[Counsel on signature page]

UNITED STATES DISTRICT COURT

NORTHERN DISTRICT OF CALIFORNIA

SAN JOSE DIVISION

CONCORD MUSIC GROUP, INC., ET AL.,

Plaintiffs,

v.

ANTHROPIC PBC,

Defendant.

Case Number: 5:24-cv-03811-EKL-SVK

JOINT DISCOVERY STATEMENT REGARDING THRESHOLD DISPUTES AS TO PLAINTIFFS' DISCOVERY REQUESTS TO DEFENDANT

Hon. Eumi K. Lee
Magistrate Judge Susan van Keulen

Plaintiffs ("Publishers") and Defendant ("Anthropic") respectfully submit this joint statement regarding two threshold disputes relating to Publishers' discovery requests to Anthropic.

Fact discovery began in March 2024 while this case was pending in the Middle District of Tennessee. This case was transferred to this District on June 26, 2024 and re-assigned to Judge Eumi K. Lee and Magistrate Judge Susan van Keulen on August 21, 2024. The Parties have not yet appeared before Judge Lee for a Case Management Conference and a new case schedule has not been entered following transfer and re-assignment; in their Joint Case Management Statement, Publishers and Anthropic request fact discovery deadlines of March 21, 2025 and January 20, 2025, respectively. Dkt. 224-1. Publishers' March 22, 2024 Requests for Production and Interrogatories ("Requests") are attached as Exs. 1-2. Anthropic's responses and objections are attached as Exs. 3-5. The Parties' joint chart is Ex. 6. The Parties met and conferred by videoconference on May 9, June 3, and August 15, 2024, and by email, but could not reach resolution.

I. Publishers' Position

Publishers are among the world's leading music publishers. Anthropic is an artificial intelligence ("AI") company valued at over $18 billion that does not pay copyright holders for the content it relies upon to build and operate its AI models. Publishers allege Anthropic systematically infringes their copyrighted song lyrics through its series of "Claude" AI models, by copying their lyrics both as the input to train those AI models and in the output those models generate. Unlike many other recent AI cases, however, this case includes extensive evidence of Anthropic's systematic distribution of verbatim and near-verbatim copies of Publishers' lyrics as output.

Anthropic has so far refused to provide relevant and proportional discovery on a wide range of key issues. To narrow the issues that must be brought to the Court, Publishers first request guidance in resolving two threshold disputes as to general objections Anthropic asserts across all Requests. These objections are unfounded and seek to improperly narrow the scope of discovery.

A. Anthropic cannot broadly exclude key Claude AI models from discovery.

First, across all Requests, Anthropic broadly refuses to provide ***any*** discovery on key Claude AI models subject to Publishers' claims, including (1) the Claude models Anthropic sells to business customers, (2) the earlier models Anthropic used to develop its current models, and (3)

the models Anthropic is currently developing. Instead, Anthropic tries to restrict discovery to only "commercially available, generally accessible Claude models"—*i.e.*, solely those models currently available to individual users through Anthropic's website and mobile applications. Ex. 5 at 5. But Anthropic's infringement through the Claude models on its website and mobile apps is just the ***tip of the iceberg***, and Publishers' claims are in no way restricted to those models only. Publishers are "entitled to seek discovery on [their] theory of the facts and the law," and are "not limited in discovery by the[ir] opponent's theory." *Wagner Aeronautical, Inc. v. Dotzenroth*, 2022 WL 198805, at *2 (S.D. Cal. Jan. 20, 2022) (quotation omitted). Anthropic has no legitimate basis for blocking relevant and proportional discovery into the AI models it sells to business customers, its earlier models, and its models in development, all of which bear directly on Publishers' claims.

Claude API models for business customers. Anthropic sells and licenses Claude AI models to "thousands" of business customers as an Application Programming Interface ("API"). *See* https://www.anthropic.com/news/claude-2. This Claude API is squarely covered by Publishers' claims, as Publishers explicitly allege that Anthropic infringes their lyrics through the training and output of its Claude API, just as it infringes through the separate versions of Claude on Anthropic's website. *See, e.g.*, Compl. ¶¶ 52, 65, 93–97, ECF No. 1. Indeed, Anthropic's counsel has admitted that Claude "API and Claude.ai [the website version] are merely different surface areas to interact with the ***same core underlying models***." S. Sampoli Email (June 14, 2023) (emphasis added). But where Anthropic's Claude API models generate separate infringing output and yield independent revenues, Anthropic cannot unilaterally exclude Claude API from discovery on a blanket basis.

Discovery as to Claude API models and outputs is critically relevant across Publishers' claims and plainly proportional, while discovery as to Claude website models clearly does not "encompass" Claude API, as Anthropic claims. For one, Publishers allege that Anthropic directly and secondarily infringes their lyrics through the specific outputs of its Claude API models, separate from Claude website models. Indeed, Publishers have already identified voluminous evidence of infringing output through Claude's API console, as well as API on Quora's Poe and DuckDuckGo. *See, e.g.*, Seymour Decl. ¶¶ 4–7, ECF No. 49 (explaining "[h]alf of the files in Exhibit B were preserved based on a prompt submitted via Claude's API"); Leonard Decl. ¶¶ 7–

11, ECF No. 50. But that is ***not*** exhaustive. While Anthropic admits licensing Claude API to "thousands" of commercial customers, it has not even disclosed the identities of the vast majority of those customers—let alone the extent to which it copied Publishers' lyrics in the training and specific unique outputs of each of those models, which goes directly to Publishers' infringement claims. Publishers' contributory and vicarious liability claims are also premised in significant part on Claude API, because API is a means by which Anthropic may (a) induce, cause, or materially contribute to the infringing conduct of others, and (b) profit from direct infringement while declining to stop it. Another critical issue is the extent to which Anthropic's alleged "guardrails" against infringing output differ between its Claude website and API models, which bears on the scope of infringement, willfulness, and harm. Anthropic also earns a substantial portion of its revenues through API, which goes to vicarious liability (*e.g.*, Anthropic's direct financial benefit from infringement), infringing profits, and damages more broadly. Anthropic's attempt to broadly exclude API from discovery would vastly distort the financial picture of Anthropic's infringement.

"[Publishers'] requests are not rendered disproportional merely because they . . . could require a large volume of responsive documents," particularly where that is due to the large scale of the infringement. *Teixeira v. BMW of N. Am.*, 2023 WL 4291798, at *5 (C.D. Cal. Apr. 24, 2023).

The fact that Anthropic has agreed to disclose limited information about a ***single*** Claude API, Quora's Poe, cannot excuse its unjustified refusal to provide ***any*** discovery as to other API. Indeed, the limited information Anthropic provided about Quora's Poe—including significant monthly usage and revenue figures—only underscores the relevance of Claude API more broadly.

Earlier Claude models. Anthropic's current Claude website and API models were built on earlier models that also copied Publishers' lyrics as input and output. Anthropic consistently describes later Claude models as a "continuous evolution" of earlier models, *see, e.g.*, Compl. ¶ 51 (quoting *Model Card and Evaluations for Claude Models* 1, ANTHROPIC (July 8, 2023), https://perma.cc/59SF-PYWA ("Model Card"))—admissions it now simply disregards. Where Anthropic made unauthorized copies of Publishers' lyrics in training its earlier models and in their output, and where it used them to build its current models, those models are plainly covered by Publishers' claims and properly subject to discovery. Such models infringe regardless of whether

they were commercialized or publicly released. *See Campbell v. Acuff Rose Music, Inc.*, 510 U.S. 569, 584 (1994) ("[T]he mere fact that a use is educational and not for profit does not insulate it from a finding of infringement[.]"). Also, where Publishers challenge aspects of the development and operation of Anthropic's models, "previous generations of the product in question [are] relevant and discoverable, regardless [of] whether they share the same characteristics as the model in question." *Green v. Toyota Motor Corp.*, 2012 WL 13065682, at *5 (N.D. Tex. July 18, 2012).

Claude models in development. During the pendency of this case, Anthropic has continued—and accelerated—the conduct that Publishers allege to be infringing. In the less than a year since Publishers filed suit, Anthropic has released ***five*** new Claude models, ECF No. 180-14, and it continues to train even more models, *see, e.g.*, ECF No. 207 at 26. Anthropic does not dispute that these new models are trained on Publishers' lyrics, just as its earlier models were. Where Anthropic has ***already*** trained these models on Publishers' lyrics, it cannot hide them from discovery just because it has not yet commercially released them. The two patent cases cited by Anthropic below, *PersonalWeb* and *Microsoft*, are inapplicable here, including because Publishers allege copyright infringement via training Anthropic has ***already*** undertaken with these models.

In sum, discovery into Claude API and other key models is plainly relevant and proportional in light of the factors considered under Fed. R. Civ. P. 26(b)(1). The discovery is needed to understand the full extent of Anthropic's infringement, through training and output, while blocking such discovery would allow Anthropic to conceal much of its infringement; the amount in controversy is high, in light of the vast scope of infringement, including through models Anthropic seeks to hide; only Anthropic has access to this information, as it does not disclose training data or a full list of customers and closely controls its models; and Anthropic is valued at over $18 billion, but fails to show any particular undue burden of responding to specific Requests. *Id.* Further, given that Anthropic objects to the Claude definition on a blanket basis—without attempting to show lack of relevance or burden of responding to any specific Requests—it fails to "specifically explain the reasons why the request at issue is objectionable" and cannot meet its heavy "burden to show that the discovery should not be allowed." *Humanmade v. SFMade*, 2024 WL 3378326, at *2 (N.D. Cal. July 10, 2024); *see also* Fed. R. Civ. P. 33(b)(4), 34(b)(2).

B. Anthropic cannot unduly restrict the timeframe subject to discovery

Second, Publishers seek discovery from 2020 to the present. Ex. 1 at 5. Anthropic, however, has broadly refused to provide ***any*** discovery prior to September 2022. *E.g.,* Ex. 5 at 6. Anthropic's arbitrary cut-off date would unduly exclude key relevant discovery from 2020, 2021, and 2022.

Anthropic has admitted it began developing and training the Claude models that are the subject of Publishers' claims well before September 2022. Indeed, Anthropic Chief Science Officer Jared Kaplan has stated in multiple sworn declarations in this case that Anthropic began developing and training its Claude models in 2020: "***After over two years of research, development, and training***, Anthropic released Claude as an application programming interface ('API') to certain business users in late 2022 and released a public conversational interface, Claude.ai, as a consumer product in January 2023." ECF No. 67-1 ¶ 8 (emphasis added); *see also, e.g.*, ECF No. 209 ¶ 8 (similar). Other Anthropic admissions likewise underscore that the company engaged in activities relevant to Publishers' claims before September 2022. *See, e.g.*, ECF No. 55 at 2 (stating Anthropic released Claude 1 and 2 models in 2023 "[a]fter years of research, development, and training"); *Model Card* (stating Claude 2, released in July 2023, was built on "2+ years of research").

Despite Anthropic's attempts to now distance itself from these admissions, they make clear that it was engaged in training its AI models—activity at the heart of Publishers' claims—dating back to 2020. Anthropic fails to show this discovery is disproportional. Fed. R. Civ. P. 26(b)(1)).

Nor can Anthropic refuse to submit to any discovery prior to September 2022 simply by claiming that its earlier AI models were "experimental," particularly given (a) the record above that these earlier models were the building blocks for its later Claude models; and (b) that Anthropic's exploitation of Publishers' works to develop and train purportedly noncommercial AI models is still directly relevant to Publishers' infringement claims, *see Campbell*, 510 U.S. at 584.

Proposed compromise: Anthropic's attempts to unduly limit discovery as to its Claude models and the timeframe for discovery should be rejected, and Anthropic should be ordered to respond to the Requests in full. In the alternative, Anthropic should be ordered to provide a list of (1) all its Claude API customers, including for each customer (a) a short description of the API "use case," *see* https://www.anthropic.com/api, and (b) the total number of user prompts and U.S.

revenues for a sample 12-month period (Apr. 2023 to Mar. 2024); (2) each earlier AI model for which Publishers' lyrics were included in the training data and/or that was used to train its current Claude models, including for each the dates development began and concluded; and (3) each model currently under development for which Publishers' lyrics have been included in the training data, including for each the date development began. Once Anthropic discloses this baseline information, the Parties can meet and confer to discuss if it may be possible to narrow any Publishers' Requests.

II. Defendant's Position

Plaintiffs seek unbounded discovery into topics that have no relevance to the actual allegations in their Complaint. Far from "refus[ing] to provide relevant and proportional discovery on a wide range of key issues," *supra* at 2, Anthropic has sought to reasonably limit discovery to the claims and defenses at issue and has consistently looked for opportunities to compromise on the present disputes. For example, while Anthropic originally objected to Plaintiffs' definition of Claude to the extent it referred to any models "other than those commercially available models identified in Plaintiffs' Complaint" (Ex. 3 at 5), it has since agreed to interpret Claude as referring to the two Claude models identified in the Complaint (Claude 1 and Claude 2), as well as five additional versions released to date (Claude 2.1, the Claude 3 model family (Haiku, Sonnet, and Opus), and Claude 3.5 Sonnet). Anthropic has also specifically agreed to produce discovery regarding Poe, a third-party product offered by Quora, based on the Plaintiffs' assertion in their preliminary injunction papers that they used this product to generate outputs containing portions of their works. Despite Anthropic's compromises, Plaintiffs continue to seek irrelevant and burdensome discovery. The Court should not grant Plaintiffs the discovery they seek.

A. Plaintiffs seek discovery related to irrelevant AI models.

Using an overbroad definition of Claude, Plaintiffs seek discovery regarding (1) never-released research AI models predating the conduct at issue in their Complaint; (2) future releases of Claude that have not been and could not yet be the subject of any allegation of harm; and (3) thousands of third-party implementations of the Claude API. None of these categories is sufficiently related to the claims or defenses in this case to warrant such overbroad discovery.

Non-commercial research models. Plaintiffs premise their lawsuit on alleged

infringement occurring during the training and in the outputs of Anthropic's commercial product, Claude, which Anthropic officially released in March 2023. *See Introducing Claude*, Anthropic (Mar. 14, 2023), https://www.anthropic.com/news/introducing-claude.[1] Each of the factual allegations in the Complaint refers to the first two versions of this commercially available AI product. *See, e.g.*, Compl. ¶¶ 50–52 (alleging that Anthropic makes Claude accessible to users via a chat interface and through an API); *id.* ¶¶ 93–96 (alleging that Anthropic profits from its chat and API Claude products); *id.* ¶¶ 70–72 (referring to *Claude* outputs and *Claude*'s capabilities). There are no factual allegations relating to models that Anthropic never made commercially available. Yet Plaintiffs seek discovery into *all* AI models that Anthropic ever developed, including research models not part of the "Claude" family and never released to the public.

This discovery has no relevance to Plaintiffs' claims. Claude 1 was not built on top of any earlier model developed by Anthropic. Instead, it was trained from scratch in October 2022 ahead of Anthropic's initial commercial product launch. Discovery into earlier research models is either irrelevant or cumulative of the discovery that Anthropic has already agreed to produce for the commercial Claude models that are the subject of Plaintiffs' complaint.

Plaintiffs' out-of-circuit case law does not translate to Anthropic's AI products. The Court in *Green v. Toyota Motor Corp.* found, in a products liability suit, that documents related to "one previous generation" of the product in question (a Toyota 4Runner) were relevant and discoverable. 2012 WL 13065682, at *5 (N.D. Tex. July 18, 2012). Anthropic's private research models are not "previous generation[s]" of Claude. They are standalone models made for an entirely different purpose and an entirely different audience. Any non-Claude models developed for research purposes before the development and launch of Claude have no relevance to the claims and defenses and are not properly subject to discovery. FRCP 26(b)(1). Plaintiffs' proposal that Anthropic search the training corpus of unreleased models for asserted works is burdensome and disproportionate absent any specific allegations in the complaint related to those models.

Future Claude Models. Anthropic originally proposed limiting the definition of Claude

[1] The model was released to select business customers in late 2022. J. Kaplan Decl., Dkt. 209 ¶ 8.

to the models identified in the Complaint, *i.e.*, Claude 1 and Claude 2. *See* Ex. 3 at 5. As a compromise, Anthropic has since offered to produce discovery related to all Claude models that have been publicly released to date. Plaintiffs nevertheless continue to insist that Anthropic produce discovery regarding all future models, notwithstanding that (1) Plaintiffs' own definition does not cover such models, *see* Ex. 1 at 2 (defining Claude as "any general purpose generative-AI large language model *created, developed, maintained, or made publicly available* by Anthropic" (emphasis added)); and (2) the collections and searches Anthropic has run pursuant to the parties' agreement do not capture discovery related to models after Claude 3.5 Sonnet.

Requiring discovery into unreleased models still in development, nearly a year after Plaintiffs filed their Complaint, would be unwarranted and burdensome; it would require Anthropic to do precisely what the parties agreed to avoid when they agreed to a search and collection cutoff : "re-run collections or searches *ad infinitum* after the date upon which Publishers served their discovery requests." 4/22/24 Pls. Ltr. at 13. Unreleased models still under development are irrelevant to Plaintiffs' claims, and discovery focused on this unlimited category of models would be disproportionate to the needs of the case. *See, e.g., PersonalWeb Techs., LLC v. Google Inc.*, 2014 WL 4088201, at *4 (N.D. Cal. Aug. 19, 2014) (holding that Google "need not produce discovery related to a version of a product which has not been released and is still being developed as "non-'live' product cannot be evaluated for infringement"); *Microsoft Corp. v. Multi-Tech Sys., Inc.*, 2001 WL 37130770, at *8 (D. Minn. Dec. 14, 2001) (similar).

Irrelevant third-party products utilizing the Claude API. Plaintiffs' argument for discovery into thousands of Claude API customers fundamentally misunderstands the technology at issue. By intentionally conflating the technical meaning of "model," Plaintiffs use their overbroad Claude definition as a pretext for seeking discovery into otherwise irrelevant third-party prompt, output, and revenue data. Anthropic has not "unilaterally" or "broadly refuse[d] to provide any discovery" regarding the Claude API as Plaintiffs suggest. To the contrary, Anthropic has offered a reasonable compromise, which Plaintiffs rejected in favor of raising this dispute.

Because the Claude API is simply a different surface area for interacting with the same underlying Claude models for which Anthropic has agreed to produce discovery, Anthropic's

proposed definition of Claude already encompasses the API product offered to its commercial customers in all relevant respects. Consistent with that definition, Anthropic will produce discovery related to the training and development of the relevant Claude models that are accessible via the Claude API. Defining Claude to also include the thousands of third-party products developed by Anthropic's business customers *using* the Claude API would be invasive, burdensome, and ultimately irrelevant to the claims and defenses in this case. Plaintiffs' conclusory allegation that "[e]ach of Anthropic's commercial customers . . . infringes Publishers' copyrights when these AI models generate output copying or relying on Publishers' lyrics," Compl. ¶ 95, is meaningless in the absence of a single factual allegation to support the implausible theory that each of Anthropic's thousands of *business customers* is using its generative AI software to generate allegedly infringing outputs of *song lyrics*. [2] *Build with Claude*, Anthropic, https://www.anthropic.com/api (noting common business use cases, e.g., coding, productivity enhancements, and customer support triage). Plaintiffs have shown no proportional relevance for their requested discovery into these businesses' uses of Claude.

Plaintiffs' claim that they have "already produced voluminous evidence of infringing output through Claude API, including Claude API through Quora's Poe and DuckDuckGo" is wrong. They have only submitted evidence of certain outputs that they allegedly elicited themselves using Poe—not DuckDuckGo or any other third-party product. Candore Decl., Dkt. 96 (identifying Poe). In light of those submissions, Anthropic agreed to produce discovery relating to Quora. Neither of Plaintiffs' cited declarations identifies any additional third-party API usage to generate allegedly infringing outputs. Seymour Decl., Dkt. 49; Leonard Decl., Dkt. 50. Plaintiffs' counsel confirmed that Mr. Seymour did not use a third-party product but rather "used the Claude API made publicly available through Anthropic's web console." *See* 8/30/24 T. Chung Email.

Plaintiffs' attempt to seek discovery into thousands of commercially sensitive third-party business customers is profoundly disproportionate to the needs of this case. Even Plaintiffs'

[2] Plaintiffs' argument that their contributory and vicarious liability claims are "premised in significant part on Claude API" is thus incorrect. *See supra* at 3.

proposed "compromise"—requiring Anthropic to turn over the names, use case descriptions, and prompt and revenue data about thousands of customers that have no connection to the allegations in this case—would be intrusive and burdensome. Absent a pleaded basis for why the data associated with these customers are relevant to Plaintiffs' claims, such discovery is inappropriate.

B. Plaintiffs seek discovery related to an irrelevant time period.

Plaintiffs' demand for discovery going back to 2020 is similarly overbroad. Anthropic was not even founded until 2021, as Plaintiffs concede. Compl. ¶ 49. Anthropic did not decide to launch a commercial AI product until September 2022, and trained Claude 1 from scratch in October 2022. Plaintiffs' reliance on general statements regarding the "years of research, development, and training" predating Claude's commercial release cannot change these bedrock facts. Besides, the "research, development, and training" in these statements relate to AI systems generally—not Claude specifically. *See Model Card, supra* at 5. Anthropic proposed the September 2022-to-present timeframe because that was when Anthropic decided to launch and proceeded to develop its first commercial product, Claude 1. Prior to that date, Anthropic focused on building experimental models unavailable to the public to support their research. *Anthropic Raises Series B*, Anthropic, https://www.anthropic.com/news/anthropic-raises-series-b-to-build-safe-reliable-ai. Plaintiffs have failed to show any relevance of these separate, unreleased pre-Claude models to the claims and defenses in this litigation, which concern the alleged training of commercially released models allegedly capable of reproducing their song lyrics. *See supra* at 7. Plaintiffs' attempt to seek discovery into the time period predating the inception of Claude should be rejected.

Proposed compromise: Anthropic proposes to interpret Claude as referring to any and all Claude models that are commercially available and generally accessible as of today, September 26, 2024—i.e., Claude 1, Claude 2, Claude 2.1, the Claude 3 model family (Haiku, Sonnet, and Opus), and Claude 3.5 Sonnet. This includes several models released after those implicated by the complaint. Anthropic will also produce discovery regarding Quora's implementation of the Claude API. Finally, Anthropic will consider producing discovery regarding specific additional business customers, to the extent within Anthropic's possession, if Plaintiffs can provide a Rule 11 basis for alleging those products have been used to infringe the works asserted in this action.

Dated: October 1, 2024

By: /s/ Timothy Chung

OPPENHEIM + ZEBRAK, LLP
Matthew J. Oppenheim
Nicholas C. Hailey
Audrey L. Adu-Appiah
(admitted *pro hac vice*)
4530 Wisconsin Ave., NW, 5th Floor
Washington, DC 20016
Telephone: (202) 480-2999
matt@oandzlaw.com
nick@oandzlaw.com
aadu-appiah@oandzlaw.com

Jennifer Pariser
Andrew Guerra
Timothy Chung
(admitted *pro hac vice*)
461 5th Avenue, 19th Floor
New York, NY 10017
Telephone: (212) 951-1156
jpariser@oandzlaw.com
andrew@oandzlaw.com
tchung@oandzlaw.com

COBLENTZ PATCH DUFFY & BASS LLP
Jeffrey G. Knowles (SBN 129754)
One Montgomery Street, Suite 3000
San Francisco, CA 94104
Telephone: (415) 391-4800
ef-jgk@cpdb.com

COWAN, LIEBOWITZ & LATMAN, P.C.
Richard S. Mandel
Jonathan Z. King
Richard Dannay
(admitted *pro hac vice*)
114 West 47th Street
New York, NY 10036-1525
Telephone: (212) 790-9200
rsm@cll.com
jzk@cll.com
rxd@cll.com

Attorneys for Plaintiffs

Respectfully submitted,

By: /s/ Joseph R. Wetzel

LATHAM & WATKINS LLP
Joseph R. Wetzel (SBN 238008)
joe.wetzel@lw.com
Andrew M. Gass (SBN 259694)
andrew.gass@lw.com
Brittany N. Lovejoy (SBN 286813)
britt.lovejoy@lw.com
505 Montgomery Street, Suite 2000
San Francisco, California 94111
Telephone: +1.415.391.0600

Sarang V. Damle
(admitted *pro hac vice*)
sy.damle@lw.com
555 Eleventh Street NW, Suite 1000
Washington, DC 20004
Telephone: +1.202.637.2200

Allison L. Stillman
(admitted *pro hac vice*)
alli.stillman@lw.com
1271 Avenue of the Americas
New York, New York 10020
Telephone: +1.212.906.1747

Attorneys for Defendant

SIGNATURE ATTESTATION PURSUANT TO CIVIL L.R. 5-1(h)

Pursuant to Civil L.R. 5-1(h), I hereby attest that concurrence in the filing of this document was obtained from all other signatories of this document. I declare under penalty of perjury that the foregoing is true and correct.

Dated: October 1, 2024

/s/ Timothy Chung
Timothy Chung