EXHIBIT 112

(FILED UNDER SEAL)

1

IN THE 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

HIGHLY CONFIDENTIAL--ATTORNEYS' EYES ONLY

VIDEOCONFERENCED and VIDEO-RECORDED
DEPOSITION of
DARIO AMODEI, Ph.D.
San Francisco, California 94111
Thursday, January 8, 2026

Reported Stenographically by:
MARY J. GOFF
CSR No. 13427
WA CSR No. 21030779
Job No. 614290
PAGES 1-158

2

HIGHLY CONFIDENTIAL--ATTORNEYS' EYES ONLY

VIDEOCONFERENCED and VIDEO-RECORDED DEPOSITION of DARIO AMODEI, Ph.D., Volume I, taken on behalf of Plaintiff Publishers, Oppenheim + Zebrak, LLP, at WilmerHale, 50 California Street, Suite 3600, San Francisco, California 94111, beginning at 1:44 p.m. and ending at 4:30 p.m., on Thursday, January 8, 2026, before MARY J. GOFF, California Certified Shorthand Reporter No. 13427 and WA CSR No. 21030779.

3

APPEARANCES:

For Plaintiff Publishers
Oppenheim + Zebrak, LLP
BY: NICHOLAS C. HAILEY
Attorneys at Law
4530 Wisconsin Avenue, NW
5th Floor
Washington, DC 20016
202-480-2999
nick@oandzlaw.com

For Defendant Anthropic PBC
WilmerHale
BY: SONAL MEHTA
Attorneys at Law
2600 El Camino Real
Suite 400
Palo Alto, California 94306
sonal.mehta@wilmerhale.com
628-235-1026

ALSO PRESENT: Natalie Naugle, Brian Israel, In-house counsel Anthropic
Ellen Stephano, Esquire, WilmerHale (remote Boston)
Videographer: Addison McNaughton

4

INDEX

WITNESS		EXAMINATION
DARIO AMODEI, Ph.D.		
Volume I		
	ATTORNEY HAILEY	6
	ATTORNEY MEHTA	--

NUMBER	DESCRIPTION	PAGE
Exhibit 246	1-29-24 Slack message Anthropic_0000605700-785	9
Exhibit 247	Transcript: Ezra Klein Pages 1-18	45
Exhibit 248	Exhibit 40 Anthropic_0000462964-968	89
Exhibit 249	Anthropic Plan for 2023 Anthropic_0000463241-255	108
Exhibit 250	Why Anthropic CEO Dario Amodei Spends So Much Time Warning of AI's Potential Dangers	136
Exhibit 251	AI Misalignment Risk Anthropic_0000478585	142
PREMARKED EXHIBITS:		
Exhibit 235	1-8-2023 email Anthropic Plan	123

5

San Francisco, California
Thursday, January 8, 2026
1:44 p.m.

THE VIDEOGRAPHER: Here begins Media Number 1 in the videotaped conception of Dr. Dario Amodei, in the matter of Concord Music Group Inc., et al., versus Anthropic PBC, in the United States District Court, Northern District of California, San Jose Division. The Case Number is 24-CV-03811-EKL (SVK).

Today's date is January 8, 2025, and the time ton the video monitor -- sorry -- January 8, 2026, and the time on the video monitor is 1:44 p.m.

The videographer today is Addison McNaughton, representing Planet Depos. And this video deposition is taking place at 50 California Street, Suite 3600, in San Francisco, California 94111.

Would counsel please voice identify themselves and state whom they represent.

ATTORNEY HAILEY: This is Nick Hailey, with the law firm of Oppenheim + Zebrak, on behalf of the Music Publisher Plaintiffs. I'm joined with my colleague by zoom, Tim Chung.

ATTORNEY MEHTA: Good afternoon, Sonal

6

Mehta, from WilmerHale, on behalf of Anthropic. With me today are Natalie Naugle and Brian Israel, who are in-house Counsel at Anthropic.

THE VIDEOGRAPHER: The court reporter today is Mary Goff, representing Planet Depos, the witness will now be sworn.

THE COURT STENOGRAPHER: My name is Mary Goff. I am a California Certified Shorthand Reporter, CSR Number 13427. Today's proceedings are being captured by stenographic means.

DARIO AMODEI, Ph.D.

Being first duly sworn or affirmed to testify to the truth, the whole truth, and nothing but the truth, was examined and testified as follows:

ATTORNEY MEHTA: You can put your hand down.

THE DEPONENT: Okay.

EXAMINATION

BY ATTORNEY HAILEY:

Q Good afternoon. Name is Nick Hailey. I represent the Plaintiff, Music Publishers, in this case.

Can you go ahead and please state your name and home address for the record?

A Are you speaking to me?

7

Q Yes.

A Can I state my name and home address for the record?

ATTORNEY MEHTA: And I'm actually just going to -- for security reasons, I don't think you need his home address. I am --

A I -- I, in fact -- no. I, in fact, don't even know it. Dario Amodei is my name.

Q (BY ATTORNEY HAILEY) And what city do you live in?

A I live in Woodside.

Q And you -- your testimony is you don't know your home address?

ATTORNEY MEHTA: You can answer the question.

A I don't -- I don't remember it. That's right.

Q (BY ATTORNEY HAILEY) And do you prefer Mr. Amodei or Dr. Amodei?

A Mr. Amodei is fine.

Q And, Mr. Amodei, you understand you're under oath today, correct?

A That is correct.

Q And we have very limited time for the deposition today and a lot of ground to cover, so

8

it's important that you listen carefully to the questions that I ask and that you answer those questions directly and succinctly.

Do you understand that?

A Yes.

Q And if you do not understand a question, please ask me to clarify. Otherwise, if you answer my question, I will assume that you understood it.

Is that understood?

ATTORNEY MEHTA: Form.

You can answer.

A Yes.

Q (BY ATTORNEY HAILEY) Is there any reason why you cannot give full, complete, and accurate testimony today?

A I mean, to the best of my ability, sure.

Q Mr. Amodei, you're Anthropic's CEO and a founder of the company, correct?

A That is correct.

Q Who are Anthropic's competitors in the AI industry?

A Who are Anthropic's competitors in the AI industry? You know, that's a subjective question.

It depends on who you -- who you list as

13

A Okay.
Q Do you see the Slack conversation in the middle of page 738 beginning with the message from Andy Stewart on February 29, 2024, at 12:02 p.m.?
A Yes, I do.
Q And he writes:
> Here's another bad one with image only as prompt. Screenshot of text. This was first try with TLB.
> Do I read that correctly?

A I do see it. There's a screenshot. But I -- the -- either the screenshot doesn't show anything or the screenshot is blank. There seems to be some error here.
Q And that was the way that the document was produced to us by Anthropic, and we can come back to that if there are questions. But --
A Okay.
Q -- please turn to page 740.
A Okay.
Q And do you see there are two messages on that page from Dario?
A That is correct.
Q And that refers to you, Mr. Amodei,

14

correct?
A I believe so, yes.
Q So you recognize this as a Slack conversation between you, Mr. Stewart, and various other Anthropic employees from February 29, 2024?
A That's right.
Q And this conversation appears to be discussing Anthropic's internal testing of CLAUDE and certain output generated by CLAUDE during that testing; is that correct?
A This looks it like from the document. It looks like that we're testing a -- we're -- a model -- bashing a model before launch, correct.
Q And what does that -- what does that mean, "bashing the model"?
A That means that we experiment with -- you know, we -- we train a model. A model can do all sorts of things.
We don't know ahead of time, you know, what all the behaviors of the models are going to be, so a bunch of people at the company get together and they interact with the model to see what it can do.
Q So you're -- you're interacting with a version of CLAUDE here before you launch that to the

15

public in order to test out the output of the model?
A In order to test the output of the model, correct.
Q And do you recall what specific issues you were testing in terms of output of CLAUDE here in this conversation?
A I -- I do not recall the specific issues particularly with -- given that these screenshots aren't visible.
Q If you turn back to page 739 --
A That's right.
Q -- do you see the second message from Jared Kaplan at the top of the page at 12:06 p.m.?
A Yes.
Q Mr. -- Dr. Kaplan writes:
> Yes, I think this is -- actually isn't all that problematic in the scheme of things. CLAUDE 2.0 will make lists of naughty words, too, without MM.
> Do you see that?

A Yes. "MM," I believe, would be multimodal. So providing images in and -- and producing words out.
Q And CLAUDE 2.0 is referring to a publicly

16

available version of CLAUDE as of February of 2024?
A That's --
ATTORNEY MEHTA: Objection --
A -- right.
ATTORNEY MEHTA: -- to form.
THE DEPONENT: Sorry.
ATTORNEY MEHTA: That's okay. You can answer.
Q (BY ATTORNEY HAILEY) And then a few messages below that you see an Anthropic employee named Diane writes:

[redacted]

Did I read that correctly?
A You read that correctly, yes.
Q Does -- do those messages clarify to you that the issues you're discussing here are output with graphic language and racial slurs?
ATTORNEY MEHTA: Objection to the form.
A No, I wouldn't -- I wouldn't say that.
I would say that the first message shows

37

A You're confusing me.
Q Mr. Amodei, you don't dispute that Anthropic has copied Plaintiffs' song lyrics from the Internet and included those lyrics in the training data to train its AI models?
A Again, you can -- you can get that information from data that was turned over and from other witnesses.
I don't, myself, look at the training data, so I don't directly know that one way or another. You can -- you can confirm that if it's true from someone else.
Q (BY ATTORNEY HAILEY) I want to turn back to Exhibit 246, which you still have in front of you, Mr. Amodei.
A Okay.
Q Do you see the Slack message from you from 12:15 p.m., on February 29, 2024?
A That's right.
Q You write:
[redacted]

38

[redacted]
A You read that correctly.
Q And it's true, as you say in this document, that Anthropic's models are unpredictable?
A All AI models are unpredictable. They're statistically-based systems.
Q Including Anthropic's CLAUDE models?
A That is correct.
Q And it's true, as you say in this Slack
[redacted]
A That is correct. I would refer back to, however, to the idea that there's a wide range of problems in models. There's very small problems.
For example, the idea that, you know, if the model said the word "hell" once, that doesn't

39

mean that we should block the launch -- the launch of the model for all the benefits that it would have all the way across the world.
The word "problems" can -- can refer to things that are extremely minor, things that on reflection are not really problems at all, or it can refer to large problems.
Q And is one of those problems outputting copyrighted material?
ATTORNEY MEHTA: Objection, asked and answered.
A Well, in this particular document -- I actually think in the context of this document, this was referring to general types of, you know, like offensive -- offensive information or information on how to do bad things.
I actually think this comment was narrowly scoped to that.
Q (BY ATTORNEY HAILEY) Let me ask you: Setting aside this specific conversation, is it a problem if Anthropic's AI models output copyrighted material?
A Again, it is -- you know, Anthropic's model should not outright copyrighted material subject to the bounds of fair use.

40

Q And what do you mean by "subject to the bounds of fair use" in the context of that answer?
ATTORNEY MEHTA: The same caution with respect to: You can share your personal views, independent of any advice of --
A Yeah --
ATTORNEY MEHTA: -- counsel.
A -- so this is -- this is not based on -- this is not based on any advice of -- this is not based on any advice of counsel.
But when we're talking about LLM outputs, you know, that is -- LLM training is obviously a very different issue. But LLM outputs is -- is just about -- you know, is just about -- is just about outputting information.
So in the normal context of copying or outputting information, there are a bunch of exceptions around the idea of fair use. And -- and you know, it's my understanding that in this context, those exceptions would -- would certainly apply.
Now, maybe there are some things that are different about LLM outputs as well and perhaps other exceptions might apply, but certainly the exceptions related to fair use would apply, as they

41

would apply in any other context.
Q (BY ATTORNEY HAILEY) So, Mr. Amodei, I'm not sure I understand your question, [sic] so I just want to clarify.
A What was not clear about my answer?
Q Give me a moment.
So do you believe, yes or no, that it's a problem if Anthropic's models output copyrighted material?
ATTORNEY MEHTA: Objection, asked and answered.
A Yeah, I'm not going to answer questions that have -- that have already been asked.
Q (BY ATTORNEY HAILEY) Are you refusing to answer the question, sir?
ATTORNEY MEHTA: Objection, asked and answered; argumentative.
A I already gave you the answer to your question.
Q (BY ATTORNEY HAILEY) And I didn't understand the question -- I didn't understand the answer, Mr. Amodei, so I'm entitled to ask a follow-up question to be sure that we understand --
A Well, you're asking the same follow-up question. It -- you know, if you want to -- if you

42

don't tell me what it is about my answer that you didn't understand, then -- then I won't be able to answer the question any differently than I did before.
Q As a factual matter, Mr. Amodei, you're aware that Anthropic's CLAUDE AI models have generated output copying copyrighted material?
ATTORNEY MEHTA: You -- to the extent you have information on that based on conversations with counsel, I direct you not to answer. If you have information on that outside of your conversations with counsel, you can answer.
A Yeah, I have -- I have no direction -- I have no information about this, other than what I have been told about this particular lawsuit.
Q (BY ATTORNEY HAILEY) You haven't -- beyond the copyrighted content at issue in this lawsuit, you're not aware that Anthropic's CLAUDE AI models have generated output copying copyrighted material?
A I --
ATTORNEY MEHTA: Same caution.
A -- I -- I don't -- yeah, I don't believe I know any details about it. Again, I run the company. I'm in charge of -- I'm in charge of the training -- the training of the models. I don't go

43

in and inspect a very specific output to models.
Q (BY ATTORNEY HAILEY) Well, you don't personally dispute that Anthropic's CLAUDE models have generated output copying the Plaintiffs' song lyrics in this case, do you, sir?
A This question was already asked to me.
Q And what was your answer?
A My -- my -- again -- again, my answer was that, you know, Anthropic data includes a lot of information. I don't -- I don't know what specifically it included. You know, everything I know about content generated that -- that belongs to plaintiff, I know through counsel in the context of this lawsuit.
Q I'm asking now, sir, about the output.
You don't personally dispute that Anthropic's CLAUDE models have generated output copying the Music Publisher Plaintiffs' song lyrics?
A I --
ATTORNEY MEHTA: Objection, asked and answered multiple times.
Counsel, you have a limited amount of time for this deposition. This is our CEO. You have asked that question at least three or four times. Let's please move.

44

ATTORNEY HAILEY: Counsel, please stop interrupting. I will raise this with the court and seek additional time on the deposition.
ATTORNEY MEHTA: He has already told you his answer to that question.
A -- I have -- I have already told -- yeah, I have -- the answer is the same for outputs as it is for training.
Q (BY ATTORNEY HAILEY) So in your 12:15 p.m. [redacted] copyright infringement as falling within those categories?
A Again, I think that this entire document was about offensive content and offensive -- offensive content and offensive words or words that could be dangerous and used to bad end.
So I think the scope of this whole comment was -- was not -- was not referring to copyrighted content one way or another. It was talking about, you know, the -- the did -- not the provenance of the content or whether you're allowed to generate the content, but the dangers of the content itself. That was the scope of his comment.

49

outputting copyrighted content, correct?
A There --
ATTORNEY MEHTA: Objection to the form.
A -- there isn't time to -- to offer every caveat or -- or -- there isn't time to offer every caveat or every piece of nuance in -- in a podcast interview.
Q (BY ATTORNEY HAILEY) And did you have that same view, that AI models shouldn't be verbatim outputting copyrighted content, when Anthropic first released CLAUDE to the public in March 2023?
A Again, that question states something to be my view that is not necessarily in -- my view and my full -- in full nuance.
My view in full nuance is that models shouldn't be outputting verbatim copyrighted content subject to the ordinary rules of fair use that -- that hold when copyright out -- content is outputted.
That -- as far as I can tell, that's been my view as long as I can remember.
Q Now, turning back to the New York Times interview, you continue at the bottom of page 15:
For things that are available

50

on the web for publicly available, our position -- and I think there's a strong case for it -- is that the training process -- again, we don't think it's just hoovering up content and spitting it out where it shouldn't be spitting it out.
Did I read that correctly?
A Yes. Followed by:
It's really much more like the process of how a human learns from experience. And so our position that that is sufficiently transformative, that this is fair use.
Again, this is a restatement -- this is a restatement and shows the consistency of the position I have that the training of models is like reading -- is like a human being reading content and using that to generate their own ideas.
Q So, Mr. Amodei, you agree that Anthropic's CLAUDE models shouldn't be spitting out copyrighted content, correct?
A That -- that question has already been asked to me, and I have already answered that question.

51

Q Mr. Amodei, that's a different question. I'm asking you: Based on what you said in this interview, you agree that Anthropic's CLAUDE model shouldn't be spitting out copyrighted content; is that correct?
A I believe actually in that statement -- in -- scoped in that specific statement, note the factual statement where I say, "We don't think it's just hoovering up content and spitting it out."
Actually, here I was making a point about the fact that if the model is not overfitting, it should not -- it should not technically do that frequently.
In other words, if the model is well trained, it should be generating content that -- different than the content that it was trained on.
I think -- I can't get back in my state of mind, because this was an interview that was almost two years ago. But I think based on this, we don't think it's just hoovering -- hoovering up content.
I think that -- I think that what I meant here was that the word "shouldn't" in this referred to technical specifications about, you know, the way in which the model is trained.

52

Q Sitting here today, sir, do you agree that Anthropic's CLAUDE models should not be spitting out copyrighted content and output?
ATTORNEY MEHTA: Objection to the form.
A That's -- that's substantially the same as the questions that I have been asked before. The -- I'll state the answer one more time for you, but I won't answer any more questions stated in substantially the same way.
My -- my belief is that subject to exceptions and exemptions around fair use, models should not be outputting copyrighted content.
Q (BY ATTORNEY HAILEY) What exceptions and exemptions surrounding fair use are you referring to specifically?
A Look, the -- the law on fair use is -- is long and voluminous. I'm not going to discuss things that have been discussed with me with counsel, but you know, my understanding is there are many -- there are many exceptions in this -- in this area. I'm not an expert in this area of law, so I couldn't give you a list.
Q Well, sir, I'm not asking you about any privilege discussions you had with counsel. But if you're volunteering in your responses and referring

57

Q Why shouldn't models be verbatim outputting copyrighted content?
A Again, I would -- I would say, again, as I have in answers to the question when I was asked it previously, that, you know, if -- if -- given the full form to expand and -- and add all nuance, as it is often not possible to do in the context of -- of a podcast interview, I would restate the statement as: I think everyone agrees the models shouldn't be verbatim outputting copyrighted content subject to the -- the exemptions to fair use that -- that apply to reproducing content in general.
Q And why is that?
A Why is that? My view -- you know, again, that is -- that is my understanding of copyright law. Again, not colored by anything that I have learned from counsel, but that's my understanding of copyright law.
Q So if copyright law didn't prohibit that, would you want CLAUDE's models to verbatim output copyrighted content?
A I mean -- you know, you're asking -- you're asking me this kind of very strange hypothetical. You know, if there have never been copyright law, like all kinds of actors in the -- in

58

the economy would behave very differently. There would be something else instead of copyright law.
It -- you know, it's very hard to answer hypothetical questions. You're going to have to be more specific.
Q Well, my question to you, sir, was why you had that view.
And your response was: It's your understanding that copyright law --
A Correct.
Q -- not colored by anything that I have learned from counsel, that's your understanding of copyright --
A That's correct.
Q -- law.
A That's correct.
Q So apart from your understanding of copyright law, is there any other reason why models shouldn't be verbatim outputting copyrighted content?
A Again, I think -- I think that's a sufficient reason.
Q I'm asking you: Are there any other reasons?
A Sure. I mean, there are lots of reasons

59

why -- why laws are made in the -- I -- I don't understand what you're asking me.
Q I'm not asking you about why the laws are made.
I'm asking you: Are there any other reasons why models shouldn't be outputting copyrighted content?
A Models shouldn't be -- models shouldn't output copyrighted content because it's -- it's against the law to do so.
Q You're aware, Mr. Amodei, that some CLAUDE users have prompted CLAUDE to output copyrighted content?
ATTORNEY MEHTA: If you have information of that independent of what you have learned through communications with counsel, you can answer.
A Again, all of my -- all of my communications about this lawsuit or similar lawsuits were -- were about -- you know, were -- were under privilege. You know, I -- you know, I don't examine or look at individual CLAUDE conversations.
Q (BY ATTORNEY HAILEY) So sitting here today as CEO of the company, you can't tell me whether some CLAUDE users have prompted CLAUDE to spit out

60

copyrighted content, apart from conversations --
A CLAUDE --
Q -- you have had with lawyers?
A -- CLAUDE -- CLAUDE users do a lot of things. It's -- it's, you know, very possible, indeed probable that across the millions of users that have used CLAUDE, you know, some users may have -- may have asked CLAUDE to output copyrighted content.
But I'm just -- I'm just being very precise here. I don't look at individual conversations myself, so strictly speaking, I don't -- I don't know of any specific user who has asked CLAUDE for information about copyrighted content.
Q Well, when you tested and developed the models internally, you -- you and Anthropic employees prompted the models to output copyrighted content; is that right?
A I think the people who were developing the safeguards against copyrighted content asked the models to output copyrighted content.
I cannot recall -- or at least I cannot recall an instance where I, myself, undertook this task, but certainly other employees at Anthropic made such tasks.

77

And there, again, I think the analogy to humans makes the most sense.

Then there's the -- the part about just outputting content. And again, my view is that all of the rules, all of the considerations -- excuse me -- all of the questions around compensation that -- that apply when a human writes content should also apply when an AI system writes content.

So, you know, again -- again, I think, you know, if -- if someone -- you know, if someone produced content, if it belongs to them, then, again, outputting it verbatim, if it isn't fair use, is -- is not something you should do, because that -- that's something that that person created.

Q And to the extent Anthropic's AI models do output copyrighted content verbatim owned by music publishers or songwriters, Anthropic doesn't pay music publishers or songwriters for that output, correct?

A Again, it's -- it's a bit of a weird form of a question. You know, it's -- it's like if you said, you know, I -- you know: I'm a company that produces a -- I'm a company that produces -- I'm a company that produces a product, and you know, it has some, like, safeguards against privacy -- or

78

piracy or stealing and someone manages to go around those safeguards and -- you know, and steal the -- steal a product, of course, we do everything we can to prevent piracy.

It's not an intent. And if we knew a specific list of things, then we could work together with the court system to make it right. And you know, we would -- we would seek to -- to -- we would seek to compensate, of course.

But this isn't an intended output of Anthropic models. This is something we're trying to prevent. This is something that happens either by accident -- although over time the -- the protections get better -- or this is something that happens because someone is adversarially trying to get around the filters and extract content.

Neither of those are things we want. And when they happen, we try to make them right.

Q But you're not disputing that CLAUDE users have, in the past, asked for lyrics to publishers' songs and that CLAUDE has generated output copying those lyrics, are you?

A I --

ATTORNEY MEHTA: Objection, form.

A -- yeah, again, you know, I'm the -- I'm

79

the -- you know, I, like, look at -- I look at high-level issues for the company. I have definitely directed folks to prevent copyrighted output from being produced and to, you know, put in place safeguards against copyrighted content.

I -- I'm not on the details. You would have to ask someone else about the extent to which copyrighted content has been generated, when it's been generated, how much adversarial activity was required to generate it.

Those are things that people on the ground look at and that I'm not familiar with the details of.

Q And again, as we talked about earlier, you don't remember the -- the dates or the details of when any of these safeguards you're referencing were actually put in place, right?

ATTORNEY MEHTA: Objection, asked and answered.

A Yeah, I mean, we had a -- we had a detailed discussion about this before. I listed the safeguards. I said some were put in place at the beginning; some were put in place later.

You gave me some dates for those safeguards and said those dates were given to you by

80

Anthropic employees. You know, I note -- you know, I assume if they -- if it's true they were given to you by Anthropic employees, then those things are true.

So those -- you know, based on -- based on what you said, those dates sound -- sound right to me.

Q (BY ATTORNEY HAILEY) But sitting here today, you don't have any independent knowledge about any of those dates or details of when --

A Yeah --

Q -- the --

A -- I don't -- I don't recall the exact date. There are -- you know, there are, like, thousands of, you know, activities and actions, hundreds of different safeguards at Anthropic -- that Anthropic puts in place.

I wouldn't remember exactly when -- you know, exactly when around three years ago a set of safeguards was put in place. Yeah, that's right.

Q Do you agree, Mr. Amodei, that Anthropic's CLAUDE AI models sometimes generate new song lyrics?

A Sorry. Could you -- could you clarify what you mean by "new song lyrics"? Like compose a new song?

81

Q Do you agree that sometimes CLAUDE users ask for new or original song lyrics or lyrics in the style of an existing song or songwriter?

A Look, I don't -- don't have detailed knowledge, but certainly something you could ask CLAUDE to do is -- you know, I could ask CLAUDE to write a poem or I could ask CLAUDE -- I could ask CLAUDE to write a poem or I could ask CLAUDE to write a -- I could ask CLAUDE to write a song.

That's certainly something that the models are capable of doing. And because it's not an existing copyrighted song, CLAUDE would allow a model to do that. CLAUDE would allow the user to do that, I would hypothesize.

Q And you agree that if a user -- a CLAUDE user used CLAUDE to write a new song, that song could compete against existing copyrighted songs?

A Well, if it were only the lyrics of a song, CLAUDE -- CLAUDE doesn't generate audio for a song. So you know, CLAUDE could write poetry that -- you know, that -- that, you know, it -- you know, someone could read instead of existing poetry. I -- you know, you can't listen to a song if CLAUDE only generates the lyrics.

Q My question was about lyrics. Do you

82

agree, sir, that -- strike that. You agree that CLAUDE users can ask CLAUDE to generate new songs that may com -- strike that again.

Mr. Amodei, you agree that CLAUDE users can ask CLAUDE to generate new song lyrics that may compete against existing copyrighted song lyrics?

ATTORNEY MEHTA: Objection to the form.

A I'm -- I'm not sure -- I'm not sure what it means for song lyrics to compete.

You know, music is listened to. You asked me these questions about whether I have ever been inspired by music. Music is inspiring because it's listened to.

I can understand what you're asking -- if you ask CLAUDE to generate a song and that song included CLAUDE lyrics, then listening to that song would -- would -- you know, you could listen to that song instead of -- instead of -- instead of what was written by -- instead of -- instead of what was written by someone else.

But -- but what exactly would you be doing? Just reading the song lyrics? Without -- without a tune to go along with it, it wouldn't be a song.

Q Sir, you understand there's a whole

83

industry of songwriters who, as their profession, write lyrics to songs?

ATTORNEY MEHTA: Objection to the form.

You can answer, if you know.

A Yeah, so I -- I don't have a sense of how the music industry -- yeah, I -- yeah, I don't actually have a sense of how the music industry works.

Q (BY ATTORNEY HAILEY) Does Anthropic compensate music publishers or songwriters at all?

A If -- so I don't know specifically. But we license a number of types of content. This includes content that is not freely available.

So for example, you know, we have licensed a number of books that were not freely available. We paid the publishers of those books for that content. If any of those books included music lyrics or song lyrics, then we would have published -- we would have compensated whoever -- whoever wrote that book, which presumably would have been -- or would have been com -- would have been compensated by, you know, whoever wrote -- wrote those music lyrics.

So I -- I -- so I know for a fact that we bought material like books. I don't know one way or

84

the other whether we bought material that includes song lyrics.

Q And the reason you paid for those books is because you couldn't find them on the Internet; is that right?

A No. I want to be clear. The reason we paid for those books is that they were not freely available on the Internet.

It's -- it's, of course, in many ways possible to obtain books that are not free -- that are not freely available on the Internet, but we choose not to do that.

We choose to buy the books and obtain them legally.

Q And you understand, Mr. Amodei, that just because something is -- can be accessed on the Internet doesn't mean it's free to be used for whatever purposes you want?

A Again, I would -- I would -- I would go back -- if we're talking about training, I would go back to this -- this same analogy with the human brain, which is that, you know, if -- if a human is legally able to read something for free, then, you know, my view, again, based on -- you know, over the last few years these legal issues are being worked

93

minute.
A That's right.
Q And just continuing here and where we have been reading, you say that:
"This is a real and important concern, especially when considered as part of an overall trend for AI progress to concentrate wealth, increase inequality, and in the long run automate labor," correct?
A Correct, yes.
Q And that was a true statement when you wrote it, correct?
A So I don't remember the details of this memo, but yeah, I -- I wrote this -- this memo.
Q You continue:
"Zooming very far out, what is happening macroeconomically is that distributed human labor is being used to train AI models by large centralized actors who concentrate the resulting profits, while in the long run making the human labor obsolete."
Did I read that correctly?

94

A Yes.
Q And that's a true statement?
A Sure.
Q And when you say "make human labor obsolete," that would include the labor of human writers, correct?
A Yeah, so I want to give a -- I want to give a broad sense here.
AI is advancing very quickly. I do have concerns macroeconomically about AI getting very good at what it's doing in -- in a way that causes it to be as good as humans at a very wide range of tasks.
But I would say certainly from the perspective of now, because it's clearer with the way the technology has developed, but I think it was even somewhat true in this case and it's compatible with this memo -- although it's not stated -- it's not stated with this memo, you know, I think this is an issue that relates to things that are much broader than just content.
I think there's a misunderstanding that the way AI models are trained is inex -- is kind of inextricably tied to content as if the models are just memorizing or spitting things out. But let me

95

give you a thought experiment.
Recently in the last couple of years AI models have gotten very good at mathematics. However, they have gotten very good at mathematics without actually training on the work of mathematicians. The reason for this is that math is about trial and error. It's about solving problems. And you can verify that your answers were correct without the need for any human involvement and without the need for any human data.
So models are and -- or if they are not already, will soon get to the point where models are better than human mathematicians at doing math.
This -- this you know, could, in fact, change the labor market for mathematicians. I guess the small labor market of mathematicians are a very, you know, narrow task. But -- but the point here is that the AI models are producing something that -- that, you know, is better than humans at what it does without having trained on data that comes from humans. And so certainly from the perspective of now -- and I can point to places where it's present -- present in the memo.
My concern is -- is not so much about content and the value of content. It's about the

96

fact that AI is -- is an artifact in the world. It's -- it's a technology. The methods for building it are very simple. I was there at the invention of many of them. But you know, they're -- they're very simple. Many people have thought of them.
And these -- these models are creating an incredible amount of productivity, but at the same time we have to think about the affect that they have on human labor in general. Again, not because they trained on a particular -- particular content, but because intrinsically the technology is capable of this.
So I guess what I would say is now -- and to some extent in the memo -- I think of this as about broader macroeconomic issues. I think the relevant issues are about things like tax policy. I think the relevant issues are about things like, you know, how we should structure our society around work.
Those are big problems. We have been alert to those problems since almost the beginning of Anthropic, and you know, I think -- I think this memo was an early attempt to grapple with those problems.
Q So, Mr. Amodei, I want to remind you: One

101

that:

"The current trend continues, and AI model training becomes an increasingly extractive concentrator of wealth and more and more industries like writing, music, video-making, learning of various human interactive skills, et cetera. Content creators grumble, and AI companies become more and more unpopular, but the trend of increasingly powerful models, increasing, automation, and increasing concentration of wealth continues."

A That's right. I'm -- I'm generally -- I'm generally concerned about -- about this set of issues, about the concentration of wealth that comes from AI.

You know, I publicly warned about the issue of job loss from AI, so I have been very transparent with the world about this.

Anthropic has an economic index that we use to keep track of these economic trends to try and document and understand these trends.

102

And I have called for -- I have even called for macroeconomic policy and taxes on AI companies in order to address these issues.

Q And you call out the music industry specifically here as -- as one of the at-risk industries, right?

A I -- I -- I'm giving several -- I'm calling -- I'm giving several examples of -- of things that AI could learn to be good at.

By the way, you know, what I had in mind here was generating audio, which is something that Anthropic doesn't do.

One reason that Anthropic doesn't generate audio and video is that, you know, we're -- we're -- you know, we're not -- generally not aiming to -- you know, we're generally not aiming to produce content in -- we're generally not aiming to produce content in creative industries.

You know, Anthropic, it's very focused on code. It's very focused on work and productivity-based activities.

So Anthropic has deliberately made -- made choices not to -- not to produce output in many creative modalities.

Q But in this internal memo from 2021, you

103

specifically called out more and more industries like writing and music.

Do you see that?

ATTORNEY MEHTA: Objection, misstates the document.

A Sorry. Could you -- could you just -- could you just ask that in a complete sentence? I didn't fully understand.

Q (BY ATTORNEY HAILEY) Sure.

You reference both writing and music in this internal memo, correct?

ATTORNEY MEHTA: Same objection.

A "In more and more industries such as." Yeah -- yeah, I give these as examples. Yes.

Q (BY ATTORNEY HAILEY) Writing and music, correct?

A That's correct, yes.

Q Now, if you go down to the next full paragraph, you write:

None of these outcomes seem particularly good for most people. An alternative that may work better for everyone is to compensate data/content producers for picture

104

their labor; not at a flat rate, but with a fraction of the profits from the model produced.

Did I read that correctly?

A That is -- that is correct.

Q So you're proposing in this Anthropic memo, one possible model to compensate data/content producers for using their content for AI training; is that right?

A No. Actually, that is not right. The way in which you stated the question makes it sound as if we are legally obligated to do so, as this -- as if this is a licensing agreement.

I did not have in mind a licensing agreement. I had in mind the idea that we would not be legally obligated to do any of this, but that we would -- we would choose to do it -- that, you know, we might consider to choose to do this because it was a good thing for society.

In other words, we did not believe that we were doing anything illegal. We believed that things can be perfectly legally, but bad for the distribution of wealth.

And so we were brainstorming possible ways to -- to address the results of a process that is

137

Cooper as part of a 60 Minutes segment about Anthropic in November 2025, correct?
A That is -- that is correct.
Q And this article, Exhibit 250, appears to be a transcript of that 60 Minutes segment, correct?
A Yeah. I mean, I -- you know, I wouldn't be able to say for sure whether it's accurate. Transcripts can have -- can have error in them. But yeah, this looks like a transcript.
Q Would you turn to page 2 for me, please.
A Page -- page 2.
Q Do you see at the top --
A That's right.
Q Do you see at the top -- towards the top of the page, Anderson Cooper says:
"You have said AI could wipe out half of all entry-level, white-collar jobs and spike unemployment to 10 percent to 20 percent in the next one to five years."
And you respond, "Yes."
A Yes.
Q And that's a true statement, that AI could wipe out half of all entry-level, white-collar jobs

138

and spike unemployment to 10 percent to 20 percent in the next one to five years?
A Well --
ATTORNEY MEHTA: Form.
A -- to be clear -- yeah, to be clear, that's a -- that's a depiction about a possible future.
So I'm not saying that that will happen. No one can know for sure that that will happen. I'm saying that there's a risk that that could happen, and I'm warning about that risk so we can decrease the likelihood that that happens.
That's my view.
Q (BY ATTORNEY HAILEY) And that's what you told Anderson Cooper in this 60 Minutes interview, correct?
A Correct.
Q And that includes affecting jobs in the music and songwriting industries, correct?
ATTORNEY MEHTA: Objection to the form.
A You know, I -- it's very hard to predict macroeconomically exactly, you know, which jobs will -- which jobs will be affected. It's very hard to predict the economy ahead of time, so I tend to think in broad terms.

139

I don't think about one particular job category. I generally think about AI as a technology that is great and wonderous in terms of, you know, being able to process information and think of new ideas and, you know, generate scientific knowledge and cures for medical information.
But precisely because it's able to do those things, it could create disruption to the job market similar to how previous technologies in the past have, but larger than previous technologies in the past have.
You know, I -- I don't think -- it's hard enough to predict the overall macroeconomic effect. I don't think I can predict which industries it will have an effect on and -- and how much it'll have an effect on those industries. That's too granular a level of prediction for me to be able to do. I'm not even sure of the broad prediction.
I'm just -- I'm just making it because I'm worried about it, and people should know that I consider it a possibility.
Q Okay. But it's certainly not your testimony that AI will not affect the music and songwriting industries, correct?

140

ATTORNEY MEHTA: Objection to the form.
A It -- you know, I don't think there's any industry where we can say AI will definitely not affect it. Although, I also don't think there's any industry where we can say we're sure AI definitely will affect it. I think we're generally uncertain.
Q (BY ATTORNEY HAILEY) Can you turn to page 8 for me, please?
A I can.
Q Do you see in the middle of the page 8 where Anderson Cooper says "that doesn't sound"?
A Let's see. I see it, yes.
Q And then if you skip down a few lines, you can see you say:
Because AI is a new technology, just like it's going to go wrong on its own, it's also going to be misused by -- you know, by criminals and malicious state actors.
Did I read that correctly?
A Oh, sorry. I'm -- I think you -- I think you lost me. Where -- could you tell me where that is?
Q Do you see where Anderson Cooper says

145

And looking back on The Cold War period, 20 to 30 percent seems like a reasonable guess for the probability of nuclear weapons causing a huge setback to civilization during that period. AGI may be more risky because it is a less stable technology. Did I read that correctly?

A That's correct.

Q And when you say "AGI" you're referring to Artificial General Intelligence, correct?

A Correct.

Q So you say in this Anthropic memo that you believe there's a 20 to 30 percent chance that AGI will lead to the destruction or permanent stunting of humanity, right?

A I -- I want to be specific here. I said that on the current trajectory. In other words, in the beginning -- in the way that companies other than us are currently behaving.

Anthropic, that probability is not a fixed amount. It's not like we're pressing a button and saying: Oh, we pressed this button, and there's a 20 -- 20 to 30 percent chance that -- that humanity

146

is destroyed.

Anthropic exists to try to change that trajectory, to try and change the development of AI in -- in a direction that reduces or eliminates the risks and increases the benefit.

The whole purpose of Anthropic's mission is to identify the risks and act in such a way that reduces those risks, to develop the technology in a different and alternative way.

Q Anthropic is developing AI models just like your competitors, OpenAI and Google and others, correct?

ATTORNEY MEHTA: Form.

A No. No. False. False. It's actually false.

Anthropic is not developing AI models that are just like its competitors. Anthropic measures its models for misalignment risk. And as of the most recent models, Anthropic models perform substantially better than other models on misalignment risk.

Anthropic is substantially more transparent, as in the Anderson Cooper interview that you described, in being honest about these -- these -- being honest about the risks of these

147

models in disclosing the risks of these models.

Anthropic has also advocated for regulation, for legislation, forcing companies to disclose the risk of these models and take safety measures against the risk of these models.

Now, that's actually a very costly action. Talking about the risk of AI models is an incredibly costly action. Because when you talk about the risks of AI models, people then blame you for this risk.

In fact, lawyers will even try to use the fact that you have talked about these risks to -- to try and catch you in "gotchas" and say: Oh, yeah, I -- I can -- I can exploit this. Because -- because you said that there was a risk, because you were honest about this risk, a lawyer can use this to try and say: Oh, you're -- you're a bad person because this risk exists.

But even though lawyers do that -- even though lawyers do that, I am still willing to talk about these risks, because it's so important to society to do it.

Q Is it your testimony under oath, Mr. Amodei, that Anthropic's CLAUDE.AI models are risk free?

148

A No. No, model can be risk free. The situation is similar to self-driving cars. There's no such thing as a risk-free self-driving car.

What there is is self-driving cars that are much safer than human drivers. Self-driving cars can be much more dangerous than human drivers. They can be equally dangerous or they can be safer.

Because the world is unpredictable -- for example, a duck could run in front of the road or there could be snow or there could be a hailstone or -- or another driver could crash into the car, creating risk that's completely exogenous and not related to the self-driving car.

So in the real messy world, there's no such thing as being risk free. It's impossible to be risk free. What we aim for is -- is to reduce risk over and over again to add more nines to the level of reliability.

There's no such thing as a risk-free world. Anthropic aims to reduce risk as much as possible. That is all that is physically possible. It is not physically possible to reduce risk to zero.

Q Turn to page 587, please, Mr. Amodei. Do you see the paragraph beginning "in closing"?

157

I, MARY J. GOFF, CSR No. 13427, Certified Shorthand Reporter of the State of California, certify;

That the foregoing proceedings were taken before me at the time and place herein set forth, at which time the witness declared under penalty of perjury; that the testimony of the witness and all objections made at the time of the examination were recorded stenographically by me and were thereafter transcribed under my direction and supervision; that the foregoing is a full, true, and correct transcript of my shorthand notes so taken and of the testimony so given;

That before completion of the deposition, review of the transcript ( ) was (XX) was not requested: ( ) that the witness has failed or refused to approve the transcript.

I further certify that I am not financially interested in the action, and I am not a relative or employee of any attorney of the parties, nor of any of the parties.

I declare under penalty of perjury under the laws of California that the foregoing is true and correct, dated this 9th day of January, 2026.

Mary J. Goff

MARY J. GOFF

158

ERRATA SHEET
PLANET DEPOS
451 Hungerford Drive, Suite 400
Rockville, Maryland 20850
888-433-3767

CASE: Concord v. Anthropic

PAGE	LINE	FROM	TO

DARIO AMODEI, Ph.D.

Subscribed and sworn to before me this _____ day of _____, 2026.

Notary Public