EXHIBIT 9

(FILED UNDER SEAL)

HC-AEO

UNITED STATES DISTRICT COURT
NORTHERN DISTRICT OF CALIFORNIA

CONCORD MUSIC GROUP, INC., et al.,
Plaintiffs,
v.
ANTHROPIC PBC,
Defendant.

Case No. 5:24-cv-03811-EKL-SVK

EXPERT REPORT OF ABHISHEK NAGARAJ, PH.D.
Signed in San Francisco, California on December 19, 2025



Abhishek Nagaraj, Ph.D.

I. INTRODUCTION

A. Qualifications

1. My name is Abhishek Nagaraj. I am an Associate Professor in the Haas School of Business at the University of California, Berkeley, and a Research Associate at the National Bureau of Economic Research ("NBER"). I am also Director of the Data Innovation and AI Lab, and a Faculty Affiliate of the Berkeley Institute of Data Science.

2. My research is in the field of applied microeconomics, with a particular focus on the economics of innovation, entrepreneurship, and the digital economy. Several of my published articles focus on digital technologies and copyright, including "Digitization and the Demand for Physical Works: Evidence from the Google Books Project" published in the *American Economic Journal: Economic Policy*,[1] "Does Copyright Affect Reuse? Evidence from Google Books and Wikipedia" published in *Management Science*,[2] and "Improving data access democratizes and diversifies science" published in the *Proceedings of the National Academy of Sciences.*[3] I was asked by the United States Copyright Office to serve as a contributing author to its volume *Identifying the Economic Implications of Artificial Intelligence for Copyright Policy: Context and Direction for Economic Research*[4] and was the lead author for the chapter titled "Developers' Access to Training Data." I am also currently working on research on the role of book data in shaping the performance of AI systems.[5]

3. I have received numerous grants and awards for my research in digital technologies, including an MIT Initiative on the Digital Economy Grant, an Institute for Business Innovation

---

[1] *See* Abhishek Nagaraj and Imke Reimers, "Digitization and the Market for Physical Works: Evidence from the Google Books Project," *American Economic Journal: Economic Policy*, Vol. 15, No. 4 (2023), available at https://www.aeaweb.org/articles?id=10.1257/pol.20210702.

[2] *See* Abhishek Nagaraj, "Does Copyright Affect Reuse? Evidence from Google Books and Wikipedia," *Management Science,* Vol. 64, No. 7 (Jul 2018), available at https://www.jstor.org/stable/48748181.

[3] *See* Abhishek Nagaraj, Esther Shears, and Mathijs de Vaan, "Improving data access democratizes and diversifies science," *Proceedings of the National Academy of Sciences,* Vol. 117, No. 38 (22 Sep 2020), available at https://www.pnas.org/doi/10.1073/pnas.2001682117.

[4] *See* Brent Lutes, Joshua Gans, et al., "Identifying the Economic Implications of Artificial Intelligence for Copyright Policy," *US Copyright Office* (2025), available at https://www.copyright.gov/economic-research/economic-implications-of-ai/Identifying-the-Economic-Implications-of-Artificial-Intelligence-for-Copyright-Policy-FINAL.pdf.

[5] *See* Abhishek Nagaraj and Stella Jia, "Cloze Encounters: The Impact of Pirated Data Access on LLM Performance," *National Bureau of Economic Research* (Mar 2025), available at https://www.nber.org/papers/w33598.

Research Grant, a Facebook Measuring Economic Impact in the Digital Economy grant, a UPenn Center for Technology, Innovation & Competition Research Award, and the Technology and Innovation Management Outstanding Dissertation Award. I have taught graduate-level classes on Entrepreneurial Strategy, the Economics of Innovation, and AI for Business Leaders. I was the co-organizer of the 2024 NBER Digital Economics and AI winter meetings, and a Panelist in the 2024 AI & Copyright Policy Roundtable organized by the United States Copyright Office. I am a contributing author on the forthcoming NBER volume on the Economics of Transformative AI.[6]

4. I hold a Bachelor of Technology in Computer Science from the College of Engineering in Pune, an MBA from the Indian Institute of Management Calcutta, and a Ph.D. in Management from the Sloan School of Management at the Massachusetts Institute of Technology. I have enclosed my CV as Appendix A to this report.

5. I am being compensated at my hourly rate of $1,200. I have not testified in deposition or trial in the last four years. I am being supported in my research for this report by staff at Keystone, an economics consulting firm. Neither my compensation nor Keystone's is affected by the outcome of this matter. My research in this area is ongoing. I reserve the right to modify or change my opinion if additional facts become available to me.

B. Assignment

6. I have been asked by Counsel for Anthropic PBC ("Anthropic") to analyze economic issues related to fair use factor 4, which is "the effect of the use upon the potential market for or value of the copyrighted work."[7] I have been asked to assess this factor in relation to Plaintiffs' claims regarding Anthropic's alleged use of lyrics from songs to which Plaintiffs hold a copyright. Specifically, Plaintiffs allege that the text from the lyrics to those Works-In-Suit were used to train a series of Anthropic's general-purpose large language models ("LLMs") known as

[6] *See generally* "Economics of Transformative AI Workshop, Fall 2025," NBER, available at https://www.nber.org/conferences/economics-transformative-ai-workshop-fall-2025; *see also* "NBER Digital Economics and AI Meeting, Spring 2024," NBER, available at https://www.nber.org/conferences/nber-digital-economics-and-ai-meeting-spring-2024.

[7] 17 U.S. Code § 107, "Limitations on exclusive rights: Fair use," ("In determining whether the use made of a work in any particular case is a fair use the factors to be considered shall include…the effect of the use upon the potential market for or value of the copyrighted work.").

lyrics displays and will often synchronize the music with the lyrics. By contrast, as a text-based general-purpose LLM, Claude does not play sound recordings of musical compositions.[135] Similarly, Genius.com provides both the lyrics and user-generated annotations to provide deeper insight into the origins of various lyrical elements in songs. For these reasons, I expect the degree of substitution between Claude.com and online lyrics aggregators to be low. However, economic theory does not rule out any negative effect; it only predicts that it is unlikely.

90. My empirical analysis therefore focuses on whether Claude's display of lyrics from songs impacts the market for online display of lyrics.

2. Empirical Analyses of Plaintiffs' Financials and Website Traffic Data Do Not Show Any Negative Impact on the Market for Online Display of Lyrics Due to Anthropic's Alleged Outputs Containing Lyrics

91. To assess whether there was harm to the market of online display of lyrics, I first conduct an analysis of Plaintiffs' revenue statements. Because of the sparse nature of the data, and because Plaintiffs only provided revenue statements that go past [redacted] three days before this report was due to be served, I have been unable to conduct a formal econometric analysis. I understand that another expert is reviewing this data from a non-econometric framework.

92. I analyze Plaintiffs' gross revenue and lyric revenue, which are available only on a quarterly or semi-annually (rather than monthly) basis. What I have seen in both sets of revenue does not conclusively show that Plaintiffs have been harmed. I reserve the right to supplement my report with further analyses in light of this new data.

93. My primary investigation looks at whether the introduction of Claude negatively impacted web traffic to other platforms and websites that display lyrics in the existing market for online display of lyrics. If Plaintiffs' theory of harm were correct, as traffic to Claude increased, I would expect to see a corresponding decrease in existing ways for users to access lyrics on the internet. To empirically test this theory, I use granular data on the usage of Claude and on two of the primary means for users to access lyrics: lyrics websites and Google search. Based on these empirical analyses, I find no evidence that either traffic to lyrics websites or lyrics-focused Google

[135] Claude has some limited multimodal uses. For example, *see* "Claude: Using images, audio, and video in practical workflows," Data Studios, 21 Aug 2025, available at https://www.datastudios.org/post/claude-using-images-audio-and-video-in-practical-workflows, accessed on 14 Dec 2025.

searches decreased as usage of Claude increased. I thus conclude that Plaintiffs' theory of harm is not empirically supported. This analysis is discussed below in further detail.

a) Plaintiffs' Revenue Data Does Not Show Any Impairment as a Result of Anthropic's Alleged Use

94. As described above, I first analyze Plaintiffs' revenue statements associated with the 499 Works-In-Suit asserted in Plaintiffs' Complaint. As I explain in further detail below, I find that the revenues reported in those financial documents show [redacted] which is inconsistent with the Plaintiffs' theory of harm.[136] However, due to the limitations in the data—including that Plaintiffs only provided revenue statements that extend beyond [redacted] three days before this report was due—I was unable to draw any meaningful conclusions about the impact of Claude on Plaintiffs' revenues.

95. For UMPG's financial information, I rely on two main datasets provided by UMPG: (1) the quarterly [redacted] document,[137] and (2) the [redacted] document,[138] which both include revenue data through [redacted] I understand that three days prior to this report being due, UMPG provided additional revenue data through [redacted].[139]

96. For Concord's financials, I rely on two datasets: (1) the [redacted] document,[140] and (2) the [redacted] document,[141] which both include data through [redacted]. I understand that three days prior to this report being due, Concord provided additional revenue data through [redacted][142]

---

[136] [redacted]. Therefore, in my analysis, I do not account for the revenue generated by [redacted].

[137] *See* UMPG 0000000616, "[redacted] The song "Baby Baby" was initially listed as owned by UMPG, but CCMG was listed as the owner when the First Amended Complaint was filed. Order to Revise Exhibit A to Plaintiffs' First Amended Complaint, 25 Apr 2025.

[138] *See* UMPG 0000000638, "[redacted]."

[139] *See* UMPG 0000025145.

[140] *See* CONCORD 0000000811, "[redacted]."

[141] *See* CONCORD 0000000815, "[redacted]."

[142] *See* CONCORD 0000016756.

97. For ABKCO's financial data, I use a lyric-specific revenue dataset, reported at a [redacted].[143] I understand that three days prior to this report being due, ABKCO provided updated revenue data through [redacted]

98. Finally, for CCMG, I use (1) [redacted] data and (2) [redacted] data that includes information through [redacted].[145] I understand that three days prior to this report being due, CCMG provided updated revenue data through [redacted].[146]

99. Using the revenue statements provided that go through [redacted] I plot the revenues for each Plaintiff by totaling the revenues reported across all songs and categories within each year. Figure 1 shows the total gross revenue over all categories, and Figure 2 shows total lyric revenue. As I interpret these plots, they do not suggest that Plaintiffs have been harmed. I reserve the right to supplement my report with further analyses in light of this data.

---

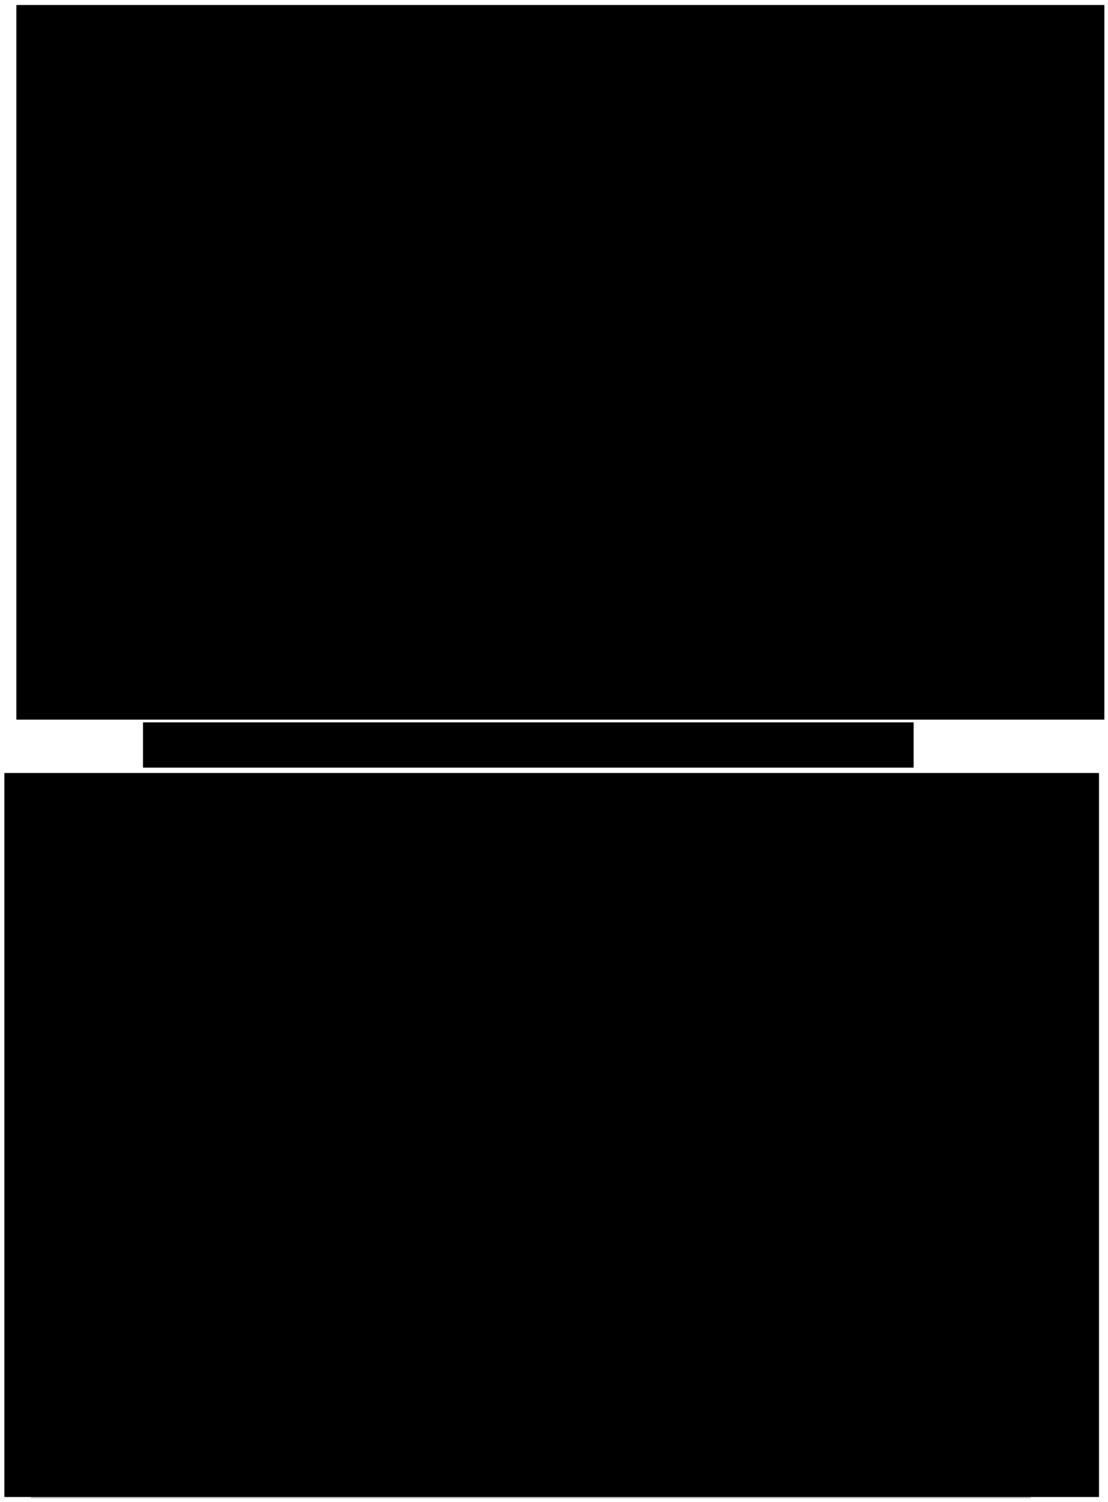
[143] *See* ABKCO 0000000611, "[redacted]."

[144] *See* ABKCO 0000001512.

[145] *See* UMPG 0000000728–793. [redacted] For example, UMPG 0000000728 is [redacted]

[146] *See* CCMG 0000001277.

[redacted]



100. To summarize, from a macro-level review of the revenue data, I observe that revenue for all four Plaintiffs [redacted] My limited conclusion is

based on the data provided through [redacted], as I have not had sufficient time to meaningfully review the updated data. Accordingly, I reserve the right to supplement my report in light of this new data.

b) My Empirical Analysis of Semrush Data on Traffic to Lyrics Websites Is Inconsistent with Plaintiffs' Theories of Harm

101. According to Plaintiffs, allowing users to obtain song lyrics directly through queries to Claude caused users to replace their existing sources of lyrics with Claude.[147] In this section, I analyze one of the primary traditional means for users to access lyrics: lyrics websites such as AZlyrics.com and Genius.com. If Plaintiffs' theory of harm is correct, then increased traffic to Claude.ai after the site's commercial launch to consumers in July 2023[148] would be associated with decreased traffic to major lyrics websites in the United States.

102. For my empirical research analyzing this hypothesis, I use daily website traffic data to the top lyrics websites in the United States that measures the number of users who visit a particular website on a given day. By using the econometric approach detailed below, I can examine the relationship between daily website traffic records for relevant lyrics websites and the traffic to Claude.ai. I include controls to account for factors such as website characteristics and the impact of exogenous events like new releases from popular artists like Beyoncé or Taylor Swift, which might be the primary driver for an increase in traffic to lyrics websites.[149] I also conduct several robustness checks for my baseline specification, and the results remain qualitatively similar. Based on my empirical study, I find no statistical evidence that increased use of Claude is

[147] Amended Complaint, ¶156 ("Anthropic's unauthorized use also undercuts the existing and potential markets for licensing song lyrics and undermines those legitimate lyrics aggregators, websites, and digital services that properly license lyrics, given that users of Anthropic's AI models do not visit the legitimate sites that compensate Publishers for the right to use their lyrics.").

[148] Claude was first launched in March 2023, but only to API users. On July 11, 2023, Anthropic released its consumer-facing product by launching Claude.ai, a public-facing website that allows consumers to directly access the Claude chatbot. *See* "Claude 2," Anthropic, 11 Jul 2023, available at https://www.anthropic.com/news/claude-2, accessed 15 Dec 2025.

[149] Specifically, to control for time-variant factors besides Claude's introduction that can impact outcome of interest, I include time-fixed effects. Similarly, to control unobserved time-invariant factors such as website traits, I also include website-fixed effects. *See* Joshua D. Angrist and Jörn-Steffan Pischke, Mostly Harmless Econometrics: An Empiricist's Companion, 2009, at Chapter 5.1.

associated with decreased use of lyrics websites. The results of my findings are therefore inconsistent with Plaintiffs' theory of harm.

(1) Identification of Lyrics Websites

103. To perform my analysis, I use a list of lyrics websites where users can access or view song lyrics. To compile that list, I consult three sources: (1) the First Amended Complaint,[150] (2) a list of the top lyrics websites from Semrush, which provides digital marketing analytics, including aggregated data on search queries, keywords, and web traffic,[151] and (3) Plaintiffs' response to Interrogatory No. 25.[152] The First Amended Complaint refers to several "popular lyrics websites" which Publishers allege Anthropic scraped from to get lyrics data. Those websites include MusixMatch.com, LyricFind.com, Genius.com, Lyrics.com, and AZLyrics.com. [153] Because these lyrics websites are also included in the data obtained from Semrush, I therefore focus only on websites from Semrush and Plaintiffs' Response to Interrogatory No. 25 that primarily provide lyrics.[154]

104. Using the information from Semrush and Plaintiffs' Response to Interrogatory No. 25, I ultimately compile a list of 26 websites that I use for my empirical study. A list of the websites is available in Table C.4 in the Appendix.

105. I narrow the final list down to 26 lyrics websites using several steps. First, using the Semrush data, I pull the top 100 websites by web traffic that appeared in Google searches mentioning "lyrics" in October 2025 for the United States. This list is provided in Table C.5. I use October 2025 data as the end date for my analysis because Semrush provides a complete dataset for a given month by the 10th of the following month, and thus this was the most recent dataset I

---

[150] *See* Amended Complaint, at ¶66.

[151] *See* "About Us," Semrush, available at https://www.semrush.com/company/, accessed on 12 Dec 2025.

[152] *See* Plaintiffs' Supplemental Objections and Responses to Defendant's Interrogatory No. 25, 18 Sep 2025.

[153] Amended Complaint, at ¶66.

[154] Note that I do not include websites that primarily contained sheet music, as the market I am evaluating is the market for online display of lyrics. I also note that several of these lyrics websites do not appear to have licenses with music publishers for the display of lyrics. *See* Alex Pham, "NMPA Targets Unlicensed Lyric Sites, Rap Genius Among 50 Sent Take-Down Notices," Billboard, 11 Nov2013, available at https://www.billboard.com/music/music-news/nmpa-targets-unlicensed-lyric-sites-rap-genius-among-5785701/, accessed on 19 Dec 2025.

could assess.[155] Semrush provides two lists—one for desktop searches[156] and one for mobile device searches.[157] In the list of top 100 lyrics websites from Semrush on desktop searches, I notice a website may appear multiple times on the list if two or more pages on the website received significant traffic. For example, both "https://www.azlyrics.com/" and "https://www.azlyrics.com/f/free.html" appeared in the top 100 list, but since they are effectively the same website, I do not find it meaningful to include both for the purposes of my empirical study. Taking into account those duplicate websites, I cut down the list of 100 websites from Semrush to 65 unique websites as in Table C.6 . I then cut this list down from 65 to 24 websites by browsing each site myself to determine if the website fits the paradigm of a lyrics website. I classify a website as a lyrics website if it is a U.S. website with a primary purpose of displaying lyrics. In Table C.7 in the Appendix, I list the 41 websites I exclude from the original list of 100 sites from Semrush, and I include my reasoning of why I remove each of the 41 excluded sites. Some of the reasons for exclusion include websites that are not from the United States, websites that are primarily for music streaming and not solely for lyric display, and other websites that are simply not lyrics websites, like Berklee.edu and Wikipedia.org. I also look at the Semrush list for mobile searches. This contains only 20 entries, with youtube.com appearing three times. All websites from Semrush for mobile searches match with websites from the desktop list, with the exception of one site which is not a lyrics website. This list can be found in Table C.8.

106. Second, I also refer to Plaintiffs' list of "URL domains and software applications" that "have executed licenses with one or more Publishers to use or display the lyrics to one or more of the Asserted Works" identified in Plaintiffs' response to Interrogatory No. 25.[158] For each of the 38 entries in Plaintiffs' list, I check whether the website is a lyrics website (i.e., a U.S. site with the primary purpose of displaying lyrics) and, if so, I add it to the list of relevant sites. Based on

---

[155] "Welcome to Trends API," Semrush Developer API, 2 Jul 2025, available at https://developer.semrush.com/api/v3/trends/welcome-to-trends-api/, accessed on 19 Dec 2025 ("The complete dataset for the prior month becomes available by the 10th day of the current month.").

[156] "Keyword Overview: lyrics (Desktop)," Semrush, available at https://www.semrush.com/analytics/keywordoverview/?q=lyrics&db=us&date=202510, accessed on 12 Dec 2025.

[157] "Keyword Overview: lyrics (Mobile)," Semrush, available at https://www.semrush.com/analytics/keywordoverview/?q=lyrics&db=us&date=202510&device=1, accessed on 12 Dec 2025.

[158] Plaintiffs' Supplemental Objections and Responses to Defendant's Interrogatory No. 25, 18 Sep 2025, 8:3-5.

this criteria, two unique websites are added for my review.[159] In Table C.9 in the Appendix, I list the 32 excluded sites, including the reason why each was removed for the purpose of my analysis. Some of the reasons for exclusion include websites that I cannot find, websites that are primarily for music streaming and not solely lyric display, and social media platforms like Instagram and Meta.

107. All combined, I identify a total of 26 lyrics websites that satisfy the criteria of being a U.S. website with a primary purpose of displaying lyrics: 24 come from the Semrush list, five of them are also mentioned in the Complaint[160], and 2 unique sites are added from my review of Plaintiffs' response to Interrogatory No. 25. A list of the 26 websites I ultimately rely upon for this study (hereinafter, "Lyrics Websites") are included in Appendix Table C.4.

(2) Traffic Data

108. In this analysis, I am measuring the relationship between the use of Claude.ai and the use of the 26 Lyrics Websites I identify through the above process. The main variable of interest in my analysis is web traffic to Claude.ai, while the explained variable is web traffic to the Lyrics Websites.

109. To collect web traffic, I use Semrush's API to collect historical daily total traffic to the Lyrics Websites and Claude.ai.[161] For the identified Lyrics Websites, I consider their historical daily total traffic from January 2022 to October 2025. I select this time period for my main analysis, as it includes a significant amount of time both before and after the introduction of Claude.ai in July 2023. I choose 2022 as the start of my sample to avoid any unintended effects on the data from the COVID-19 pandemic.[162]

---

[159] There are six websites listed in Interrogatory No. 25 that fit the criteria of being a website that displayed lyrics online. Four of these were already accounted for in the Semrush list, so I do not include them again in my ultimate list.

[160] Amended Complaint, at ¶66.

[161] *See* "Trends API," Semrush, available at https://www.semrush.com/kb/930-trends-api, accessed on 12 Dec 2025. Semrush partners with clickstream data providers to build a panel of over 200 million real, anonymized internet users around the world. They then process this data through their Neural Network algorithm to produce accurate website traffic estimates. *See* "Data and Metrics," Semrush, available at https://www.semrush.com/kb/997-semrush-data, accessed on 18 Dec 2025.

[162] Raj Chetty et al., "The Economic Impacts of COVID-19: Evidence from a New Public Database Built Using Private Sector Data," *The Quarterly Journal of Economics*, Vol. 139, No. 2 (May 2024), at pp. 829–889, available at https://doi.org/10.1093/qje/qjad048. ("We evaluate the empirical value of the new data by using them to analyze the economic effects of COVID-19, focusing on the period from March 2020 to December 2021—covering both the decline in economic activity and the recovery to baseline spending levels.").

110. For Claude.ai data, I use Semrush to obtain daily web traffic to Claude.ai, which is the main website offering Claude models to the public.[163] To the best of my knowledge, before Claude.ai's launch on July 11, 2023, Claude was accessible only through (i) partnerships with Slack, Quora (Poe), Juni Learning, DuckDuckGo, Robin AI, Notion, Zoom, and AssemblyAI,[164] and (ii) through an API for developers.[165] For my empirical study, I do not evaluate these modalities for several reasons. First, developer API tools and some of the platforms with Claude integration (like Juni Learning, an online education platform) are not relevant to my analysis because their predominant use case is related to coding or other business needs, and as I stated before, it is unlikely that enterprise users are using Claude to search for lyrics. Second, to the extent there are consumer-facing API integrations, like Quora (Poe) and DuckDuckGo, where a user could conceivably request lyrics, I expect that use case to be relatively small compared to use through Claude.ai. And in any case, if users on Poe were using Claude to request lyrics, I would also expect them to use Claude.ai to request lyrics, which my study accounts for.

111. In my main specification, I consider data from the July 11, 2023 launch onward, to account for traffic to Claude.ai. As one would expect, the traffic to Claude.ai is zero before July 11, 2023.[166] This can be also seen in Figure C.5 in the Appendix. Furthermore, I set traffic to Claude.ai at zero after [redacted]. This is because Anthropic bolstered their guardrails as it relates to a broad set of music lyrics on this date which, as I will explain, dramatically reduced the possibility that Claude.ai could substitute for song lyrics.

---

[163] With the second version of Claude, the product Claude.ai (the public-facing version of Claude) was also launched on July 11, 2023. "Claude 2," Anthropic, 11 Jul 2023, available at https://www.anthropic.com/news/claude-2, accessed on 12 Dec 2025.

[164] *See* "Introducing Claude," Anthropic, 14 Mar 2023, available at https://www.anthropic.com/news/introducing-claude, accessed on 8 Dec 2025; *see also* "Meet Claude," Anthropic, 12 Jun 2023, available at https://web.archive.org/web/20230612193739/https://www.anthropic.com/product, accessed on 13 Dec 2025.

[165] *See* "Introducing Claude," Anthropic, 14 Mar 2023, available at https://www.anthropic.com/news/introducing-claude, accessed on 8 Dec 2025; *see also* "Meet Claude," Anthropic, 12 Jun 2023, available at https://web.archive.org/web/20230612193739/https://www.anthropic.com/product, accessed on 13 Dec 2025. For my empirical study, I do not evaluate these modalities for several reasons. First, developer API tools and some of the platforms with Claude integration (like Juni Learning, an online education platform) are not relevant to my analysis because their predominant use case is related to coding or other business needs, and as I stated before, it is unlikely that enterprise users are using Claude to search for lyrics. Second, to the extent there are consumer-facing API integrations, like Quora (Poe) and DuckDuckGo, where a user could conceivably request lyrics, I expect that use case to be relatively small compared to use through Claude.ai. And in any case, if users on Poe were using Claude to request lyrics, I would also expect them to use Claude.ai to request lyrics, which my study accounts for.

[166] There is a small amount of traffic to Claude.ai on July, 10 2023 which could be due to internal testing, which I ignore since Claude.ai was officialy launched on July 11, 2023.

112. As I understand it, Anthropic developed guardrails to prevent the generation of potentially infringing outputs prior to the first release of Claude in March 2023, and Anthropic has continued to bolster and improve those guardrails over time to evolve to respond to new and varied threats.[167] Based on the information I have reviewed, Anthropic's guardrails currently consist of: (1) Prompt Shield, which modifies a user's prompts when the system detects a user's "attempts to generate an output that would violate Anthropic's Terms of Use or Acceptable Use Policy,"[168] and (2) Content Compare System ("CCS"), a tool that blocks certain outputs from displaying that output if it detects "that the output would reproduce an amount of a work in the hash corpus above a certain threshold."[169] I understand that Prompt Shield was first implemented in its first-party API product to prevent lyric display for a large number of songs on or around [redacted] This is in addition to the CCS guardrails that were bolstered on [redacted], when Anthropic [redacted].[171]

113. Anthropic's guardrails are highly effective. Based on the materials I reviewed, Anthropic had (and continues to have) multiple layers of defenses against reproducing copyrighted material.[172] In response to prompts from users, Claude returned error messages saying that it could not reproduce lyrics because of copyright concerns.[173] My review of the record is consistent with what I understand from Dr. Bilmes. Averaged across the available models at issue in this case and

---

[167] Defendant Anthropic's PBC's Amended and Supplemental Responses and Objections to Plaintiffs' First Set of Interrogatories (Nos. 9, 11, 12), 12 Nov 2025, 12:1-2.

[168] Defendant Anthropic's PBC's Amended and Supplemental Responses and Objections to Plaintiffs' First Set of Interrogatories (Nos. 9, 11, 12), 12 Nov 2025, 12:24-13:1.

[169] Defendant Anthropic's PBC's Amended and Supplemental Responses and Objections to Plaintiffs' First Set of Interrogatories (Nos. 9, 11, 12), 12 Nov 2025, 13:21. I further understand that Anthropic employs many other tools in addition to its guardrails to protect against copyright infringement, including the Terms of Service and Usage Policy. I focus on the guardrails I have identified for the purposes of my empirical studies because they are directly intended to prevent lyric regurgitation in Claude's outputs.

[170] Defendant Anthropic's PBC's Amended and Supplemental Responses and Objections to Plaintiffs' First Set of Interrogatories (Nos. 9, 11, 12), 12 Nov 2025, 13:6-7.

[171] Defendant Anthropic's PBC's Amended and Supplemental Responses and Objections to Plaintiffs' First Set of Interrogatories (Nos. 9, 11, 12), 12 Nov 2025, 14:3-4.

[172] I understand Dr. Bilmes will be discussing guardrails implemented in Claude further in his report.

[173] *See* Second Kaplan Declaration, at ¶¶ 52-54. I understand from counsel that for one of the 499 Works-In-Suit, the incorrect song was added to the [redacted] Plaintiffs asserted infringement of "Joy to the World" by Three Dog Night, but the work that was added to the [redacted] in November 2023 was the holiday song, "Joy to the World." I further understand from counsel that the correct song was added to the [redacted] December 2025. This does not meaningfully impact my analysis because the traffic analysis considers [redacted] as the event date.

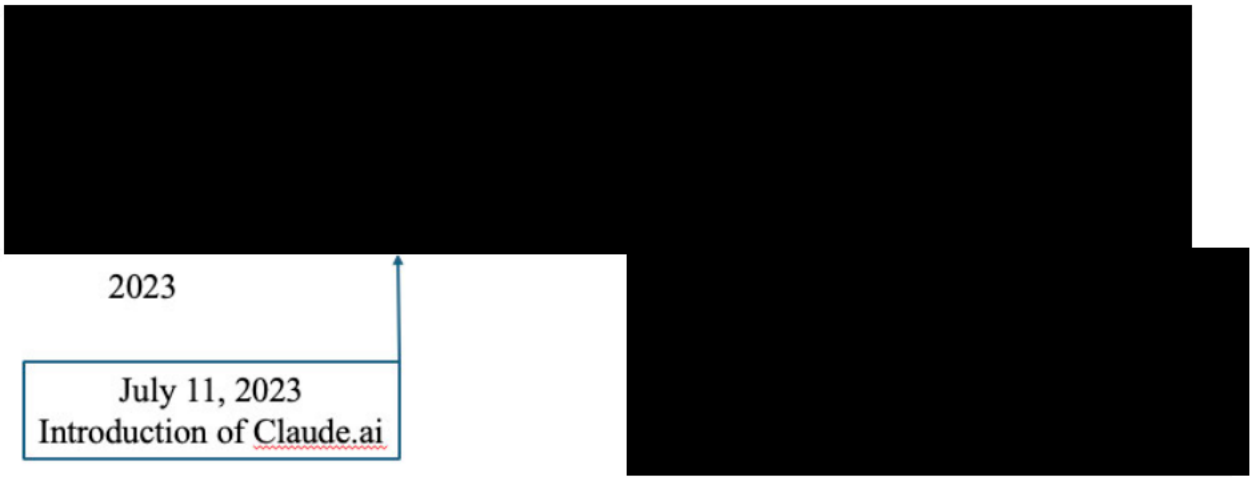
across the three prompt variations, I understand that in his study, Dr. Bilmes finds that Claude did not respond to prompts requesting lyrics with a response containing more than five consecutive words of a song in more than 94.4% of the responses. Moreover, when prompted for song lyrics, all Claude models respond less than 1% of the time with an output that matches more than 10% of the lyrics.[174] In Dr. Bilmes' study, none of the attempts showed Claude including more than 25% of a song's lyrics as a single consecutive match in response to a prompt.

114. Given the guardrails' effectiveness against reproducing full lyrics in Claude's outputs, I would expect there to be no, or very few, instances where Claude has produced partial or full sets of lyrics in outputs after the guardrails were bolstered. Dr. Bilmes has confirmed this assumption. As a result, once the guardrails were implemented and bolstered, it would have been unlikely that users could use Claude to receive the entire or even partial reproduction of lyrics in outputs. Moreover, I understand from Dr. Bray that even *prior* to the imposition of guardrails, user attempts to find lyrics were very rare. Accordingly, I assign a value of zero to traffic to Claude.ai in the post-guardrail period.

115. For purposes of this lyrics websites analysis, I use [redacted], the date Prompt Shield was first implemented, as the effective guardrail activation date. I select the date Prompt Shield was implemented, as opposed to the date when CCS was bolstered, because Prompt Shield affected all users' lyrics requests, whereas CCS [redacted] And because I do not have song-by-song traffic data for lyrics websites, it makes more sense to apply the guardrail date that would affect a broader set of traffic data. Therefore, I consider that after Prompt Shield was implemented in December 2023, the Claude.ai traffic was zero. Figure 3 below illustrates the timeline and analysis period.

[174] The same number for Claude Haiku 3.5 is ~1.2%.

Figure 3: Illustration of Traffic Information Considered in the Analysis



(3) Methodology

116. The goal of my analysis is to assess the relationship between lyrics website traffic (on the one hand) and traffic to Claude.ai (on the other). Economists can use a statistical technique called a regression analysis to estimate how much impact a change in one variable (in this case, visits to Claude.ai) has on another variable (in this case, visits to lyrics websites). The simplest form of regression—and the one I use here—is a linear regression, which estimates a coefficient that tells me how much changes in the independent variable (i.e., Claude.ai traffic) alter the dependent variable (i.e., website traffic).[175] This allows me to see how traffic to Claude.ai is related to website traffic. If Plaintiffs' theory of harm is true, I should see a significant—in econometrics terms, "a statistically significant"—drop in traffic to lyrics after the introduction of Claude.ai on July 11, 2023.

117. However, in this case, it is not enough to draw an imaginary line that implies that as Claude.ai traffic goes up, traffic to lyrics websites goes down. There are factors that impact both

[175] The goal of a regression model is to explain an outcome variable, y, in terms of an explanatory variable, x. For example, one may want to explain how much of the difference in wages can be explained by education. To answer this question, economists can assume that there is a linear relationship between education (x) and wages (y), so that as education goes up, wages also increase. Consider a simple equation of $Y = b_0 + b_1X + U$, where U represents all other factors that I cannot observe. The coefficient b_1 corresponds to the slope of this equation, telling how much 1 unit of increase in education correspond to changes in wages. *See* Jeffrey M. Wooldridge, Introductory Econometrics A Modern Approach, at Chapter 2 and Chapter 3.

Claude.ai traffic and lyrics websites traffic at the same time. For example, during the winter, when people are more likely to be at home to avoid the cold, more people may surf the internet, which may result in more traffic to both Claude.ai and lyrics websites. Therefore, to net out the impact of these factors that I cannot observe in my data, but which nevertheless impact the website traffic, I put in place a number of controls called "fixed effects" that are well known in economic literature.[176]

118. One control I apply is intended to account for factors that remain constant over time, such as website popularity or other website traits. For example, trends in traffic for more popular websites may be similar to the growing popularity of AI tools for unrelated reasons. To account for these factors, I use website-fixed effects in my analysis. For example, one website-fixed effect would control the role that a website's popularity has on the website's traffic. In this study, I apply a website-fixed effect as a variable representing, for example, genius.com—observations that belong to genius.com are assigned a value of 1, and all other observations are assigned a value of 0. This control accounts for the vast variation in popularity across different websites and essentially asks whether websites experienced any drop in traffic accounting for the average level of traffic they receive on any given day.

119. Another control I use is intended to account for time-dependent factors, meaning factors that are not constant over time, such as the release of a new song from a popular artist, or the season, all of which may impact web traffic. All websites, including Claude.ai and lyrics websites, are affected by such time-dependent factors—for example, during the winter, people may spend more time online than during the summer. To control for these time-dependent factors, I use a variable I call a "month-year fixed effect" for each month-year pair available in the data. For example, one of the fixed effect variables in my study is for May 2023, meaning the month-year fixed effect takes the value of 1 for observations belonging to May 2023, and 0 for all other month-year combinations. Altogether, these time-independent and time-dependent variables

[176] *See* Joshua D. Angrist and Jörn-Steffan Pischke, Mostly Harmless Econometrics: An Empiricist's Companion, 2008, at Chapter 5.1.

account for unobserved variations across units and over time.[177] The details of my empirical approach are provided in Appendix C.

(4) Results

120. Table 1 presents the results of my regression model, which looks at the relationship between weekly traffic to lyrics websites and weekly traffic to Claude.ai and provides a coefficient for the relationship between traffic to lyrics websites and that to Claude.ai. A coefficient is a number estimated by the regression model that measures the relationship between an independent variable and the dependent variable, holding all other variables constant. In this case, the coefficient measures how Claude affected traffic to lyrics websites, all other things being equal.[178]

121. As shown in the chart below, this number for the regression I conduct (0.005) is statistically insignificant, and hence one cannot conclude that Claude.ai has an effect on traffic for lyrics websites. This means that the very small positive number I see is not statistically different from zero at any conventional levels of significance (either 5% or 1%).[179] In other words, the coefficient is not meaningfully different from zero, and I cannot statistically differentiate the effect from zero.[180] Based on these results, I do not find any economically or statistically significant effects from Claude.ai's launch in July 2023 on online lyrics consumption, as measured by traffic

[177] Month–year fixed effects are widely used in the economics literature to control for common changes over time that affect all units at once, such as seasonal patterns or broad shifts in overall activity. For example, *see* Erik Brynjolfsson, Danielle Li, Lindsey Raymond, Generative AI at Work, *The Quarterly Journal of Economics*, Vol. 140, No. 2 (May 2025), at pp. 889–942, https://doi.org/10.1093/qje/qjae044.

[178] All traffic variables are expressed in logarithms to account for the large scale and variation in volumes across websites and over time. Following standard econometric practice, traffic variables are expressed in logarithms because they are strictly positive and vary widely in scale across observations. Log transformations can narrow the range of variables with large magnitudes and reduce sensitivity to extreme values. Jeffrey M. Wooldridge, Introductory Econometrics: A Modern Approach, at Chapter 6 ("Another potential benefit of using logs is that taking the log of a variable often narrows its range. This is particularly true of variables that can be large monetary values… Narrowing the range of the dependent and independent variables can make OLS estimates less sensitive to outlying (or extreme values).").

[179] *See* Jeffrey M. Wooldridge, Introductory Econometrics: A Modern Approach, at p. 110. ("We must first decide on a **significance level** ("level" for short) or the probability of rejecting H_0 when it is in fact true. For concreteness, suppose we have decided on a 5% significance level, as this is the most popular choice. Thus, we are willing to mistakenly reject H_0 when it is true 5% of the time.").

[180] In more technical terms, this means that I cannot reject the null hypothesis that Claude has no impact on lyrics websites. *See* Jeffrey M. Wooldridge, Introductory Econometrics: A Modern Approach, at p. 115.

to popular lyrics websites.[181] More details about my methodology and additional robustness test can be found in Appendix C

Table 1: Regression Result for U.S. Lyrics Traffic on Claude Traffic

		Dependent Variable: Log Lyrics Traffic
	[A]	[B]
[1]	Log Claude Traffic	0.005
[2]	(p-value)	0.584
[3]	Observations	5,174
[4]	R Squared	0.844

Sources: Semrush Traffic Data
Data covers weekly traffic to 26 lyrics websites and Claude.ai for full weeks contained in January 2022 to October 2025. Claude.ai traffic is set to zero after Prompt Shield ([redacted]). P-value is calculated using the wild cluster bootstrap method.
[B][1] Estimate of regressing log lyric traffic on log Claude.ai traffic.
[B][2] p-value calculated using wild cluster bootstrap method shows that the coefficient at [B][1] is statistically insignificant.

c) My Empirical Analysis of Google Data on Searches for Lyrics Is Inconsistent with Plaintiffs' Theories of Harm

122. Similar to Section III.B.2.b) above, in this section I empirically test the effect of Claude usage on a second traditional means for users to view lyrics: Google searches. An advantage of this analysis is that I can access song-by-song data, thereby allowing me to focus specifically on the 499 songs at issue in this case.

123. Similar to Section III.B.2.b) above, I focus on Claude.ai and its release in July 2023 to test this claim. I also treat the effective usage of Claude for lyrics text as zero after Anthropic bolstered its guardrails. However, unlike with the empirical study involving lyrics websites, my Google search analysis is able to consider specifically the songs at issue in this case. Because the Google search empirical study is a song-by-song analysis, and CCS is a guardrail that blocks certain outputs [redacted]

[181] Results in Appendix C for robustness checks are also consistent with this conclusion. In any alternative specification, I do not find any significant impact.

[redacted], the relevant date I consider for the bolstering of guardrails is therefore [redacted].[182]

124. For this study, I obtain weekly data from Google Trends indicating search interest for each song in the WIS chart compiled by counsel for purposes of conducting my analysis as presented in Table C.1. I describe this process in further detail below.

125. In order to empirically test Plaintiffs' theory of harm, I need to account for two critical features of Google searches for songs:

a. First, interest in an individual song is likely affected by the overall popularity of that song, and events such as a new release by the same artist, news coverage about the artist, use of the song in a popular TV series or movie, the time since the song's release or other factors unrelated to Claude.

b. Second, not all Google searches for a song are searches for lyrics; users may also be trying to (for example) find a music video.

126. To account for these factors, I obtain Google Trends index data to measure interest for two types of Google searches. First, the "song index" consists of searches likely related to the audio of a song, which cannot be reproduced by Claude because Claude does not output audio. The second is the "lyrics index," which consists of searches likely related to the lyrics text of a song (text that Plaintiffs allege could be reproduced by Claude).[183] The indices, ranging in value from 0 to 100, are relative to each other and over time, so I can observe the evolution of search interest for "lyrics" queries relative to "song" queries. I also use the data showing weekly user visits to Claude.ai as a measure of the potential for Claude users to have substituted search on Google with use of Claude.

127. If Plaintiffs' theory of harm is correct, I expect to find that as Claude usage increases, the lyrics index decreases relative to the song index, because more people would be turning to Claude to find lyrics. Therefore, I estimate an econometric model that examines whether

[182] Defendant Anthropic PBC's Amended and Supplemental Responses and Objections to Plaintiffs' First Set of Interrogatories (Nos. 9, 11, 12), 12 Nov 2025, 14:3-4 ("Separately, on or around [redacted], Anthropic added the lyrics to the then-500 Works-In-Suit to [redacted]

[183] All Google Trends data are obtained from the subscription-based tool SearchAPI. *See* "Google Search API," SearchAPI, available at https://www.searchapi.io/, accessed on 12 Dec 2025. I discuss the specific terms I use to acquire the song index and lyrics index for each song below.

weeks in which more users visited Claude.ai correspond to weeks with diminished search interest for lyrics-focused searches relative to song-focused searches. Again, my goal with this analysis is to compare how traffic to Claude.ai affects the search interest for "lyrics" as compared to searches for "songs." Doing so allows me to assess whether Plaintiffs' theory of harm is correct (i.e., that internet users were turning to Claude.ai to find lyrics rather than using Google Search to do so).

128. In my main specification, I consider the timeline of January 2022 through October 2025.[184] I choose to start from 2022 to avoid any unintended effects from the COVID-19 pandemic, as I did for my previous analysis.

129. In the remainder of this section, I detail the specific search terms I use to create the lyrics index and song index for each of the 499 songs in the WIS Chart in Table C.1

(1) Construction of Lyrics Index and Song Index Using Google Trends

130. Google Trends is a product developed by Google which provides information on actual search queries.[185] For each term queried, Google Trends returns a time series index that ranges from 0 to 100, representing the relative popularity of the term over time. Figure C.1 in Appendix C provides an illustration of Google Trends data for searches for Work-In-Suit 454 "tiny dancer"[186] for the period January 2022 through October 2025 in the United States.[187] Further details on Google Trends are also provided in Appendix C.

131. In my analysis, I consider the lyrics of the 499 Works-In-Suit, and I analyze search interest specific to each song. To obtain the weekly song index and lyrics index for each of these songs, I use two main criteria for selecting my sample of queries. First, a query must identify a relevant use of the song title, i.e., it must relate to one of the 499 Works-In-Suit. Second, I need to distinguish lyrics-focused searches from non-lyrics-focused searches.

132. Given that song titles vary in how long, unique, and recognizable they are, I need a way to assess whether users were actually searching for the Works-In-Suit when they put a song

---

[184] Specifically, my estimation is at the weekly level starting January 2, 2022, and ending October 25, 2025.

[185] *See* "FAQ about Google Trends Data," Trends Help, available at https://support.google.com/trends/answer/4365533?hl=en, accessed on 12 Dec 2025.

[186] I.e., searches for the song "Tiny Dancer" by Elton John, Work-In-Suit 454. For all searches, I use lowercase letters.

[187] My sample begins the week of January 2, 2022 and ends the week of October 19, 2025.

title into the search bar. When a song has a simple title, a generic search for "title" may lead a user to find unrelated subjects, and the search therefore would need to include the artist's name to retrieve more relevant results. However, such a problem would not arise for very popular and uniquely named songs whose titles are well-known. To illustrate this, consider the two songs: "Maps," Work-In-Suit 14, by Maroon 5, and "Tiny Dancer," Work-In-Suit 454, by Elton John. A search for "maps" will not only include searches relevant to the song but also include unrelated subjects in cartography or geography. Hence, the inclusion of the artist's name, e.g., "maps maroon 5," will improve the relevance of the search by requiring the artist's name to be a part of the search results. On the other hand, a search for "tiny dancer" is already likely to refer to the Work-In-Suit, and users may be less inclined to include the artist's name. Therefore, the inclusion of the artist's name is an additional search criterion that could result in undercounting actual searches for the song because the search criterion requires the artist's name to be a part of the search results.

133. To address this tradeoff empirically, I run the statistical analysis (whose details I describe in the next subsection on methodology, see Section III.B.2.c)(3) below) in three ways with different combinations of queries on Google Trends:

- First, I run the analysis by comparing search queries on Google Trends for lyrics and songs using the song title alone. That is, I compare the Google Trends search query results for "maps lyrics" and "maps song." (The results are presented in Table 2, where this group is represented by column B, "Title.")
- Second, I run the analysis again comparing Google Trends search queries for lyrics and songs using *both* the song title and the artist. That is, I compare the Google Trends search query results for "maps maroon 5 lyrics" and "maps maroon 5 song." (See Table 2 below. This group is represented by column C, "Title Artist.")
- Finally, I create a "mixed" group based on a measurement I developed of whether a song is more like "maps" (where adding the artist name in the query would likely be helpful) or more like "tiny dancer" (where adding Elton John's name is more likely to be limiting). I discuss this measurement, which I call the "Index Gap," in greater detail in Appendix C. For this mixed group, if the Index Gap determines that a song is more like "maps", I use the Google Trends indices for "maps maroon 5 lyrics" and "maps maroon 5 song" (that is, both title and artist) in my sample. If the song is more like "tiny dancer", I use the Google Trends indices for "tiny dancer lyrics" and "tiny dancer song" in my sample. Because this specification balances the trade-offs discussed above, this is my preferred specification. (See Table 2 below. This group is represented by column D, "Mixed.")

134. In Appendix C, I provide further details of how I construct the three sets of queries described above and include a discussion of queries with low or no search interest.

(2) Traffic to Claude Results

135. I next describe the data I use to measure traffic to Claude.ai for the Google search empirical analysis. To measure traffic to Claude.ai, I use the same method as in my analysis of visits to lyrics websites above (see Section III.B.2.b).[188] Because this Google search relies on a song-by-song analysis, unlike the web traffic study, this time I use [redacted]. I use this as the effective date that the guardrails were strengthened, after which Claude.ai traffic is set to zero. Because CCS is a guardrail that blocks certain outputs from being produced [redacted] I find that it is the more relevant guardrail when conducting a song-by-song analysis, as I am doing here. For ease of comparison, I further rescale this measure to a range of 0 to 100 to be on the same scale as the Google Trends indices between July and [redacted]. I additionally examine the robustness of my results using the [redacted] guardrails which have been previously described.[189]

136. As an alternative specification, I also consider Google Trends search interest for Claude as another measure of Claude usage. Appendix C explains the construction and details of this measure. As before, I impute zero after [redacted] to mute the relationship between Google searches for "claude ai"[190] and Google searches for lyrics. The corresponding Claude interest variable is illustrated in Figure C.5.

(3) Methodology

137. As I did in Section III.B.2.b) above, I use the statistical technique of regression analysis to estimate the effect of visits to Claude.ai on Google search interest for lyrics (for which I use the lyrics index) on Google Trends. In my regression model, I include as a benchmark the

---

[188] As in the previous analysis, I take July 11, 2023 as the release of Claude.ai, which is the start date of the traffic.

[189] As noted previously, I understand from Counsel that for one of the 499 Works-In-Suit, the incorrect song was added to the [redacted]. Plaintiffs asserted infringement of "Joy to the World" by Three Dog Night, but the work that was added to the [redacted] in November 2023 was the holiday song, "Joy to the World" (Work-In-Suit 273). I further understand from Counsel that the correct song was added to the [redacted] in December 2025. This does not impact my analysis as I exclude this song from my analyses except when I estimate robustness analysis with [redacted] guardrails.

[190] Since "claude" is a common name with relevance prior to the inception of Claude, I used "claude ai" instead in the primary specification. As with Claude traffic, I also rescale the "claude ai" index to 0-100 after imputing 0.

search interest for the songs themselves (for which I use the song index). If Plaintiffs' theory of harm is correct, then Claude usage should cause a reduction in the lyrics index relative to the song index.[191] This would be a consequence of people turning away from searching for lyrics on Google—they would instead be turning to Claude.ai—while continuing to turn to Google when searching for the song otherwise. Accordingly, the song index is an appropriate control to include in my model in order to account for general interest in the song.

138. Furthermore, as discussed in Section III.B.2.b) above regarding traffic to lyrics websites, I control my results for factors that may affect the Google Trends indices data, including characteristics unique to individual songs (such as its genre, style, or producer), which I will call "song-fixed effects," and patterns affecting people's musical interests over time (such as interest in holiday music between Thanksgiving and the New Year, or love songs on Valentine's Day), which I previously define as "month-year fixed effects." I include song-fixed effects in my model to account for any characteristics that are unique to individual songs,[192] and I include month-year fixed effects to account for potential trends impacting search interest for songs over time.[193] I include further details of the regression model in Appendix C.

(4) Results

139. Based on my empirical study and consistent with economic theory described above, I do not find a *statistically significant* effect of Claude usage on search interest for lyrics relative to interest in song searches in my preferred specification. This leads me to conclude that the launch of Claude.ai in July 2023 did not have any economically or statistically significant impact on online lyrics consumption measured in Google search trends, which is inconsistent with Plaintiffs' theory of harm.

---

[191] Lyrics were accessible by Claude until the implementation of guardrails but song audio was not.

[192] For example, there is a variable for "7 Rings" by Ariana Grande which takes the value of 1 if the observation belongs to this song and 0 for all other songs. This variable is applied to each song and absorbs song-specific characteristics.

[193] As discussed in Section III.B.2.c)(3) above, song-level and time-fixed effects account for unobserved variations across units and over time. The model includes month-year fixed effects to flexibly control for common seasonal patterns and other unobserved time-varying factors, such as broad shifts in overall online activity, that affect all websites over time. Month–year fixed effects are widely used in the economics literature to control for common changes over time that affect all units at once, such as seasonal patterns or broad shifts in overall activity. For example, *see* Erik Brynjolfsson, Danielle Li, Lindsey Raymond, Generative AI at Work, *The Quarterly Journal of Economics*, Vol. 140, No. 2 (May 2025), at pp. 889–942, https://doi.org/10.1093/qje/qjae044.

140. Table 2 summarizes the results of my empirical analysis. As I note above, I run my regression with several alternate specifications that use combinations of the "Title" version and "Title Artist" version of the indices. As shown in Table 2, the "Title" specification in column B uses the "title lyrics" and "title song" query indices. The "Title Artist" specification in column C uses the "title artist lyrics" and "title artist song" query indices. The "Mixed" specification in column D uses either the "title" or the "title artist" query indices for songs depending on the average gap between the search indices for "title" and the "title artist," as described in Appendix C. As I state above, the title-artist index gap signifies the relative difference between search queries of song titles compared to search queries of song titles with the artist. If the gap is less than 20 or less—meaning the title alone is more likely to be unique or recognizable—then I use the "Title" query indices. If the gap is more than 20—meaning the title is more generic and unlikely to identify the correct song—I use the "Title Artist" query indices. I consider the Mixed specification to be the most robust since it relies on the title-artist index gap measurement to address the nuances of how users use Google search to query for song lyrics—sometimes users will search for just the song name if it is a recognizable song title, while other times users will search for the song plus the artist name.[194]

141. In Table 2, "Lyrics" is an "indicator variable" where I have assigned a value of zero for the song indices and one for the lyric indices.[195] Claude Traffic, as described above, is the traffic to Claude.ai. The interaction term (Lyrics x Claude Traffic) is the term of interest in my analysis, as it shows the effect of Claude usage on the lyric index relative to the song index.

142. If Plaintiffs' theory of harm were correct, I would expect the estimated coefficient of the interaction terms to be negative and statistically significant, because users would be shifting away from Google searches for lyrics (relative to their searches for songs) and relying on Claude for lyrics instead. However, across all three approaches (detailed in paragraph 140 above) that I use to compare the relative frequency of lyrics-focused searches to song-focused searches, none

[194] With an index gap of 20, 112 songs use the "Title" indices and 382 songs use the "Title Artist" indices. In Appendix C, I test the robustness of my results to an index gaps of 10, in which 65 songs use the "Title" indices and 429 songs use the "Title Artist" indices, and to an index gap of 30, in which 173 songs use the "Title" indices and 321 songs use the "Title Artist" indices.

[195] A "indicator variable" is a variable that is equal to either zero or one to distinguish two groups. Consider a variable for "student." if this variable is equal to one then the observation belongs to a student; if this variable is equal to zero, then the observation does not belong to a student.

of the estimated coefficients of the interaction term is negative and statistically significant (reading left to right across row [1]: 0.019, 0.002, and 0.003).

143. In sum, these results do not indicate that the introduction of Claude caused a meaningful impact on search interest for lyrics on Google. None of the three specifications shows a negative and statistically significant effect of Claude.ai traffic on web traffic for Google searches for lyrics. In the Mixed specification, I find an estimate that is statistically insignificant, meaning that no effect of Claude usage on the lyric index is detected that is statistically different from zero. I similarly find a positive and statistically insignificant effect in the "Title Artist" specification. The effect I find in the "Title" specification is positive and statistically significant. However, the magnitude of this estimate is too small to be interpreted as having an economically meaningful effect. Additionally, as described earlier, the title specification may include searches for unrelated topics such as geography for "Maps," making the Mixed specification our preferred model. As I do not find an effect of the introduction of Claude.ai on Google searches for lyrics, I do not see any empirical evidence supporting Plaintiffs' theory of harm.

Table 2: Baseline Model with Alternate 'Title' and 'Title Artist' Specifications

		Title	Title Artist	Mixed (20)
	[A]	[B]	[C]	[D]
[1]	Lyrics x Claude Traffic	0.019***	0.002	0.003
[2]		0.003	0.002	0.002
[3]	Lyrics	11.438***	3.092***	3.500***
[4]		0.068	0.049	0.051
[5]	Claude Traffic	-0.021***	-0.002	-0.006
[6]		0.005	0.003	0.003
[7]	Observations	198,204	198,204	198,204
[8]	R Squared	0.596	0.187	0.245

Sources: WIS Chart provided by Counsel; Google Trends data, accessed via SearchAPI.

Table results are obtained from estimates of my baseline regression model using 498 WIS provided by counsel and include song fixed effects and month-year fixed effects. "Joy to the World" (Work-In-Suit 273) is excluded from estimation because it was [redacted], as explained in text. Observations are at the week-WIS level. The sample period is full weeks contained in January 2022 – October 2025. Lyrics is an indicator that equals 1 for lyrics queries and 0 for song queries. Claude Traffic is imputed 0 after [redacted], and rescaled to a range of 0-100. Standard errors are provided in the rows below coefficient estimates. Standard errors are heteroskedasticity-robust, and significance stars are reported at the 10% (*), 5% (**) and 1% (***) level.
[B] Title specification includes indices of queries "[title name] lyrics" and "[title name] song".
[C] Title Artist specification includes indices of queries "[title name] [artist name] lyrics" and "[title name] [artist name] song"
[D] Mixed (20) specification includes indices of queries "[title name] lyrics" and "[title name] song" if the Index Gap is 20 or less, otherwise includes indices of queries "[title name] [artist name] lyrics" and "[title name] [artist name] song".

IV. CONCLUSION

144. To summarize, after evaluating the impact of Anthropic's alleged conduct across three potential markets, I conclude that the alleged conduct did not harm any real or potential market for data sufficient to train a general-purpose LLM, any real or potential market for lyric training data, or the existing market for online lyrics display.

145. First, a market for all of the data needed to train a general-purpose LLM did not exist at the time of the alleged use and is unlikely to emerge in the future. The scale and diversity of data required, combined with the prohibitive transaction costs of licensing from numerous rightsholders, make this market economically unviable. Recent licensing deals do not alter this conclusion; these agreements represent rare cases where the value of specific data exceeds acquisition costs, not a broader market trend.

146. Second, while I continue to believe the relevant market should be a market for data needed to train a general-purpose LLM, I also find that the market for the use of song lyrics to train a general-purpose LLM did not exist at the time of the alleged use and is unlikely to exist in the future. Lyrics provide minimal value for LLM training relative to the high costs of acquisition, making such a market economically impractical.

147. Since neither of these markets existed at the time of the alleged use and neither is likely to emerge in the future, I conclude that Plaintiffs cannot have been economically harmed in such markets.

148. Finally, I conclude that Anthropic's alleged conduct did not harm the existing market for online lyrics display. Economic principles suggest that harm is unlikely because the degree of substitution between typical LLM use cases and online lyrics services is too low to affect demand. Empirical analyses conducted after Claude's commercial release confirm no measurable impact consistent with Plaintiffs' theory of harm.

149. Combined, my conclusion is that Anthropic's conduct did not harm any existing or potential market for the data needed to train a general-purpose LLM, any existing or potential market for lyric training data, or the existing market for the display of online lyrics. Thus, Anthropic's conduct had no adverse effect on the potential market for or value of Publishers' copyrighted works.