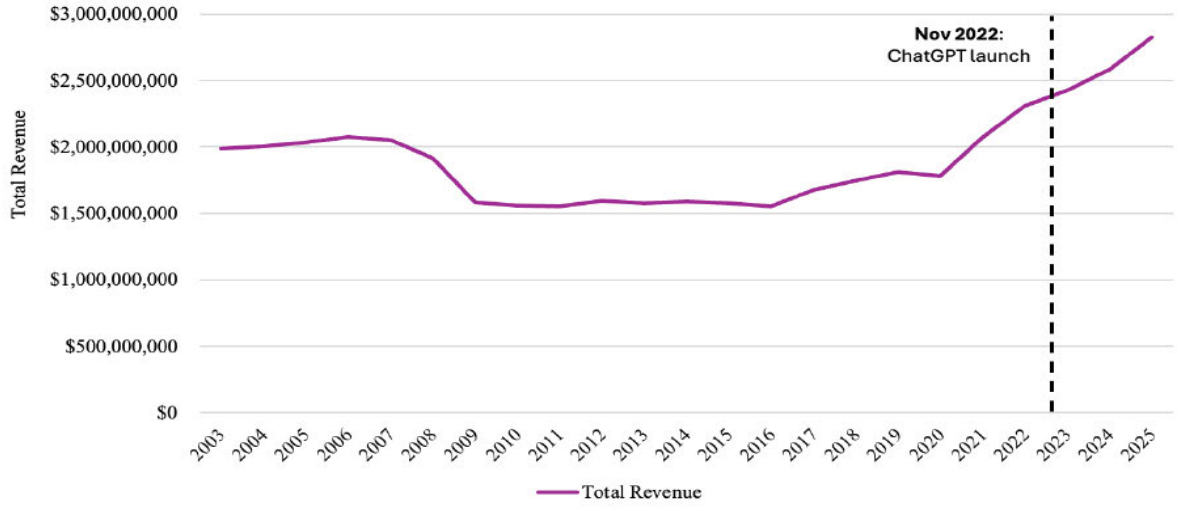
[14] Goldfarb NYT Opening Report ¶ 131.
[15] Goldfarb NYT Opening Report ¶¶ 131-133.

20. Accordingly, a ChatGPT output generated years after an Asserted Work's publication date would have limited opportunity to divert meaningful traffic from an article whose demand had [redacted]; the time gap between acquiring an article for training and releasing a trained model further limits that possibility. [16]

(2) Financial Performance

21. In my opinion, The Times's business performance does not show harm from OpenAI's products. Specifically, The Times's revenue grew steadily throughout this period, including after ChatGPT's public release. From 2020 through 2025, its revenue increased from $1,783,639,000 to $2,824,918,000, an increase of approximately 58%. Figure 3 illustrates the upward trajectory of The Times's revenue since 2003. Revenue continued to grow after ChatGPT's public release on November 30, 2022, including through 2025.[17]

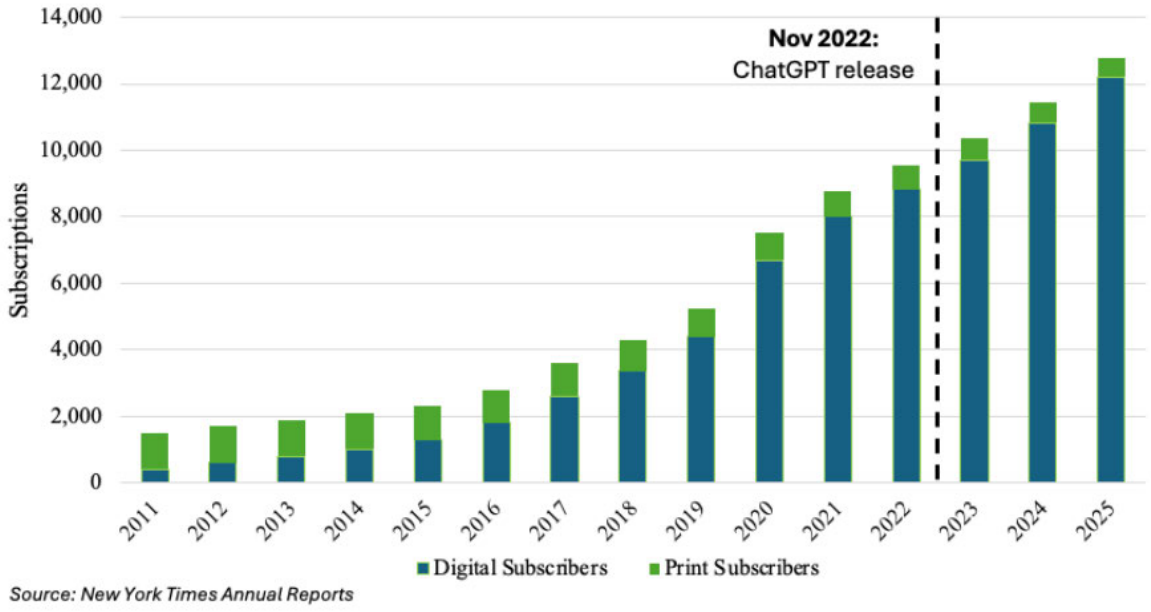
Figure 3: The Times's Revenue 2003-2025 (not adjusted for inflation)



[16] Goldfarb NYT Opening Report ¶¶ 134-139.
[17] Goldfarb NYT Opening Report ¶ 144.

22. The Times's subscriber base and licensing revenue [redacted] over this period, with subscriptions remaining strong after ChatGPT's release. From 2020 through 2025, The Times's subscriber base increased by approximately 70%. The Times added more than 1.4 million digital-only subscribers in 2025 and ended that year with approximately 12.78 million total subscribers. Its licensing revenue [redacted] pre-ChatGPT levels, and internal forecasts placed the licensing margin at approximately [redacted]. Internal measures also reported that subscribers were reading [redacted]. My analysis of The Times's financial data did not show a negative impact from OpenAI's products.[18]

Figure 4: The Times's Digital and Print Subscription Growth 2011-2025



[18] Goldfarb NYT Opening Report ¶¶ 144-147, 153-161.

(3) Web traffic datasets, methodology, and findings

23. For The Times, I used two traffic datasets produced by the Plaintiff: (1) Session Data, which reported monthly sessions[19] on The Times's website and app; and (2) Pageview Data, which reported hourly pageviews for each article during the first 14 days after publication. Both datasets began in 2018, with session data extending through February 5, 2025, and the pageview data covering articles published through February 20, 2025. Both datasets, described below, are consistent in showing that The Times's traffic [redacted] ChatGPT's release.[20]

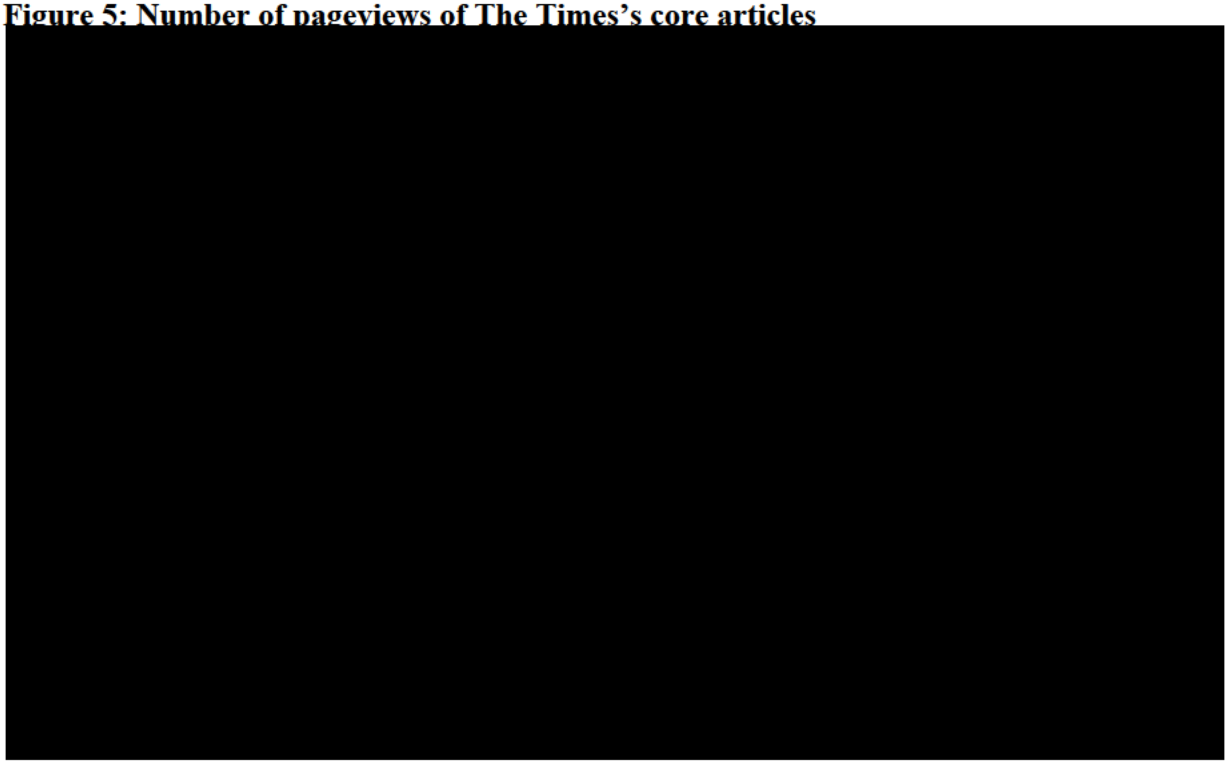
24. In the combined traffic analysis, core news sessions showed [redacted] at ChatGPT's November 30, 2022, release, while pageviews [redacted]. The figure below shows that the number of pageviews of The Times's core news articles by U.S. visitors [redacted] and the [redacted] suggests that ChatGPT has not caused a decrease in The Times's news pageviews.[21]

[19] A session is defined as activity ending after [redacted] of inactivity.

[20] Goldfarb NYT Opening Report ¶¶ 131, 154-155.

[21] Goldfarb NYT Opening Report ¶¶ 156-160.

Figure 5: Number of pageviews of The Times's core articles



25. The Times's session data classified each session by [redacted] "Core sessions" referred to sessions that [redacted]. I excluded [redacted] sessions from the focused analyses because [redacted], was not a product for which ChatGPT offered [redacted] In the product-level analysis, core news sessions [redacted] after ChatGPT's November 30, 2022, release. Overall core sessions were [redacted] after accounting for COVID-related fluctuations; traffic from [redacted] began [redacted] before ChatGPT's release, while direct visits [redacted] after ChatGPT's release.[22]

[22] Goldfarb NYT Opening Report ¶¶ 167-169, 201, 209.

26. My analysis also shows that referrals to The Times from Bing.com [redacted] after ChatGPT's public release on November 30, 2022, after the introduction of Bing Chat in 2023, and after the October 31, 2024, rollout of ChatGPT Search, rather than showing a ChatGPT-related [redacted].[23]

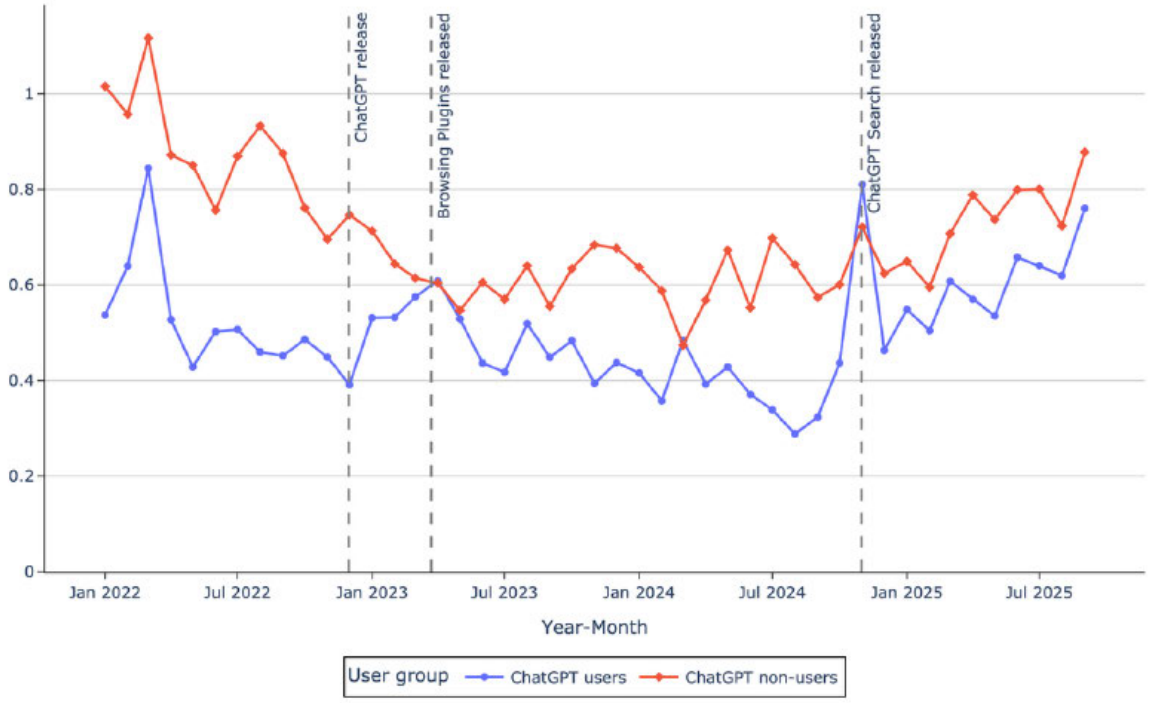
(4) Comscore audience measurement and site-visitation analysis

27. I used Comscore data to examine user-level Internet activity. Comscore is a third-party audience-measurement service that tracks activity associated with a unique machine identifier (an identifier for a particular computer or other device) and provides separate desktop and mobile datasets. For each machine-month observation, Comscore assigns a projection weight based on characteristics including age, gender, and household income, among other factors. Applying those weights projects the panel's observed activity to estimates representative of U.S. Internet users. For The Times, I defined a relevant user as one who visited a property of The Times at least once between January 2022 and September 2025. I defined a ChatGPT user as a relevant user who visited the ChatGPT website or app at least once after ChatGPT's November 2022 release; all other relevant users were ChatGPT non-users. As the figures below illustrate, the weighted average number of Times sessions per user for ChatGPT users relative to non-users showed no divergence consistent with substitution, and ChatGPT users did not reduce their visits to The Times after using ChatGPT.[24]

[23] *See* Goldfarb NYT Opening Report ¶¶ 168-209; Goldfarb NYT-DNP Rebuttal Report ¶¶ 112-115.
[24] *See* Goldfarb NYT Opening Report ¶¶ 210-213, 225-227, 376-378.

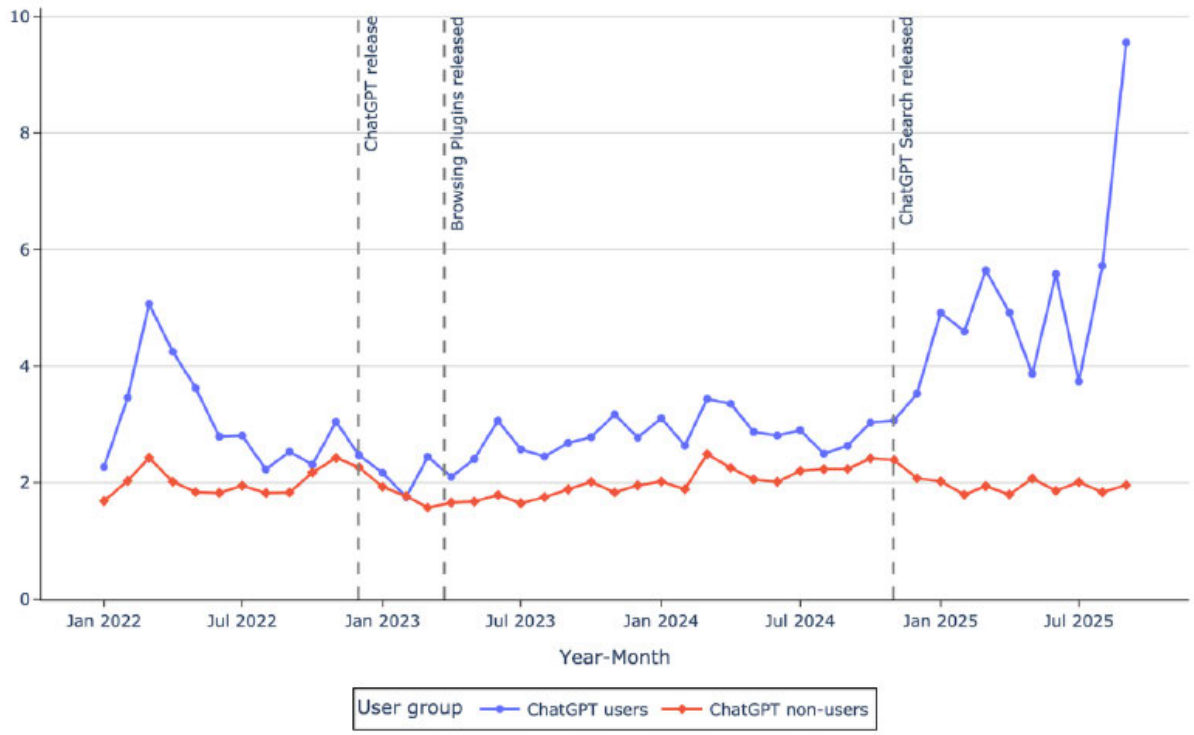
28. I first asked whether each person's average monthly visits to The Times changed after the release of ChatGPT. I found that, relative to non-users, ChatGPT users did not exhibit a systematic decline in their sessions to The Times after ChatGPT's public release in either the desktop or mobile data.[25]

Figure 6: Average weighted number of The Times sessions per relevant desktop user and ChatGPT usage status



[25] See Goldfarb NYT Opening Report ¶¶ 214-227.

Figure 7: Average weighted number of The Times sessions per relevant mobile user and ChatGPT usage status



29. In a separate first-use comparison, the weighted average number of The Times's sessions remained similar before and after users' first observed ChatGPT use in both the desktop and mobile data, with no sustained decline after first use. If ChatGPT substituted for The Times's traffic, I would have expected sessions to be lower after the first use than in the immediately preceding period. The figures demonstrate no such pattern.[26]

[26] Goldfarb NYT Opening Report ¶¶ 221-223.

Figure 8: Average weighted number of The Times sessions per relevant ChatGPT desktop user

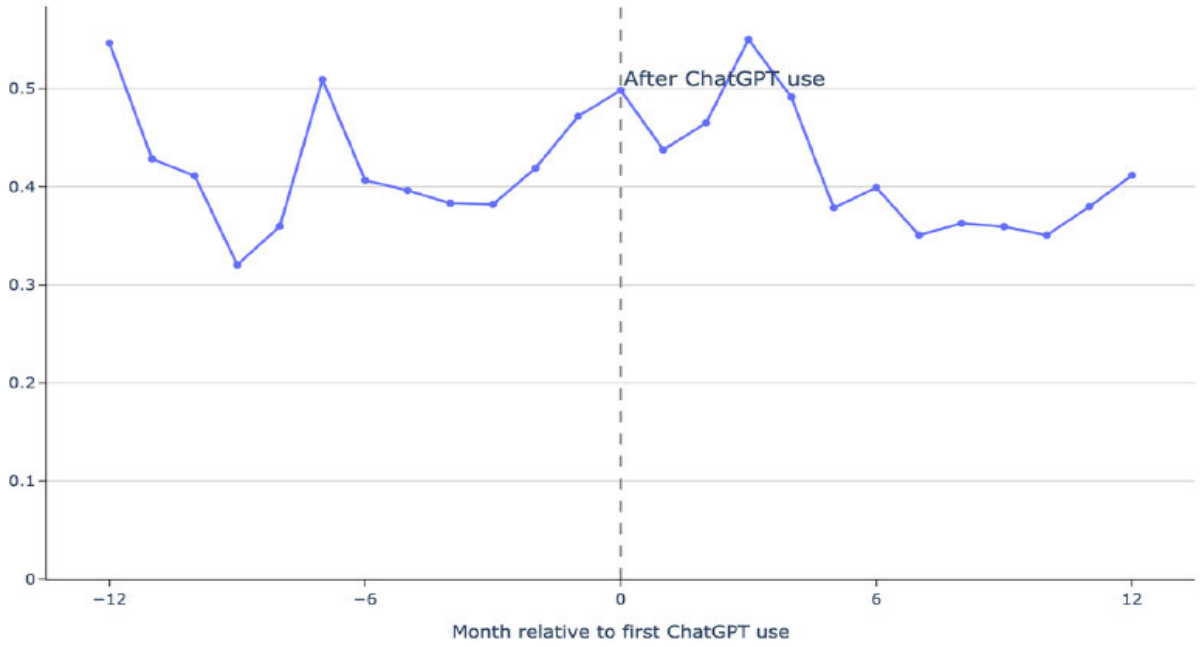
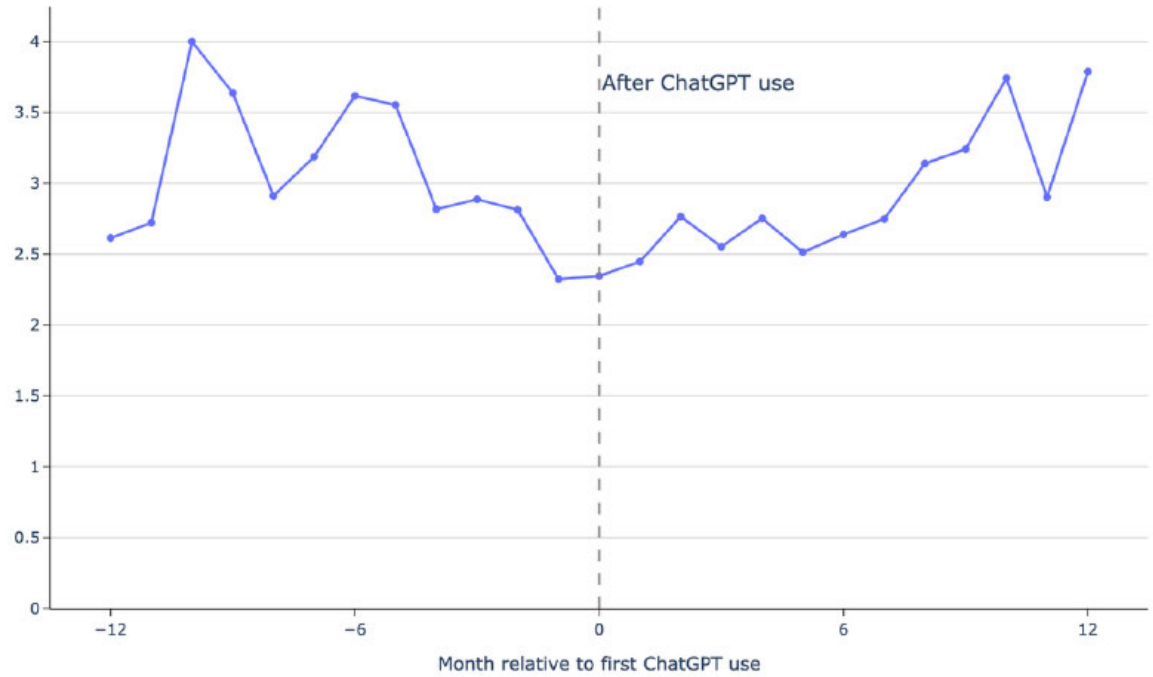


Figure 9: Average weighted number of The Times sessions per relevant ChatGPT mobile user



(5) Difference-in-differences and statistical methods, results, and checks

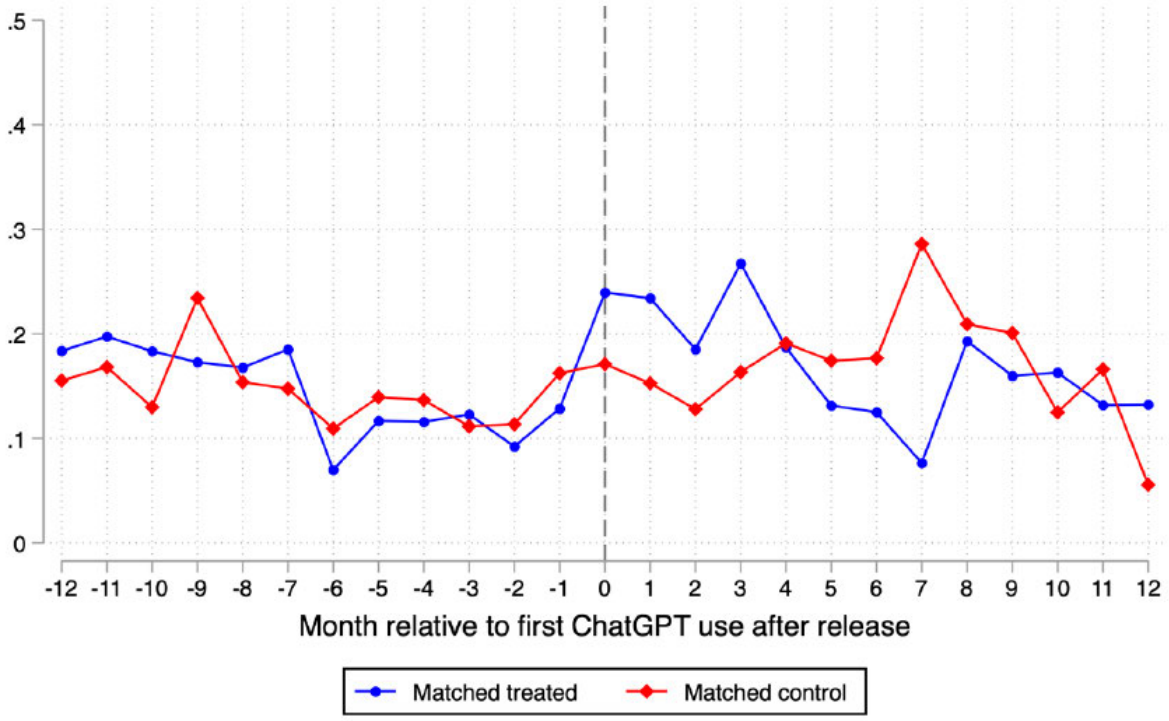
30. My principal econometric design was a difference-in-differences ("DiD") analysis of monthly sessions to The Times. I compared the change in visits to The Times before and after each user's first use of ChatGPT with the change over the same period for users who did not use ChatGPT. Subtracting those two changes helped remove effects common to both groups, such as seasonality, major news events, and other time-specific changes. For this piece of my analysis, I limited the ChatGPT user group to users who visited ChatGPT in at least two different months, ensuring that they are not one-time or incidental users. I also estimated an event-study specification, which reported results month by month relative to each user's first observed use, so that I could examine both pre-adoption trends and post-adoption effects rather than assume one immediate effect.[27]

31. I found no statistically significant reduction in visits to The Times at any point through one year after ChatGPT adoption. I also extended the analysis to 24 months after adoption for users who first used ChatGPT before October 2023; the post-adoption estimates were generally statistically indistinguishable from zero or positive and showed no negative pattern. My analysis included periods that, on average, reflected later versions of ChatGPT such as ChatGPT Search.[28]

[27] *See* Goldfarb NYT Opening Report ¶¶ 228-236.
[28] *See* Goldfarb NYT Opening Report ¶¶ 237-239, 379-385; Goldfarb NYT-DNP Reply Report ¶¶ 235-236.

32. In a separate matching analysis, I paired each ChatGPT adopter with a non-user who had comparable pre-adoption behavior and demographic characteristics. The matched groups tracked one another before adoption, and the analysis found no statistically significant reduction in The Times sessions after adoption.[29]

Figure 10: Average number of The Times sessions per user



33. I tested the robustness of my findings using alternative samples and methods, including extending the pre-adoption matching window from six to twelve months, broadening the analysis to include Comscore users who did not visit The Times, examining adopters with at

[29] *See* Goldfarb NYT Opening Report ¶¶ 240-246, 386-391.

least twenty-four months of follow-up, and testing mobile data and alternative specifications. The results continued to show no negative impact on The Times sessions.[30]

ii. Daily News Plaintiffs

(1) Daily News Plaintiffs' Asserted Works

34. For the Daily News Plaintiffs, I analyzed the eight publisher properties: New York Daily News, Chicago Tribune, Orlando Sentinel, Sun-Sentinel, San Jose Mercury-News, Denver Post, Orange County Register, and Pioneer Press. DNP's Asserted Works list contained approximately 2,592,766 unique URLs with publication dates from January 1985 through January 2024.[31]

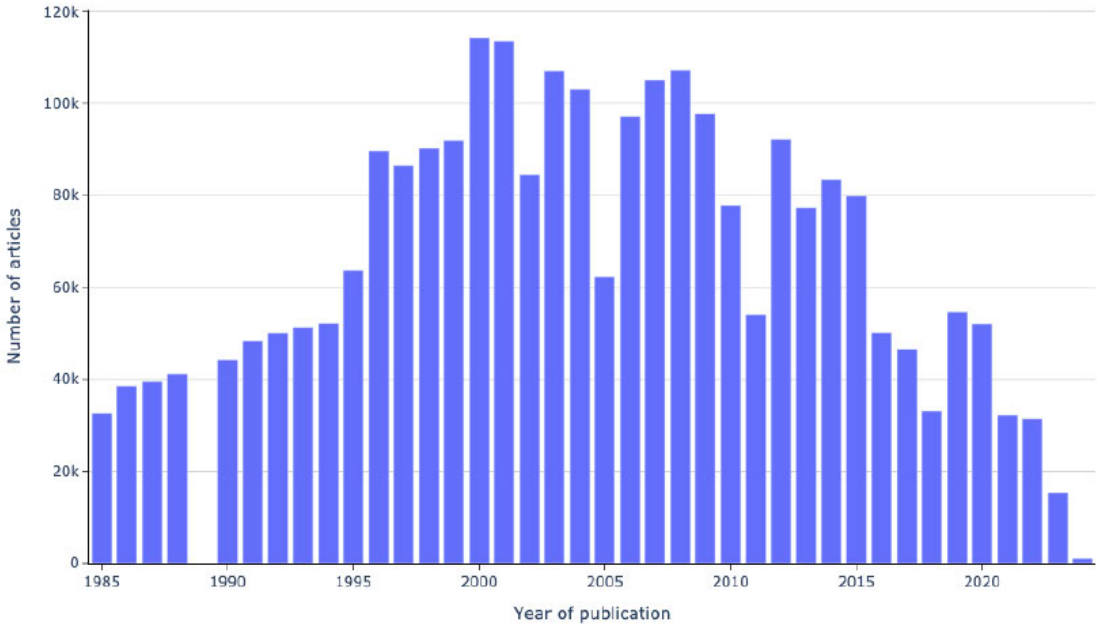
35. The aggregate publication year distribution is shown in the figure below, which presents the number of Daily News Plaintiffs' Asserted Works by year. As with The Times, these figures show that DNP's Asserted Works are overwhelmingly composed of older content published before OpenAI's LLMs existed. 31.6% of the articles were published before 2000, 87.4% before OpenAI was founded in November 2015, and 99.2% before ChatGPT was released in November 2022; 1% were published in or after September 2022, and all works were published before GPT-4o was released in May 2024.[32]

[30] *See* Goldfarb NYT Opening Report ¶¶ 379-385; Goldfarb NYT-DNP Reply Report ¶¶ 220-236.

[31] Goldfarb DNP Opening Report ¶¶ 35–36, 130.

[32] Goldfarb DNP Opening Report ¶¶ 35-36, 130-131.

Figure 11: Number of Asserted Works by publication year



36. For the individual DNP publishers, the Asserted Works counts and publication-date ranges were:

a. Chicago Tribune: 1,205,350 (January 1985–June 2023)

b. Denver Post: 89,269 (December 2006–January 2024)

c. Mercury News: 130,331 (January 2007–January 2024)

d. New York Daily News: 525,341 (September 1995–June 2023)

e. Orange County Register: 35,603 (January 2006–April 2015)

f. Orlando Sentinel: 424,822 (July 1999–December 2022)

g. Pioneer Press: 47,532 (January 2007–October 2023)

h. Sun Sentinel: 134,518 (November 1992–April 2022).[33]

37. DNP produced traffic data for its eight properties, but most Asserted Works predated the period covered by the article-level traffic data. For the article-level traffic data used in the decay analysis, data was available for all eight Plaintiffs from August 2018 through October 2025. Publication dates could be extracted for 67% of the URLs. I restricted the plotted sample to articles published after August 2018 through December 2021, before ChatGPT's release, to examine the general pattern of traffic decay.[34]

38. Across the DNP properties, traffic to individual articles was [redacted]. Specifically, approximately [redacted] of pageviews occurred [redacted] and approximately [redacted] occurred [redacted]. This means that even for the Asserted Works published three months or more before ChatGPT's release, approximately [redacted] of their expected article traffic had already occurred before ChatGPT became available. As a result, relatively little remaining traffic to those works could potentially have been diverted by ChatGPT after its release.[35]

[33] Goldfarb DNP Opening Report ¶ 348, Table 4.
[34] Goldfarb DNP Opening Report ¶¶ 127-134.
[35] Goldfarb DNP Opening Report ¶¶ 133-135.

Figure 12: Cumulative share of pageviews over month, by publication year



39. Given the lag between acquiring an article for training and releasing a model trained on it, DNP's works would generally have exhausted the value associated with hot or breaking-news demand before OpenAI allegedly used them to train its LLMs or released a product that could generate allegedly substitutable outputs.[36]

(2) Financial Performance

40. DNP's financial data did not show ChatGPT-related harm. Both DNP's revenue and subscriptions were [redacted] ChatGPT's release, reflecting a longer-term industry shift rather than a new AI-driven effect. DNP's operating revenue [redacted], with the [redacted]

[36] Goldfarb DNP Opening Report ¶¶ 136-138.

[redacted] before ChatGPT's launch. The subscription [redacted] was underway by [redacted], before ChatGPT's release. Print subscriptions [redacted], while digital-only subscriptions [redacted]. Cancellation data identified [redacted] as stated reasons for subscriber cancellations.[37]

Figure 13: DNP Subscribers July 2020 to 2024



41. DNP's subscription mix [redacted] over the period analyzed, while total subscription revenue [redacted]. Total DNP subscriptions have been

[37] Goldfarb DNP Opening Report ¶¶ 142-156.

[redacted] the launch of ChatGPT.[38] Importantly, there is no [redacted] [redacted] the release of ChatGPT, suggesting that [redacted] [redacted] rather than one driven by the introduction of generative AI or OpenAI's LLMs.

Figure 14: DNP Subscription Revenue July 2020 to 2024

[redacted]

(3) DNP's web traffic datasets, methodology, and findings

42. Across the eight properties, visitors and sessions [redacted] during the COVID period and [redacted] By early 2022, both measures had [redacted]. The data

[38] Goldfarb DNP Opening Report ¶¶ 148-151.

therefore indicated that [redacted] could not be attributed to ChatGPT.[39]

Figure 15: Monthly count of visitors of all DNP properties[40]

[redacted]

Figure 16: Monthly session count of all DNP properties

[redacted]

[39] See Goldfarb DNP Opening Report ¶¶ 133-141, 161-165.

[40] DNP properties include the aggregated results for the eight Daily News Plaintiffs: Chicago Tribune, The New York Daily News, Denver Post, Orange County Register, Orlando Sentinel, San Jose Mercury News, Pioneer Press and Sun Sentinel.

Figure 17: Monthly visitor counts of each DNP property

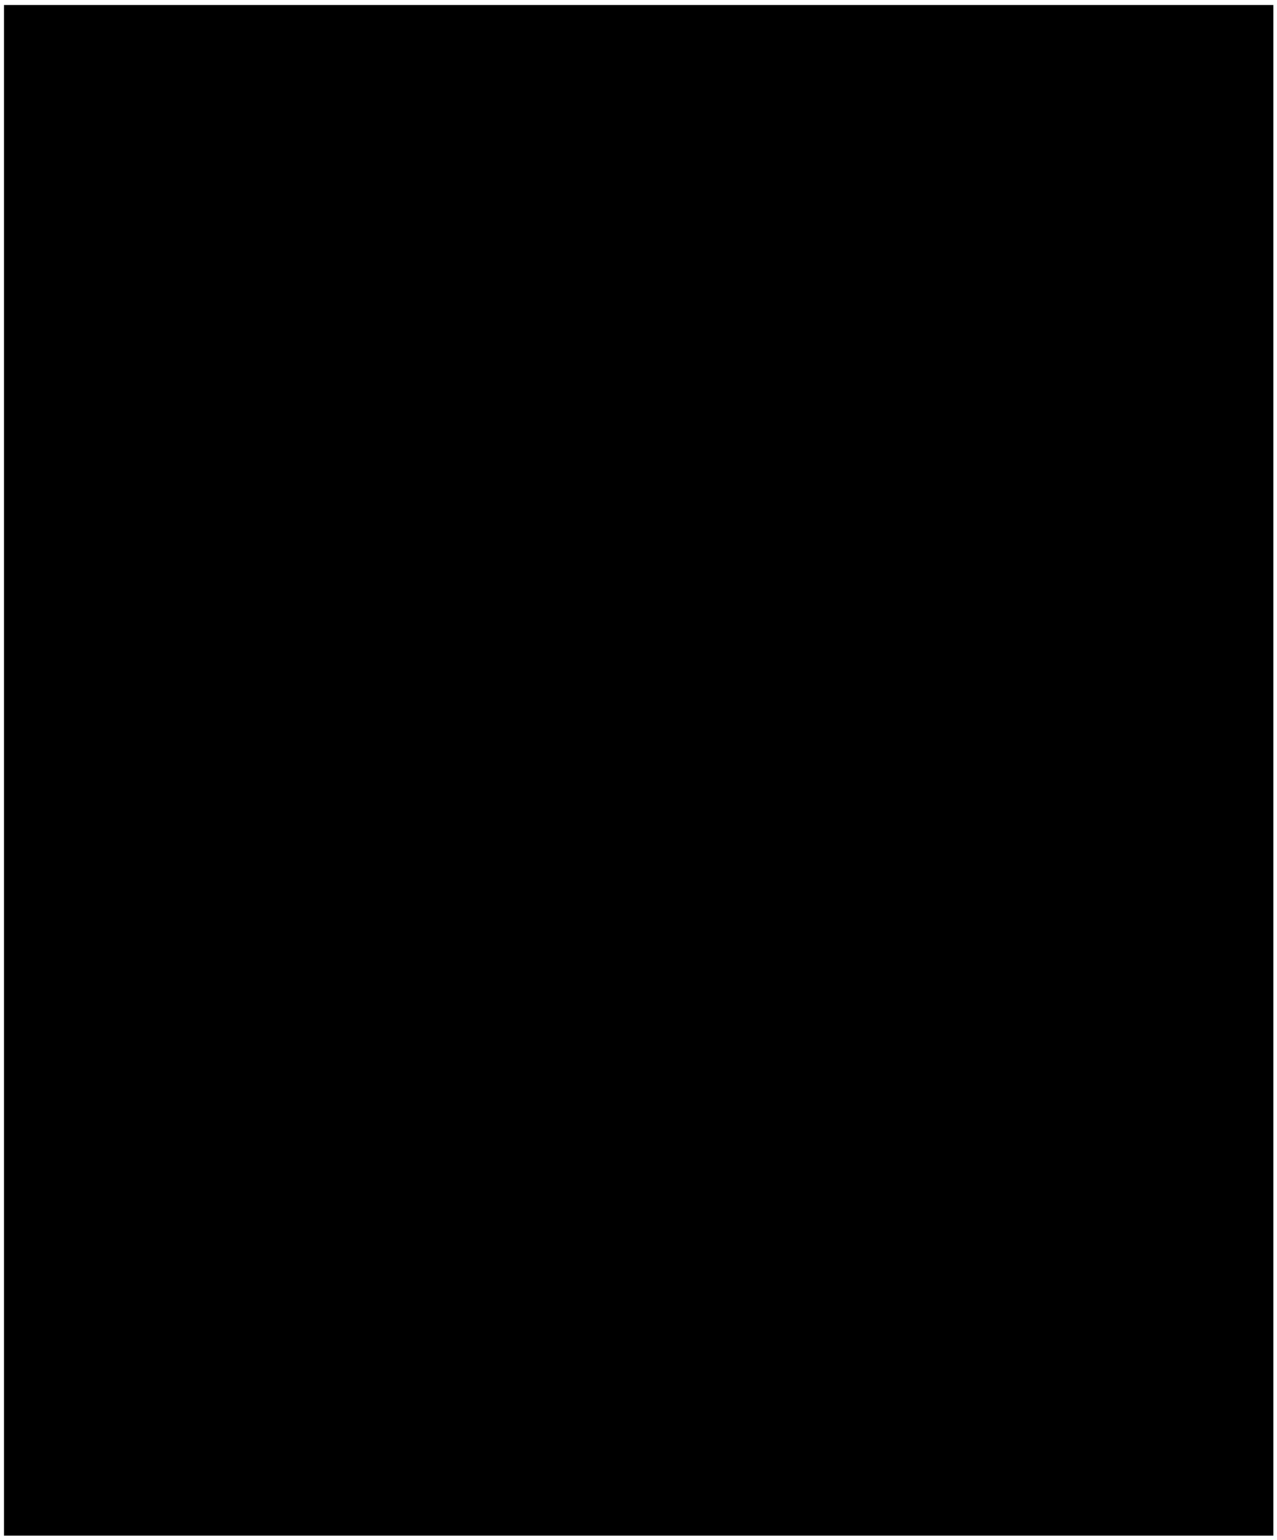
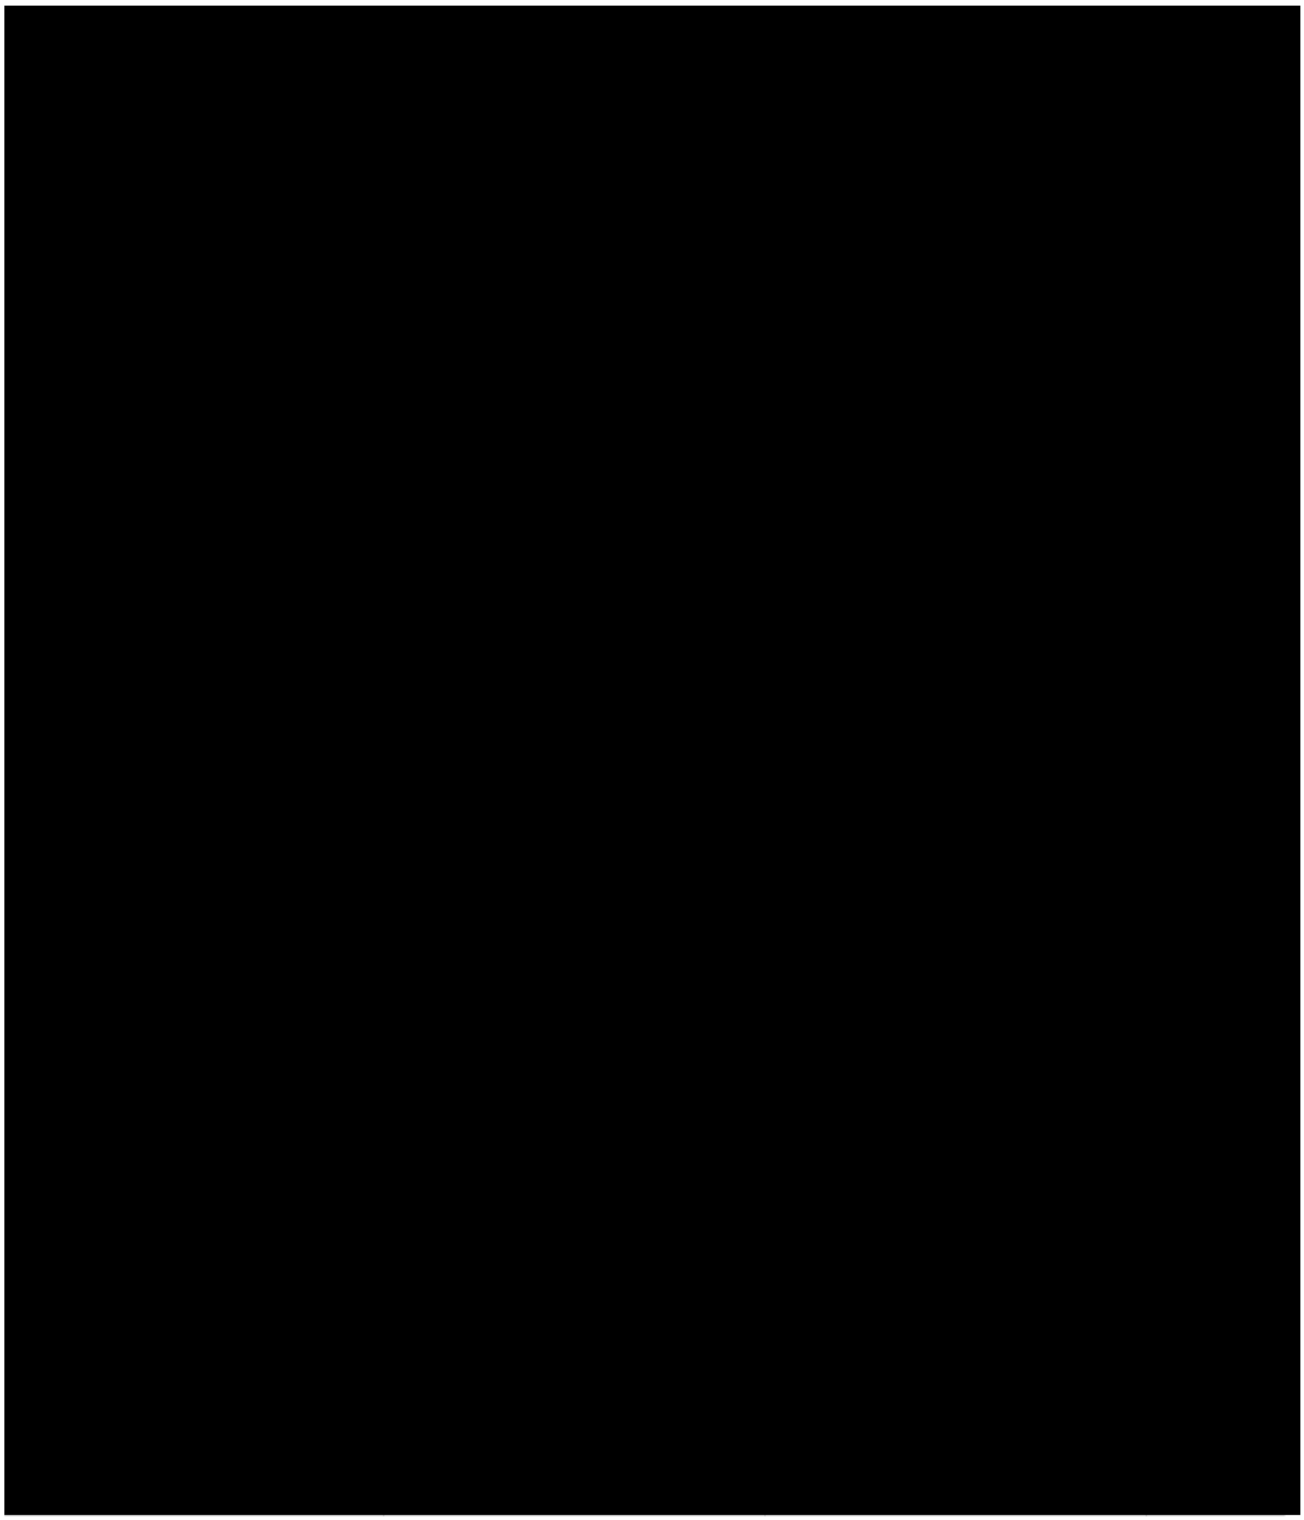


Figure 18: Monthly session counts of each DNP property



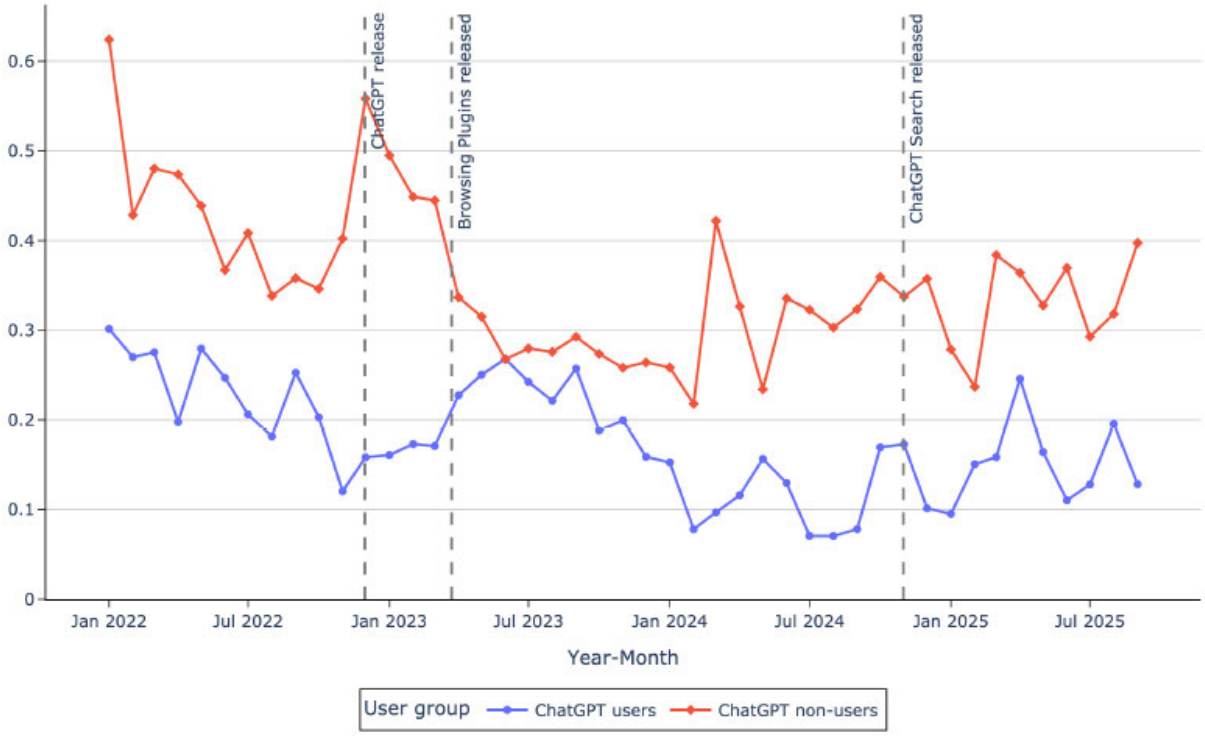
43. The DNP referral traffic from [redacted] and [redacted] before ChatGPT's release. The [redacted] demonstrates that there is no indication that the DNP referral traffic trends were negatively affected by ChatGPT or Bing Chat.[41]

(4) Comscore audience measurement and site-visitation analysis

44. I conducted the same Comscore analysis for the Daily News Plaintiffs as described above for The Times. The Daily News Plaintiffs' properties are analyzed jointly in the data. My principal DNP user-level analysis used Comscore desktop data, and I applied the same method to Comscore mobile data. I defined ChatGPT users as users who had visited the ChatGPT website after ChatGPT's release and used each user's first observed ChatGPT visit as the adoption date. For the econometric sample, I required at least six months before and after first ChatGPT use. As with The Times, the following analysis did not show a post-release divergence between the two groups.[42]

[41] Goldfarb DNP Opening Report ¶¶ 184, 188; Goldfarb NYT-DNP Rebuttal Report ¶¶ 110-111, 113.
[42] *See* Goldfarb DNP Opening Report ¶¶ 191-204, 355-356; Goldfarb NYT-DNP Reply Report ¶¶ 254-255.

Figure 19: Average weighted number of DNP desktop sessions per user and ChatGPT usage status

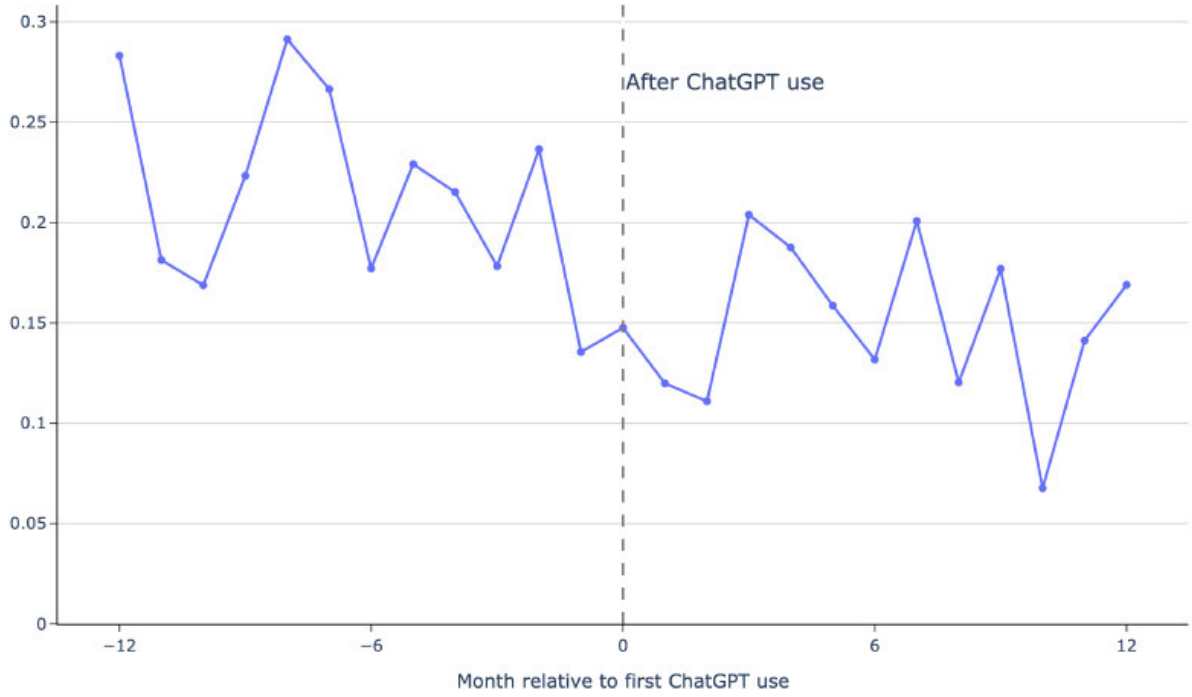


45. After ChatGPT's release, sessions for ChatGPT non-users generally remained higher. Figure 19 does not show a clear break at the time of ChatGPT's release in which ChatGPT users began to decline relative to non-users. Instead, DNP sessions for ChatGPT users increased shortly after the release while sessions for non-users decreased. If ChatGPT had substituted for DNP traffic, I would have expected the gap between the groups to widen starting around the release. Figure 19 does not show that divergence, which is inconsistent with a substitution effect.[43]

[43] *See* Goldfarb DNP Opening Report ¶¶ 195-201.

46. Figure 20 below shows DNP visits around each ChatGPT user's first observed ChatGPT use. When I aligned activity to each adopter's first use, users observed for at least six months before and after adoption showed a decline that was already underway before adoption and continued afterward at a slower rate. The data did not show a discrete post-adoption drop, which was inconsistent with a substitution effect. I therefore did not interpret the gradual continuation of the pre-existing decline as evidence that ChatGPT caused it.[44]

Figure 20: Average weighted DNP desktop sessions per month per ChatGPT user



(5) DiD and econometric methodology, results, and robustness checks

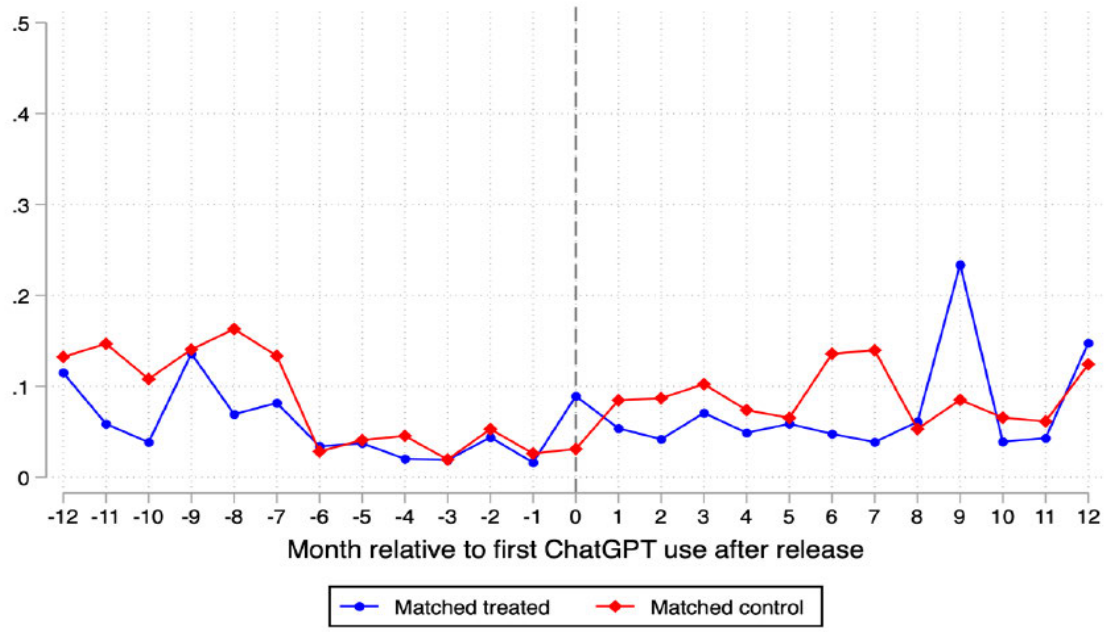
47. The principal DNP DiD analysis used the same before-and-after comparison as described above for The Times. I found no statistically significant reduction in DNP visits at any

[44] *See* Goldfarb DNP Opening Report ¶¶ 195-201.

point through one year after adoption, including periods after the release of ChatGPT Search. I also extended the analysis to 24 months after adoption for users who first used ChatGPT before October 2023; the post-adoption estimates were either positive or statistically indistinguishable from zero and did not indicate a negative effect.[45]

48. I also matched adopters (treated users) and non-adopters (control users) by cohort using demographic information and comparable pre-adoption visit patterns. The matched groups tracked one another before adoption and did not show a sustained post-adoption widening in which treated users fell below controls. Overall, Figure 21 provides no evidence that beginning to use ChatGPT reduced visits to DNP in this closely matched sample.[46]

Figure 21: Average number of DNP sessions per user



[45] *See* Goldfarb DNP Opening Report ¶¶ 214-216, 357-362; Goldfarb NYT-DNP Reply Report ¶¶ 235-236.
[46] Goldfarb DNP Opening Report ¶¶ 217-223, 363-369.

49. Robustness checks comparing users with different levels of ChatGPT visit intensity, using alternative specifications, mobile data, including users who did not visit DNP, excluding other LLM users, and using longer pre-adoption periods for matching likewise did not produce a negative DNP effect.[47]

iii. Ziff Davis

(1) Ziff Davis's Asserted Works

50. I understand that Ziff Davis has alleged that OpenAI used millions of copies of Ziff Davis's copyrighted works to train its LLMs, but that Ziff Davis did not produce the URLs associated with its Asserted Works list or article-level traffic data. It produced a list of a small number of individual articles and identified 20 properties. I also understand that the copyrighted works at issue in this case fall within three categories: (1) Technology & Shopping, (2) Gaming & Entertainment, and (3) Health & Wellness.[48]

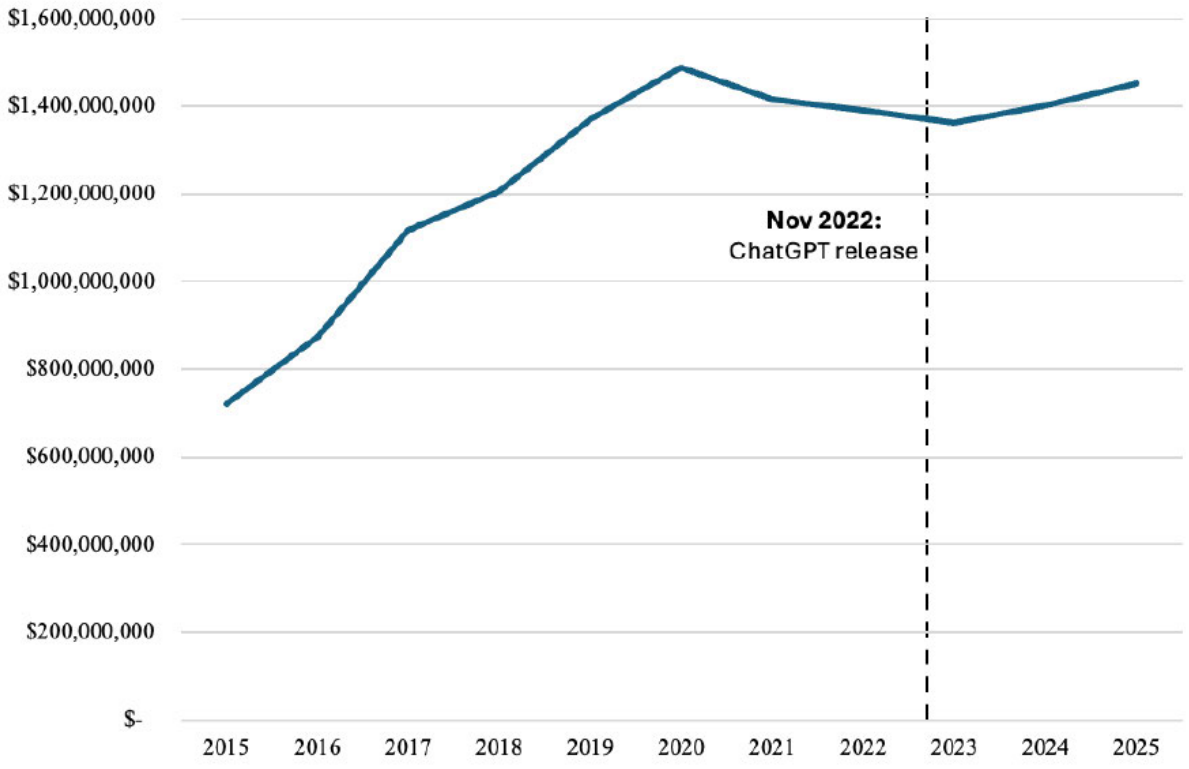
51. I could not perform an article-level traffic-decay analysis for Ziff Davis with the available produced Ziff Davis data. Without the Asserted Works URLs and article-level traffic data, I could not calculate publication-date distributions or match individual Asserted Works to pageviews to estimate cumulative post-publication traffic decay. I used Ziff Davis's produced property-level data to analyze property-level traffic as detailed below but was unable to analyze the age or lifetime traffic of individual Asserted Works as I did with other Plaintiffs.

[47] *See* Goldfarb NYT-DNP Reply Report ¶¶ 220-229, 253-261.
[48] Goldfarb ZD Opening Report ¶¶ 22, 25, 109, 110.

(2) Financial Performance

52. Ziff Davis's overall revenue does not show a negative impact from ChatGPT; rather, revenue grew substantially after 2015 and has remained stable or increased in the years following ChatGPT's release. Ziff Davis displayed strong revenue growth from approximately $720 million in 2015 to nearly $1.5 billion in 2020. Ziff Davis's overall revenue then remained at approximately $1.4 billion. In 2025, Ziff Davis reported approximately $1.45 billion in revenue, a 3.5% increase over 2024, and nearly $290 million in free cash flow. Four of its five segments reported revenue growth in the fourth quarter of 2025.[49]

Figure 22: Ziff Davis's Revenue 2015-2025 (not adjusted for inflation)



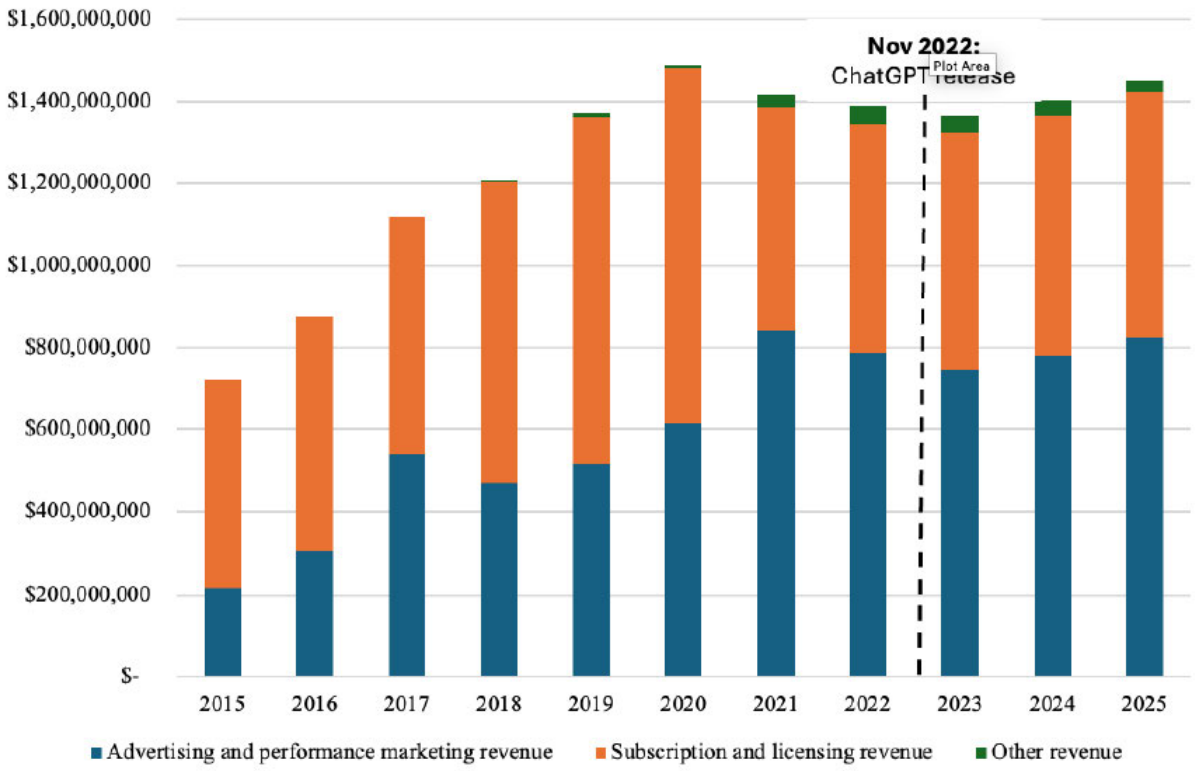
[49] See Goldfarb ZD Opening Report ¶¶ 81, 115.

53. Because Ziff Davis completed acquisitions and divestitures during the period, I also examined its asset-turnover ratio, calculated as total revenue divided by total assets. That measure provided context for consolidated revenue because acquired businesses were reflected in consolidated assets; in years with large acquisitions, the ratio could decline temporarily before the acquired businesses contributed a full year of revenue. Despite these acquisitions, the ratio showed no decline correlated with ChatGPT's release. Ziff Davis completed approximately $415 million in acquisitions in 2019, $141.1 million in 2021, and $218 million in 2024. The asset-turnover analysis showed no observable decline after ChatGPT's launch and corroborated my conclusion that Ziff Davis's business performance did not show harm from OpenAI's products.[50]

54. As shown in the figure below, Ziff Davis's advertising and performance-marketing revenue did not show a sustained post-ChatGPT decrease, and its subscription-and-licensing revenue remained stable. I therefore did not find a negative impact on either major revenue category.[51]

[50] *See* Goldfarb ZD Reply Report ¶¶ 184-188.
[51] *See* Goldfarb ZD Opening Report ¶¶ 114-116.

Figure 23: Ziff Davis's Revenue by Source 2015-2025 (not adjusted for inflation)



(3) Web traffic datasets, methodology, and findings

55. Ziff Davis produced traffic data for a subset of its web properties and a separate IGN App property, but only BabyCenter and What to Expect had data extending before ChatGPT's release. I therefore used those two properties for the comparison of periods before and after ChatGPT's release. As discussed below, both properties showed [redacted] [redacted] ChatGPT's release.[52]

[52] Goldfarb ZD Opening Report ¶¶ 119-130, 315-316.

56. BabyCenter and What to Expect showed [redacted] ChatGPT's release. BabyCenter's [redacted] was concentrated among [redacted] sessions, while What to Expect [redacted]; the [redacted] for What to Expect was [redacted].[53]

57. I examined the monthly session counts for BabyCenter and What to Expect. My review of the available Ziff Davis internal data showed no evidence of an impact on traffic resulting from the launch of ChatGPT. Rather, traffic to these Ziff Davis properties had been [redacted] the launch of ChatGPT.[54]

Figure 24: Monthly session counts for BabyCenter

[redacted]

[53] Goldfarb ZD Opening Report ¶¶ 119-130.
[54] Goldfarb ZD Opening Report ¶¶ 122-123.

Figure 25: Monthly session counts for What to Expect



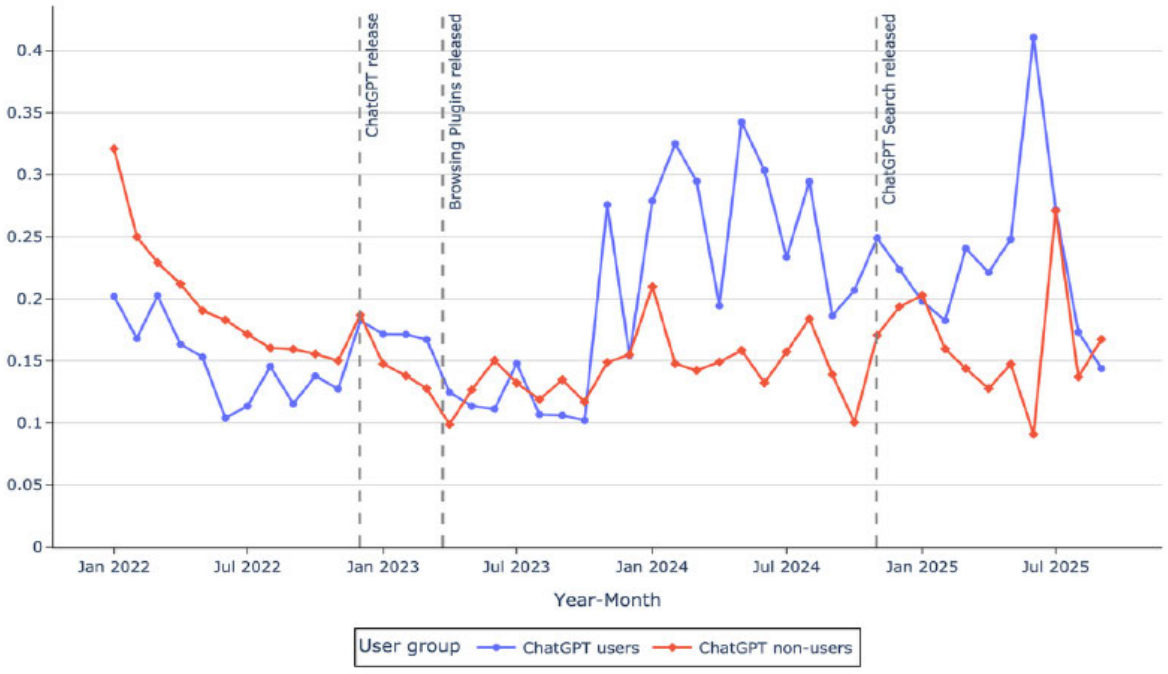
58. Ziff Davis's own data and public statements likewise show that search referrals from Bing.com [redacted] rather than [redacted] after Bing implemented Generative AI into its search experience. Ziff Davis also reported that organic referral traffic from Bing [redacted] year over year from [redacted], while Bing's total visits [redacted] I also reviewed Ziff Davis's public statements describing generative search as having a very positive impact on traffic referrals and that there was no reported traffic dislocation or negative effects from generative AI in 2024. These statements concerned Generative AI in Microsoft's Bing products, and not solely ChatGPT, but they provide context for evaluating Ziff Davis's generalized claims about AI-related traffic harm.[55]

[55] Goldfarb ZD Rebuttal Report ¶ 90; Goldfarb ZD Opening Report ¶¶ 178-180.

(4) Comscore audience measurement and site-visitation analysis

59. Using Comscore's desktop traffic data, I compared average weighted sessions for ChatGPT users and non-users by Ziff Davis's Technology & Shopping, Gaming & Entertainment, and Health & Wellness lines of business, and by vertical: news, review-and-recommendation, and forum-like properties. The figures below do not show a clear break at the time of ChatGPT's release in which ChatGPT users began to decline relative to non-users. If ChatGPT were substituting for these properties, I would have expected a divergence between the two groups of users starting around ChatGPT's release. None of the segmentations showed a differential post-release decline—that is, a decline for ChatGPT users relative to non-users.[56]

Figure 26: Average weighted number of Technology & Shopping sessions per relevant user and ChatGPT usage



[56] Goldfarb ZD Opening Report ¶¶ 139-145.

Figure 27: Average weighted number of Gaming & Entertainment sessions per relevant user and ChatGPT usage

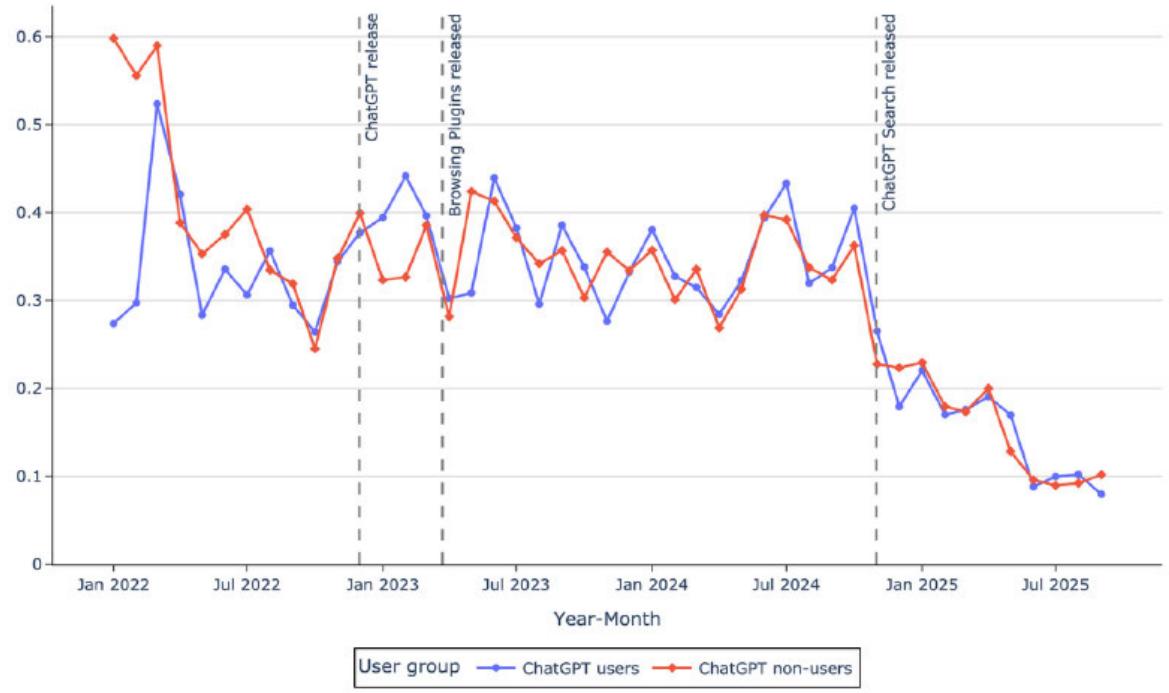


Figure 28: Average weighted number of Health & Wellness sessions per relevant user and ChatGPT usage

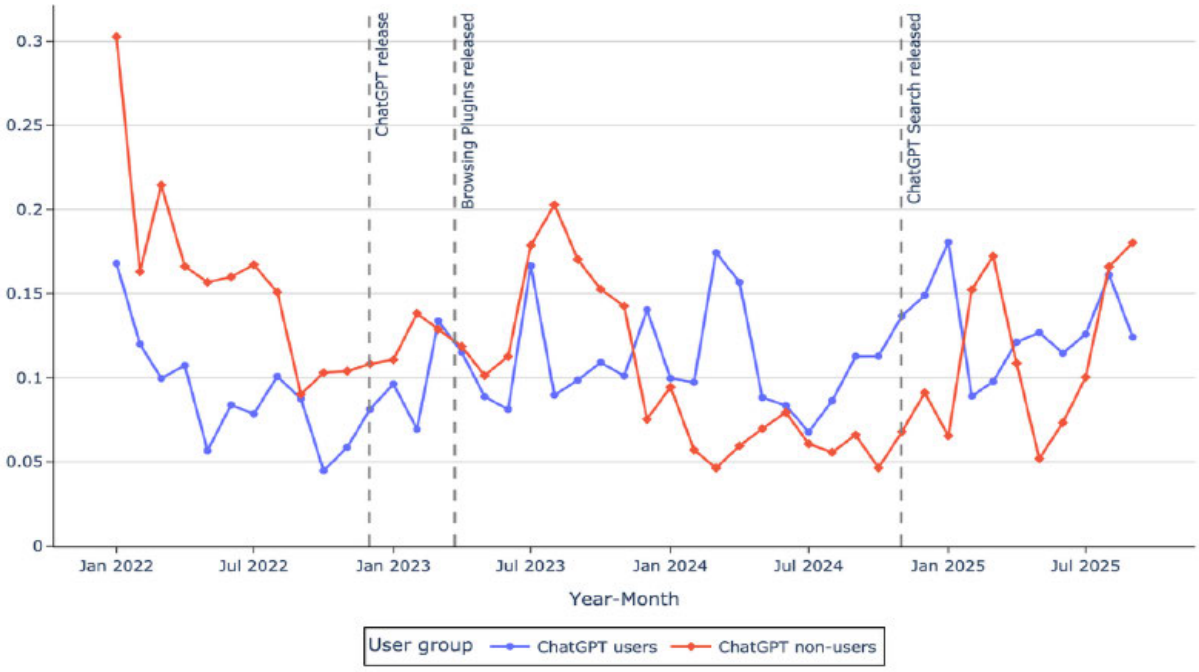


Figure 29: Average weighted number of sessions for vertical news properties per relevant user and ChatGPT usage

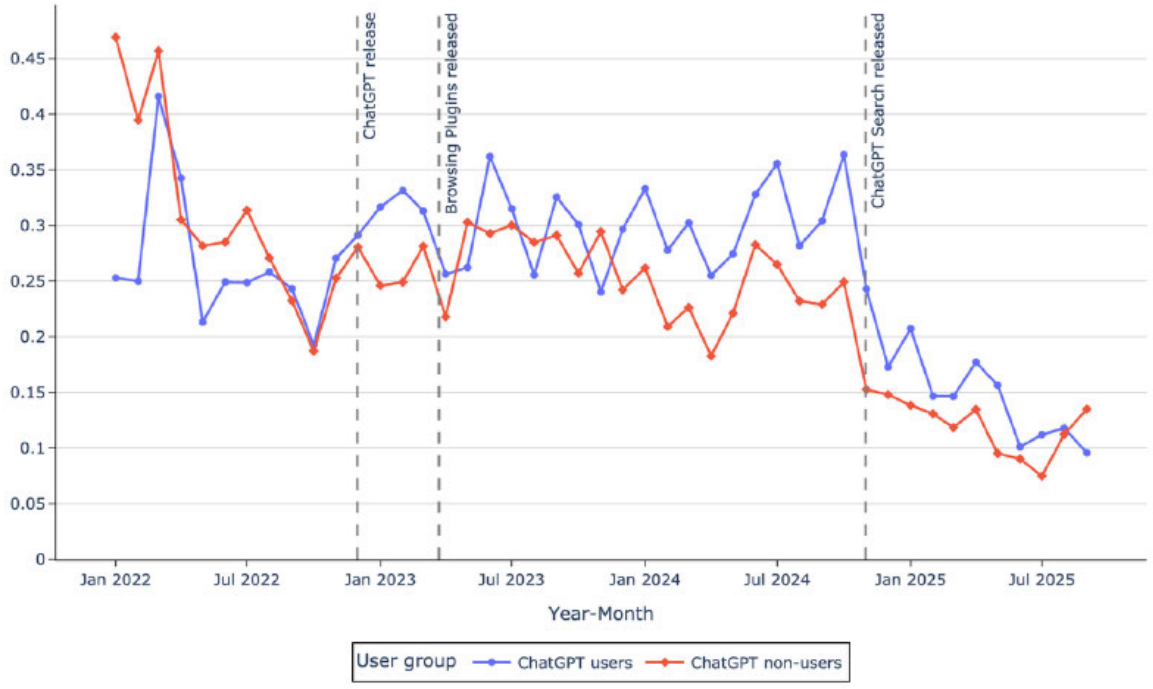


Figure 30: Average weighted number of sessions for reviews & recommendations properties per relevant user and ChatGPT usage

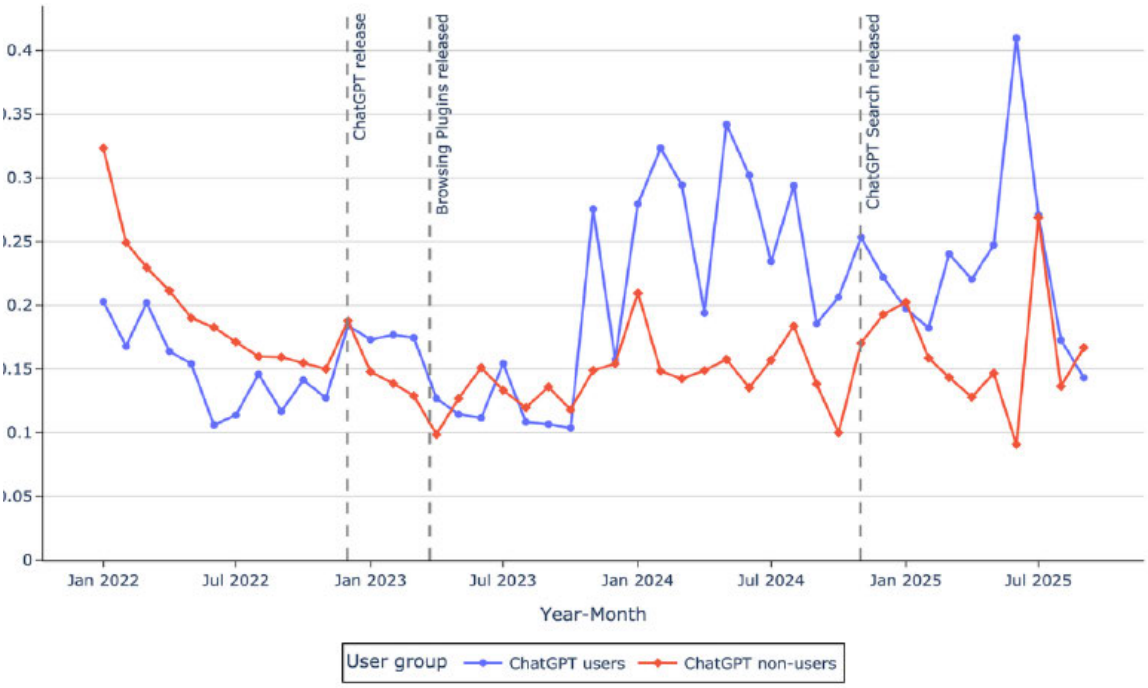
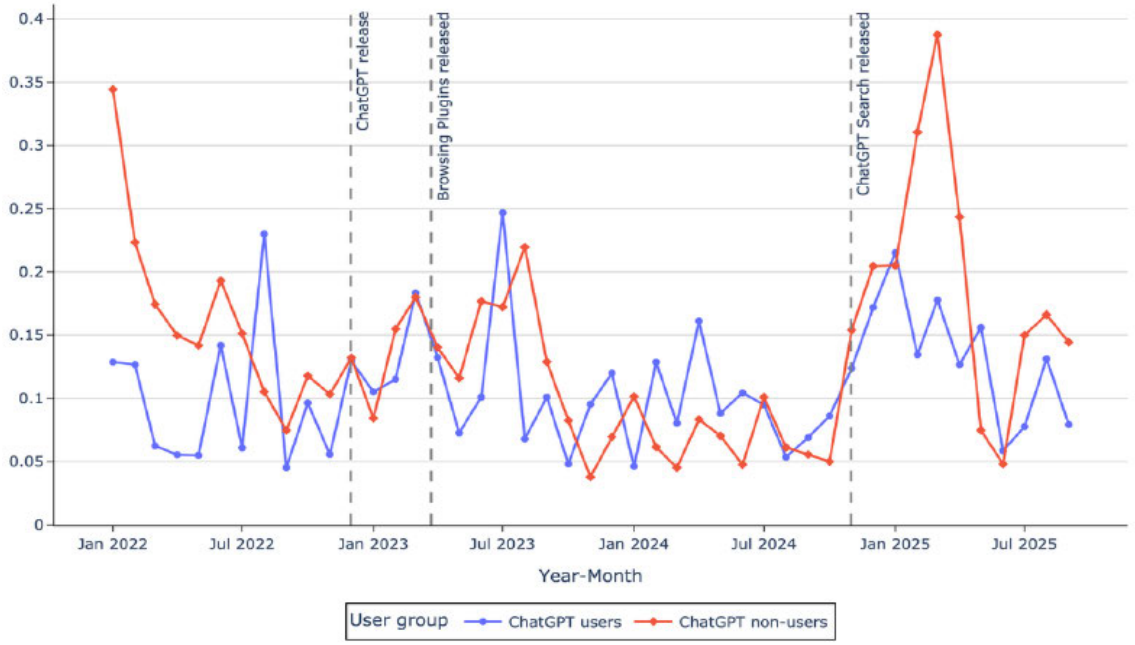


Figure 31: Average weighted number of sessions for forum-like properties per relevant user and ChatGPT usage

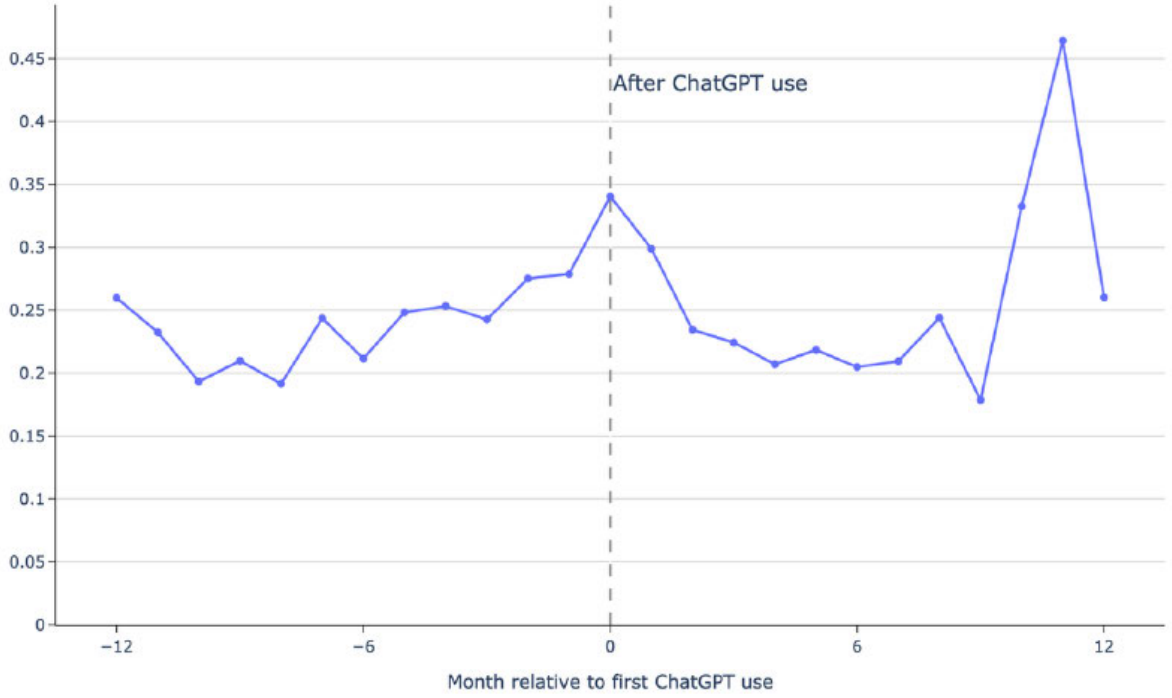


60. In sum, the data consistently showed that ChatGPT's release did not affect traffic to Ziff Davis properties. This conclusion held regardless of how the data was segmented: whether by line of business or property type.

61. A separate first-use event series—a month-by-month record around first use—for 8,151 users likewise showed no statistically significant post-adoption decline. If ChatGPT had substituted for Ziff Davis's traffic, I would have expected traffic to Ziff Davis to be lower after a user's first ChatGPT use than it was in the period immediately before ChatGPT adoption. The pattern shown in the figure below does not display such a decline.[57]

[57] *See* Goldfarb ZD Opening Report ¶¶ 146-149.

Figure 32: Average weighted number of Ziff Davis sessions per relevant ChatGPT user



(5) Difference-in-differences (DiD), results, and robustness checks

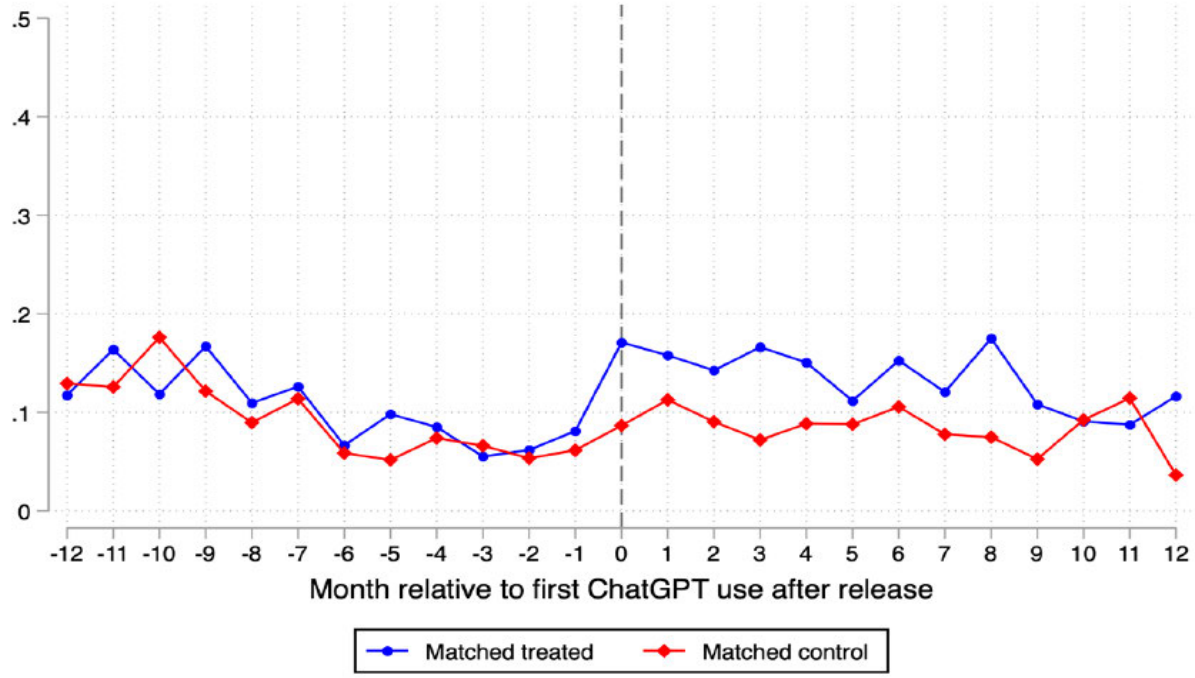
62. The principal Ziff Davis DiD analysis compared monthly sessions before and after each user's first observed ChatGPT use, with people who did not use ChatGPT serving as the contemporaneous control group. It showed no statistically significant reduction through one year after adoption.[58]

63. In a separate matching analysis, I matched each adopter to a non-user within the adopter's monthly cohort using demographic information and comparable pre-adoption visit patterns. After adoption, ChatGPT users more often met or exceeded the comparison non-user

[58] *See* Goldfarb ZD Opening Report ¶¶ 151-166, 319-324.

group's session levels, and some estimates were statistically significant. That pattern supported my conclusion that Ziff Davis sessions did not decrease relative to matched non-users.[59]

Figure 33: Average number of Ziff Davis sessions per user



64. Robustness analyses across usage quartiles and the top 10% of the most intensive ChatGPT users, using alternative specifications, excluding users of Copilot, Gemini, Claude, and Perplexity, and expanding the analysis to include Comscore users who did not visit Ziff Davis likewise did not show a sustained negative effect or a statistically significant reduction in Ziff Davis sessions.[60]

[59] Goldfarb ZD Opening Report ¶¶ 168-172, 324-330.
[60] *See* Goldfarb ZD Opening Report ¶¶ 167-172, 325-330; Goldfarb ZD Reply Report ¶¶ 164-183.

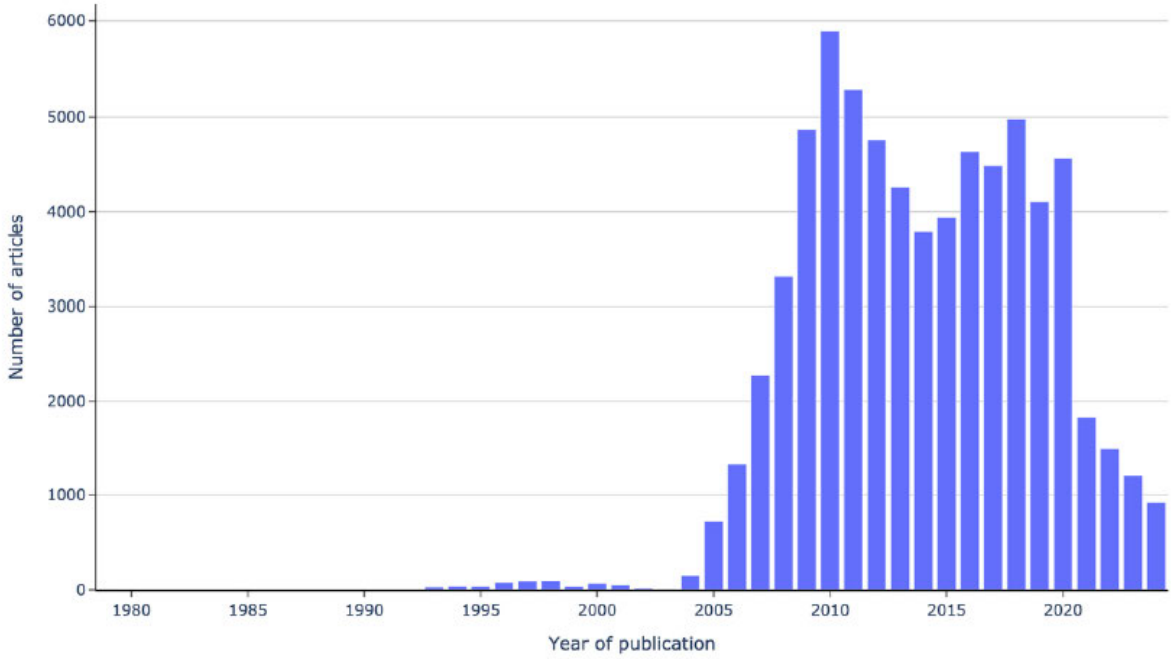
iv. Center for Investigative Reporting (Mother Jones and Reveal)

(1) CIR's Asserted Works

65. For CIR, I analyzed the Mother Jones and Reveal individual Asserted Works separately. Mother Jones's list contains 69,300 unique URLs. Reveal's list contains 3,680 URLs, of which 3,462 are unique, and I understand that CIR has not asserted copyright infringement claims for any Reveal Asserted Works.[61]

66. Mother Jones's individual Asserted Works ranged from April 1979 through September 2024, and 92% were published in 2020 or earlier, as demonstrated in the figure below.[62]

Figure 34: Mother Jones's individual Asserted Works by publication month



[61] Goldfarb CIR Opening Report ¶¶ 29-30, 116-117.
[62] Goldfarb CIR Opening Report ¶ 118.

67. For Mother Jones's individual Asserted Works, I mapped 40,797 of the 40,809 works published during the August 2012–July 2023 period to the monthly Universal Analytics pageview data, along with approximately 5,000 works published before August 2012. Because CIR did not produce data that directly isolated traffic to the Asserted Works, the matches were made to page-level URLs, and the Universal Analytics data reported only the top 5,000 pages by pageviews each month, including some non-article pages. This matched work showed that traffic to Mother Jones's Asserted Works was [redacted] For works published from 2013 through 2021, more than [redacted] of cumulative pageviews were recorded by approximately [redacted] after publication, and the [redacted].[63]

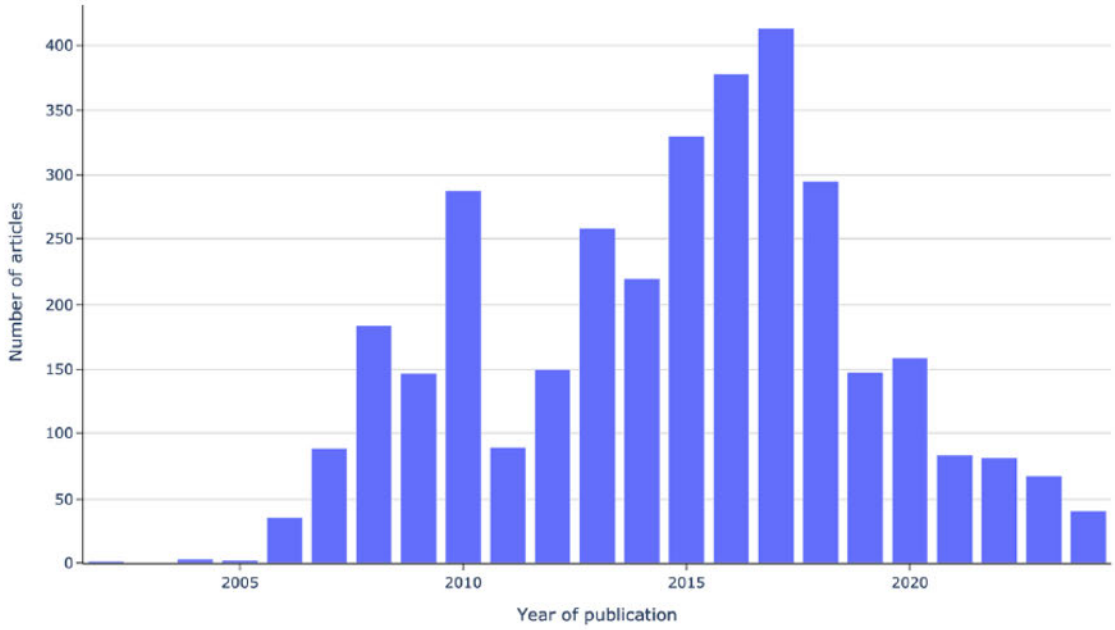
Figure 35: Cumulative share of pageviews of Mother Jones's individual Asserted Works over time



[63] Goldfarb CIR Opening Report ¶¶ 121–122.

68. Reveal's Asserted Works ranged from October 2002 through September 2024, and 92% were published in 2020 or earlier.[64]

Figure 36: Reveal's individual Asserted Works by publication month



69. I was able to map pageview data for 1,834 works published by Reveal between April 2015 and July 2023. This mapping showed a similar pattern of [redacted] for Reveal's Asserted Works. For Reveal's Asserted Works published from 2016 through 2021, more than [redacted] of cumulative pageviews were recorded within approximately [redacted] publication, indicating [redacted].[65]

[64] Goldfarb CIR Opening Report ¶ 119.
[65] Goldfarb CIR Opening Report ¶¶ 119, 124-125.

Figure 37: Cumulative share of pageviews of Reveal's Asserted Works over time



70. Taken together, these analyses show that the older Mother Jones and Reveal works had [redacted] available to be displaced. There is necessarily a time gap between acquiring an article for training and releasing a model trained on it; as a result, most of the traffic value of these works had been realized before ChatGPT's release or before a model trained on an article could be released.[66]

(2) Financial Performance

71. Overall, CIR's revenue performance was [redacted] the release of ChatGPT. After the 2024 merger of Mother Jones and Reveal, CIR's reported revenue [redacted]

[66] Goldfarb CIR Opening Report ¶ 126.

[redacted] because the post-merger figures combined the revenues of both publishers.[67]

72. CIR's post-merger financial results reflect continued growth across its combined operations. After the merger, CIR reported approximately $23.1 million in revenue for February–December 2024 and approximately $23.6 million for January–October 2025. Its December 2025 materials projected [redacted]. CIR also recorded [redacted], which CIR attributed to [redacted].[68]

Figure 38: CIR Total Revenue 2018-2025

[redacted]

[67] Goldfarb CIR Opening Report ¶¶ 129-130.
[68] Goldfarb CIR Opening Report ¶¶ 129–136.

Figure 39: Reveal Revenue 2018-2023



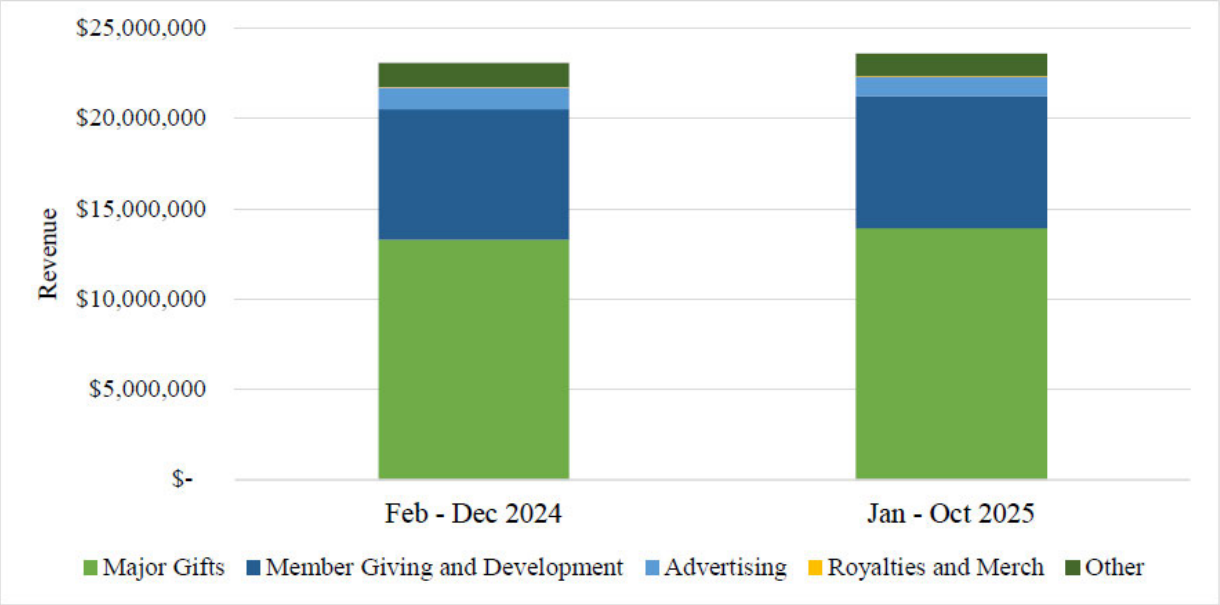
73. Mother Jones's subscription and member-development revenue [redacted] ChatGPT's release. After hovering at approximately [redacted] around the release of ChatGPT, Mother Jones's subscription revenue from member circulation [redacted] from January to October 2025, while member-development revenue [redacted]. From 2018 through 2023, [redacted] of Mother Jones's annual revenue, while [redacted] accounted for [redacted] and that share [redacted].[69]

[69] Goldfarb CIR Opening Report ¶¶ 132–138.

Figure 40: Mother Jones Revenue 2018-2023

[redacted]

Figure 41: CIR Revenue



(3) Web traffic datasets, methodology, and findings

74. For CIR, I analyzed data produced by CIR for Mother Jones and Reveal. CIR produced aggregate and page-level Universal Analytics and Google Analytics data. I analyzed

the monthly Universal Analytics data and the daily Google Analytics series separately because they differed in frequency and their totals were not directly comparable. This data allowed me to evaluate users, pageviews, sessions, referral sources, and traffic to individual Asserted Works. Depending on the metric and dataset, the available series for pageviews ran from August 2012 through November 2025 for Mother Jones and from April 2015 through November 2025 for Reveal. The resulting analysis, discussed below, is [redacted] unrelated to ChatGPT.[70]

75. Across the aggregate and page-level traffic data, I found [redacted] ChatGPT's release or at later updates. Mother Jones's earlier traffic data showed a [redacted], followed by [redacted]. Reveal's [redacted], while its post-ChatGPT-release traffic data were [redacted]. These patterns are not consistent with a new, persistent traffic decline attributable to the release of ChatGPT or its updates.[71]

[70] *See* Goldfarb CIR Opening Report ¶¶ 139-145, 305-318.
[71] Goldfarb CIR Opening Report ¶¶ 149-167, 326-333.

Figure 42: Mother Jones Monthly Total Pageviews (August 2012-November 2025)

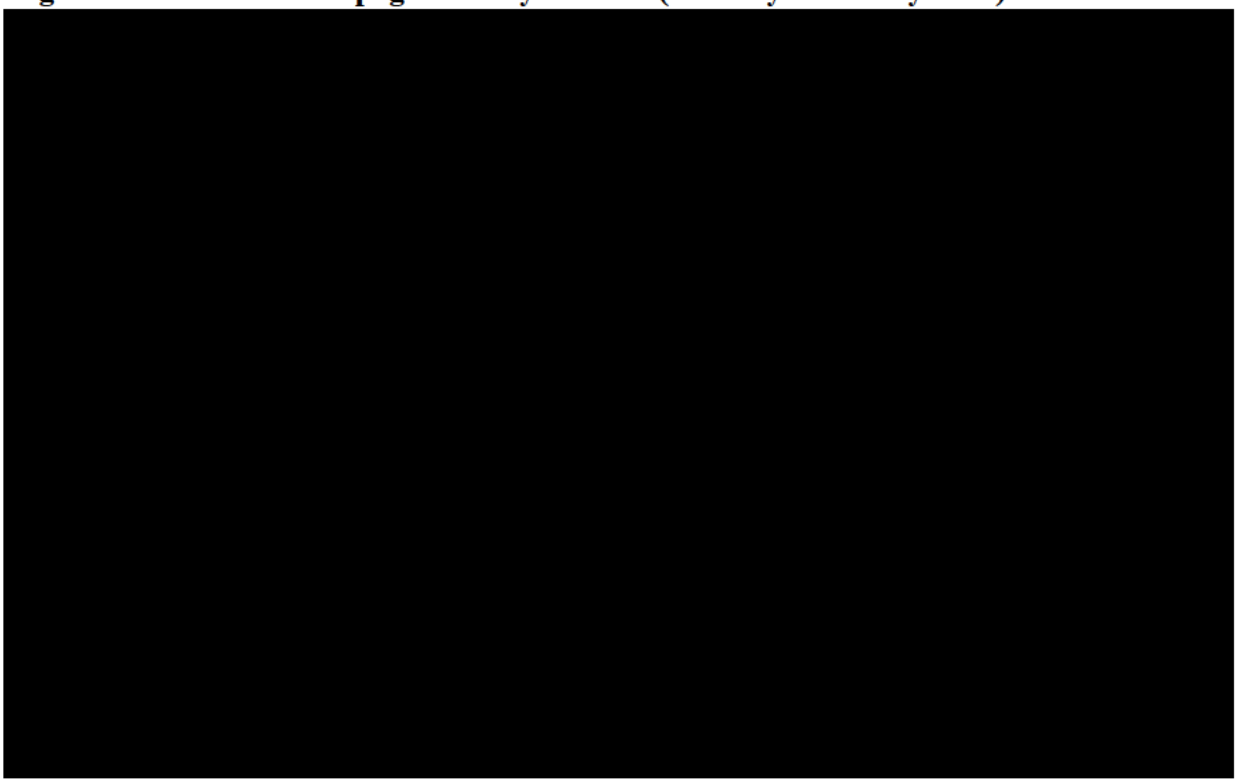


Figure 43: Reveal Monthly Total Pageviews (April 2015-July 2023)



76. The same conclusion appeared when I separated traffic by source. Direct, search, and social referral traffic [redacted] ChatGPT, and none showed [redacted] ChatGPT's launch. The patterns illustrated by the figures below are consistent with longer-term shifts in audience and referral dynamics that were unrelated to ChatGPT rather than with a new substitution effect caused by ChatGPT.[72]

Figure 44: Mother Jones pageviews by channel (January 2018-July 2023)



[72] Goldfarb CIR Opening Report ¶¶ 168-183.

Figure 45: Reveal pageviews by channel (January 2018-July 2023)



77. For CIR, the Mother Jones and Reveal online publications do not have a sufficiently large number of visitors for Comscore's sample of such visitors to be of an appropriate size to conduct the kind of statistical analysis that I performed for The Times, DNP, and Ziff Davis.

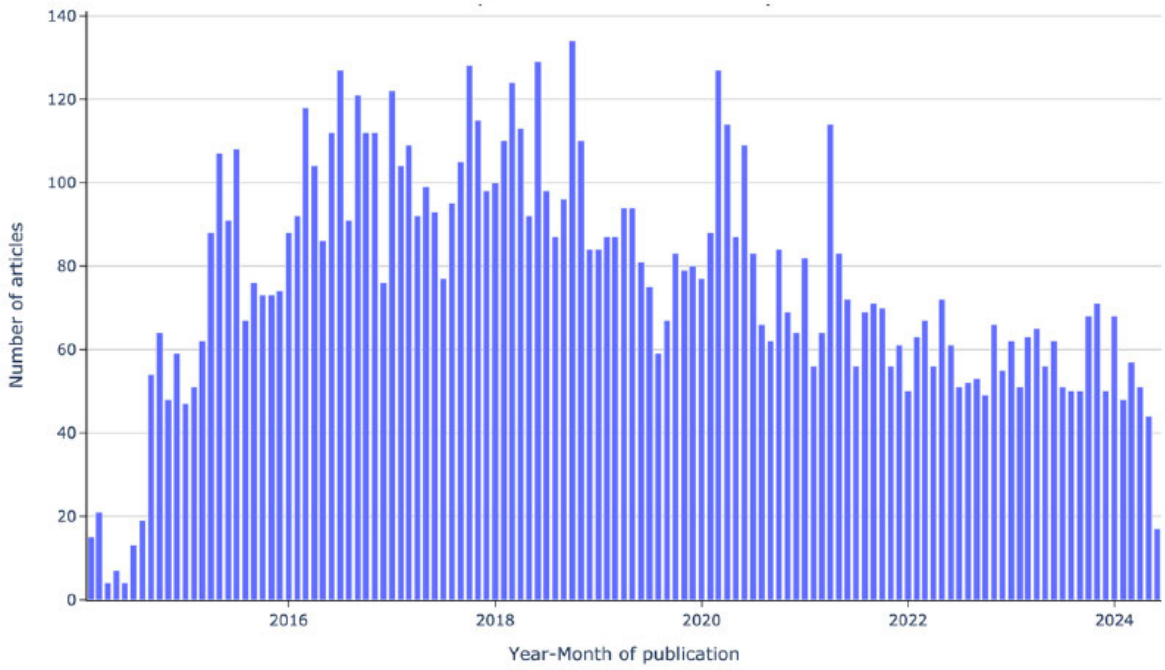
v. The Intercept

(1) The Intercept's Asserted Works

78. The Intercept's Asserted Works list contained 9,511 works with publication dates ranging from February 2014 to June 2024. As with the other News Plaintiffs, most of these works predate OpenAI's LLMs. 11.3% of the articles were published before OpenAI was

founded in November 2015, 88.4% before ChatGPT was released in November 2022, and 99.4% before GPT-4o was released in May 2024.[73]

Figure 46: Number of The Intercept Asserted Works by publication date



79. For The Intercept, I used Parse.ly data produced by The Intercept covering January 2017 through October 2025 and matched the Asserted Works to pageviews for a traffic decay analysis. This analysis showed that most reader engagement with The Intercept's Asserted Works [redacted] For the works I analyzed, the share of cumulative pageviews recorded in the month of publication [redacted].[74]

[73] Goldfarb Intercept Opening Report ¶¶ 32-34; 161-163.
[74] Goldfarb Intercept Opening Report ¶¶ 136-137, 161-162.

Figure 47: Cumulative Share of Pageviews of The Intercept Asserted Works over Month



80. Because 88.4% of The Intercept's Asserted Works were published at least one month before ChatGPT's release, the number of Asserted Works with meaningful traffic potentially affected by a later ChatGPT output was sharply limited. The traffic pattern therefore indicates that most engagement was realized near publication and does not support a ChatGPT-related reduction in traffic to the Asserted Works.[75]

(2) Financial Performance

81. The Intercept derives the overwhelming majority of its revenue from contributions and grants. Although web traffic [redacted] internal donation data showed that individual member contributions [redacted] the ChatGPT launch [redacted]

[75] Goldfarb Intercept Opening Report ¶¶ 161-164.

[redacted], showing that [redacted] The Intercept's financial performance.[76]

82. Although web traffic to The Intercept's Asserted Works [redacted], internal donation data shows that [redacted] over the same period. As shown below, individual contributions by members [redacted] launch of ChatGPT [redacted], despite the [redacted].[77]

Figure 48: The Intercept contributions 2018-2025

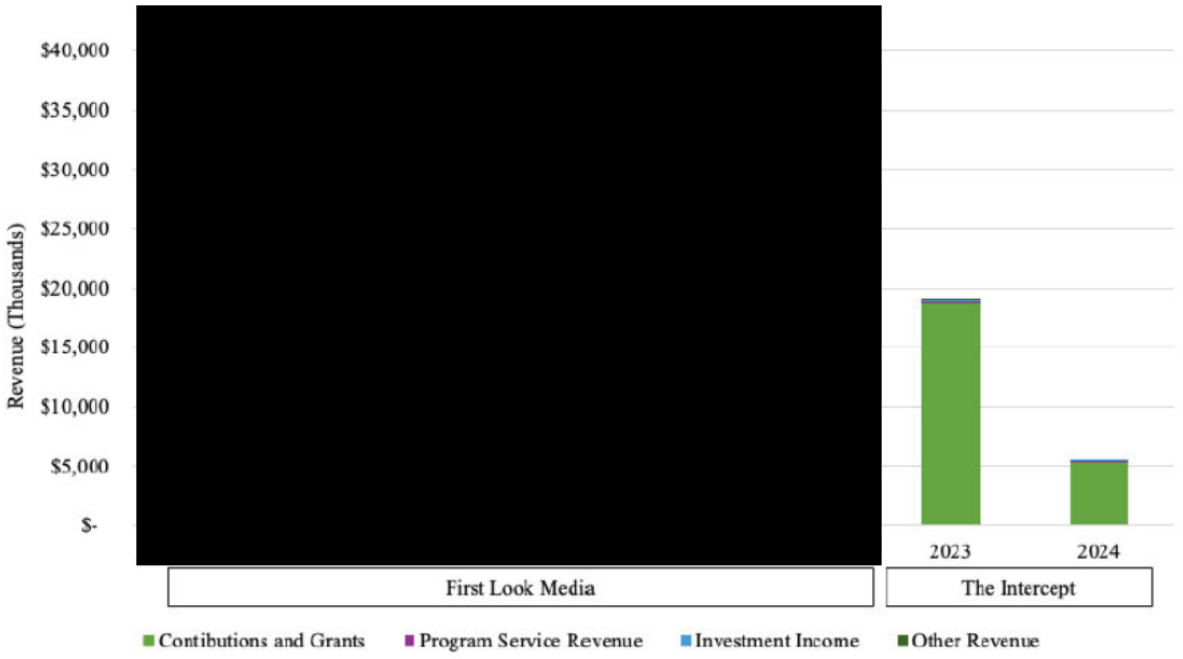
[redacted]

83. The Intercept's total revenue [redacted] since 2018, but [redacted] reflects organizational changes, including its 2023 spin-off from the First Look Institute, rather

[76] Goldfarb Intercept Opening Report ¶¶ 123-134.
[77] Goldfarb Intercept Opening Report ¶¶ 123-134.

than any effect of ChatGPT. Since 2018, The Intercept's total revenue has [redacted], from approximately [redacted], when it operated under the First Look Institute umbrella, to about $19 million in 2023, the year of its spin-off, and further to roughly $6 million in its first year of independent operation. That [redacted] was punctuated by [redacted] pre-dated ChatGPT's launch, and was not shown by the data to have been induced by ChatGPT.[78]

Figure 49: The Intercept revenue 2018-2024



84. The Intercept's funding model was relatively insulated from short-term traffic fluctuations: more than [redacted] of its revenue in the period analyzed came from contributions [redacted], there was [redacted] associated with the Asserted Works, and its [redacted] was followed by [redacted]

[78] Goldfarb Intercept Opening Report ¶¶ 123-131.
[79] Goldfarb Intercept Opening Report ¶¶ 130-134.

(3) Web traffic datasets, methodology, and findings

85. For The Intercept, I used data produced by The Intercept from Parse.ly, a web-analytics service, covering January 2017 through October 2025, and focused on pageviews. Because The Intercept's revenue was primarily from [redacted] rather than [redacted], I treated pageviews as an indicator of demand and engagement, not as a direct measure of financial return. The pageview and publication-volume data discussed below reflect a gradual decline that pre-dates ChatGPT.[80]

86. The Intercept's website traffic showed a [redacted], and publication volume [redacted][81]

Figure 50: The Intercept's number of published Asserted Works and pageviews per month

[redacted]

[80] *See* Goldfarb Intercept Opening Report ¶¶ 108-110, 122, 136-137.
[81] Goldfarb Intercept Opening Report ¶¶ 135-159, 162-164.

(4) Aggregate analytics and site-visitation analysis

87. As demonstrated below, the pageview analysis showed [redacted] to the Asserted Works after the launch of ChatGPT. The [redacted] The Intercept's Asserted Works across [redacted]. My analysis attributed those [redacted] to factors unrelated to the launch of ChatGPT.

Figure 51: Monthly Pageview Counts by Traffic Type for The Intercept's Asserted Works



88. Like CIR, The Intercept's online publication did not have a sufficiently large number of visitors for Comscore's sample of such visitors to then be of an appropriate size to

conduct the kind of statistical analysis that I performed for The Times, DNP, and Ziff Davis. Therefore, I used aggregate pageview and referral-source data to assess traffic and engagement.[82]

F. Other Causes and Failure to Isolate OpenAI's Effect

89. My reports identified several alternative causes of traffic and revenue changes that predated or were independent of ChatGPT. For The Times and DNP, the Plaintiffs themselves attributed traffic changes to [redacted] Each of these factors has independently affected web traffic. Ziff Davis's data likewise pointed to [redacted]

90. Several Plaintiffs' experts offered opinions about traffic, substitution, licensing, and market harm. Ms. Stricchiola compared Bing Chat and Bing Search click-through rates. Mr. Malackowski relied on that comparison and broader assertions about AI substitution and licensing. Professors Bell and Owen relied on industry studies, surveys, and anecdotes concerning AI-generated news and search products such as Google's. Dr. McDonald estimated

[82] *See* Goldfarb Intercept Opening Report ¶¶ 136-160, 284-291.
[83] Goldfarb NYT Opening Report ¶¶ 167-209; Goldfarb DNP Opening Report ¶¶ 166-190; Goldfarb ZD Opening Report ¶¶ 123-134; Goldfarb CIR Opening Report ¶¶ 129-146, 168-183; Goldfarb Intercept Opening Report ¶¶ 123-160.

time-trends for The Intercept's and Mother Jones's web traffic using regressions. Dr. Waehrer estimated a correlation between ChatGPT growth and Ziff Davis's search traffic. I evaluated each expert's method using available traffic, financial, and user-level data. The material limitation across these opinions was twofold: they did not isolate a change in demand for the Asserted Works and they did not isolate the harm caused by OpenAI from other factors driving trends in the Plaintiffs' traffic and other business metrics. As detailed below, none of these experts conducted any analysis that isolated OpenAI's alleged use of the Asserted Works as the cause of any change in traffic, subscriptions, revenue, or licensing value for the Asserted Works.

91. Jessie Stricchiola, an expert retained by The Times, DNP, and Ziff Davis to opine on search-engine optimization, compared click-through rate ("CTR"), calculated as clicks divided by impressions, for The Times's, Daily News Plaintiffs', and Ziff Davis's domains in Bing Chat and Bing Search during a limited period. She used the comparison to opine that AI chat interfaces divert referral traffic from publishers. Her analysis measured clicks per impression rather than total traffic or referral volume, did not evaluate changes associated with any metric over time as AI chatbots became available, used Microsoft's Bing products rather than OpenAI's products even though OpenAI data was available, did not account for changes in impression volume, and did not test whether users who otherwise would have clicked in Bing Search instead used ChatGPT. She also extrapolated from limited and inconsistent data without controlling for search engine changes or other confounding factors and ignored OpenAI data that contradicted the conclusions she drew from the Microsoft data. Her analysis showed at most that

two differently designed products (Bing Search and Bing Chat) can generate different CTRs; it did not establish that users substituted ChatGPT for the Asserted Works or that OpenAI caused a decline in traffic, subscriptions, revenue, or other business metrics.[84]

92. Mr. James E. Malackowski relied on Ms. Stricchiola's CTR comparison, industry articles and surveys, ChatGPT usage counts, internal statements, and broad theories about AI substitution and licensing. He compared broad ChatGPT informational-query counts with selected website visits, but he did not conduct an empirical traffic analysis linking particular queries or Asserted Works to displaced visits or use a but-for test. Those materials did not show that users chose OpenAI instead of the Asserted Works or that OpenAI caused any decline in traffic, subscriptions, revenue, or licensing value.

93. Professor Emily Bell relied on industry studies, surveys, and small user studies concerning AI-generated news and Google AI Overviews. She used those materials to opine that generative AI products may substitute for news publishers' content and cause traffic, reputational, and other market harm. But she did not conduct any analyses of The Times's traffic or advertising revenue and did not isolate any effect attributable to OpenAI. The materials she cited did not establish that any user substituted ChatGPT for The Times's works.

94. Professor Taylor Owen relied on counts of AI-generated news websites, survey responses and anecdotes concerning Google AI Overviews, AI search, and ChatGPT features. He

[84] Goldfarb NYT-DNP Rebuttal Report ¶¶ 101-115; Goldfarb ZD Rebuttal Report ¶¶ 80-93; Goldfarb ZD Reply Report ¶¶113-115.

used those materials to opine that generative AI could substitute for original reporting and divert search traffic from Daily News Plaintiffs. But he did not analyze DNP's traffic or revenue data, identify displaced visits, subscriptions, or revenue, or isolate an effect attributable to OpenAI. Therefore, his materials did not establish substitution for the Asserted Works or OpenAI-caused harm.

95. Professors Bell and Owen and Mr. Malackowski also relied on effects attributed to Google AI Overviews and broader changes to traditional search as evidence of harm. Google integrated AI Overviews into its existing search engine, whereas ChatGPT is a separate OpenAI product. Plaintiffs' experts did not identify a mechanism, data, or analysis tying ChatGPT to Google's referral CTR or showing that Google's effects were the same as ChatGPT's. Evidence of a Google product's effect therefore did not establish substitution for the Asserted Works or harm caused by OpenAI.[85]

96. Dr. McDonald's regression analysis addressed The Intercept and CIR, using Mother Jones data for the CIR analysis, and was intended to attribute historical and long-running traffic declines to AI-driven diversion, but it did not reliably isolate any ChatGPT effect. Dr. McDonald treated time as a proxy for structural forces and assumed for purposes of his analysis that AI-driven traffic displacement was the most significant such force, even though the regressions did not include a measure of ChatGPT use and did not account for other factors such

[85] Goldfarb NYT-DNP Reply Report ¶¶ 115-118; Goldfarb ZD Reply Report ¶¶ 98-106; Goldfarb CIR-INT Reply Report ¶¶72-73.

as COVID-19, search and social-platform changes, publication volume, and site migration. His own caveat was that he could not rule out a correlation between AI growth and harm but that his analysis could not establish harm or its magnitude.[86]

97. Dr. Waehrer's correlation analysis does not reliably establish any impact of ChatGPT on Ziff Davis's traffic. It regresses monthly search sessions for two of Ziff Davis's properties on overall ChatGPT sessions but does not isolate the effect of ChatGPT or establish what traffic would have been in the absence of ChatGPT. Accordingly, the analysis measures correlation rather than causation, does not test whether ChatGPT users substituted ChatGPT for Ziff Davis's content, and does not address the confounding factors that Dr. Waehrer himself acknowledges can affect traffic, including seasonality, COVID-era disruptions, site changes, marketing, competitive developments, and search-algorithm updates.[87]

I declare under penalty of perjury under the laws of the United States of America that the foregoing is true and correct, and that I have personal knowledge of these facts and am competent to testify thereto if called upon to do so.

Dated: 09/04/2026
Toronto, Canada

By: _____
Avi Goldfarb

[86] Goldfarb CIR-INT Reply Report ¶¶ 108-132.
[87] Goldfarb ZD Reply Report ¶¶ 155-163.