UNITED STATES DISTRICT COURT

NORTHERN DISTRICT OF CALIFORNIA

FEDERAL TRADE COMMISSION, et al., Plaintiffs, v. UBER TECHNOLOGIES, INC., et al., Defendants.	Case No. 25-cv-03477-JST (TSH) **DISCOVERY ORDER** Re: Dkt. No. 243

The Court held a hearing today concerning the parties' discovery status report at ECF No. 243. This order follows.

A. Request for Production 11

The FTC's RFP 11 sought "All video recordings, Including compilations of video recordings, of consumers or individuals interacting with the Uber One Enrollment or Cancellation Process." The FTC requests a July 6, 2026 deadline for Uber to complete its production in response to this RFP, and Uber agrees to that date. Accordingly, the Court **ORDERS** Uber to complete its production in response to RFP 11 by July 6.

B. Requests for Production 33(c) and (d)

On March 11, 2026, the Court ordered Uber to complete its non-custodial production in response to RFPs 33(c) and (d) within 45 days. ECF No. 150. Because the 45th day was a Saturday, that means Uber's deadline was April 27, 2026. In the status report, the FTC requested that the Court order Uber to confirm whether it produced all non-custodial documents in response to RFPs 33(c) and (d), and Uber confirmed that it did by the April 27 deadline. There was further discussion of this matter at the hearing. At this time, there is no dispute for the Court to resolve.

C. Privilege Logs

Uber proposes a July 24, 2026 deadline for Uber's first privilege log (expected to include approximately half of the documents) and a July 31, 2026 deadline for Uber's final log (to include the remaining privileged custodial documents). The FTC asks the Court to enter Uber's proposal. Accordingly, the Court **SO ORDERS**.

D. Requests for Production 8 and 10 (Enrollment and Cancellation Flows from Foreign Jurisdictions)

There was discussion at the hearing about the fact that Uber did not produce the entry points for the five foreign jurisdictions. Uber argued that it considers the entry points to be different from the enrollment flows. From a business perspective, that could be true. However, the FTC defined "enrollment" and "enrollment process" in its RFPs to include "all immediately preceding or immediately subsequent steps that form part of that process, flow, or experience (such as any elements that could influence the customer's behavior during, or understanding of, the process, flow, or experience)." The Court thinks that definition includes the entry points. Accordingly, the Court **ORDERS** Uber to produce the entry points for the five foreign jurisdictions. The Court **ORDERS** Uber to complete that document production within three weeks.

E. Uber's TAR Protocol

The FTC says that Uber made its first custodial document production of about 18,000 documents and represented that its production is 25% complete. The FTC extrapolates from this that Uber will likely produce around 72,000 custodial documents. The FTC thinks "[t]his is a shockingly low figure," which means there is a problem with Uber's TAR protocol. Uber says the FTC's math is wrong because its 25% estimate excluded Slack materials, and that the more likely end result is a production of around 156,000 documents.

The FTC has two requests. First, it asks the Court to change the recall rate for Uber's TAR protocol from 75% to 85%. Second, it asks the Court to order Uber to produce a random sample of the training documents marked non-responsive.

As to the first request, the Court has given Uber a July 13, 2026 deadline to complete its

custodial document production. ECF No. 233. Changing the recall rate 11 days before the production deadline is not practical or feasible. ECF No. 243-1 (Supplemental Declaration of Jeff Grobart ¶¶ 8-10).

The FTC's second request, however, has merit. Sometimes when parties use TAR they exchange training sets, so each side can see what the other is calling responsive. And then when they've completed production, they exchange validation information to show they've met the agreed upon metrics. When that happens, the parties can then bring to the Court any disputes about responsiveness. Here, that kind of transparency does not exist. Uber's document reviewers train the TAR model with their responsiveness calls, but the FTC does not see what those calls are. Allowing the FTC to review a random sample of the training documents that Uber's document reviewers marked as non-responsive would let the FTC see if the model is being trained improperly, and if it is, the Court could order Uber to retrain it.

Both sides cite *Winfield v. City of New York*, 2017 WL 5664852 (S.D.N.Y. Nov. 27, 2017). The FTC cites it in support of its request for a random sample of non-responsive training documents so it can determine if the TAR model has been appropriately trained. Uber cites the same case for the proposition that courts have held that documents used to train a TAR model are work product.

Winfield similarly dealt with an objection to how the defendant was coding documents as responsive or non-responsive for purposes of training its TAR model. In that case, the Court found "that Plaintiffs have presented sufficient evidence to justify their request for sample sets of non-privileged documents from the documents pulled from the 50 custodians." It then ordered the defendant "to provide to Plaintiffs a sample of 300 non-privileged documents in total from the HPD custodians and the Mayor's Office. These documents should be randomly pulled from the corpus of non-responsive documents." *Id*. The Court further ordered the defendant to "provide Plaintiffs with a random sample of 100 non-privileged, non-responsive documents in total from the DCP/Banks review population." *Id*. Thus, the case cited by Uber ordered the relief the FTC seeks here. *Winfield* does not stand for the proposition that the training documents coded as non-responsive are work product, as it ordered random samples of them produced.

Also, think about Uber's work product argument for a moment. If the non-responsive documents are work product because producing them would reveal counsel's thought processes, then the responsive documents would also be work product for the same reason. Under Uber's reasoning, every document review should result in no documents being produced. That doesn't make any sense.

Further, aside from its work product argument, Uber argues that it is not standard practice to exchange documents used to train a TAR model. But in reality, "where the parties do not agree to transparency, the decisions are split and the debate in the discovery literature is robust." *Rio Tinto PLC v. Vale, S.A.*, 306 F.R.D. 125, 128 (S.D.N.Y. March 2, 2015); *see also Winfield*, 2017 WL 5664852 at *10 ("Courts are split as to the degree of transparency required by the producing party as to its predictive coding process.").

Here, the parties disagree on whether Uber's 10% responsiveness rate means the search terms are broad (Uber's view) or Uber's responsiveness calls are too narrow (FTC's view). There is no real way to answer that question without data. Accordingly, it is appropriate to require Uber to produce a random sample of 300 documents that were marked as non-responsive to allow the FTC to test Uber's responsiveness calls.

At the hearing, there was discussion about the denominator from which the 300 documents should be drawn. Uber suggested it should be the output of the TAR model rather than the coding decisions that were the input. Also, if it is the training documents, there needs to be some specification of which training documents, as continuous active learning means that every day there are more training documents. The FTC requested that the random sample be taken from the initial seed set that was first used to train the model. The FTC explained that it understands that there were later coding decisions made to further train the model, but it wants to see the initial responsiveness calls to see if there were problems right from the beginning. That sounds like a good idea. The FTC has doubted Uber's 10% responsiveness rate from the beginning, and that rate has held fairly constant, suggesting that the human reviewers have likely been consistent in their responsiveness calls. An additional benefit to sampling the initial seed set is speed. If the sample were drawn from the final output of the model, or the final set of training documents, it

could not be drawn until after custodial document production is completed by July 13. By contrast, Uber represents it can produce a random sample drawn from the initial training set by July 10.

Accordingly, the Court **ORDERS** Uber to produce a random sample of 300 non-privileged, non-responsive documents from the initial seed set no later than July 10, 2026.

F. Future Discovery Status Reports

For future discovery status reports, the Court **ORDERS** the moving party to send its portion (including any declarations and exhibits) by 9:00 a.m. Pacific time three days before the deadline, the responding party to send its responses (including any declarations and exhibits) by 5:00 p.m. the day before the deadline, and the parties to file the statement by noon on the day of the deadline.

G. Interrogatories 9, 10 and 11

There does not seem to be a dispute for the Court to rule on concerning Uber's interrogatories 9-11 to the States because the States have agreed to amend their responses by around July 6, 2026.

H. Interrogatory 12

Uber and the States have a dispute about interrogatory 12. They cite Alabama's response as being representative of the States' responses. Rog 12 asked: "State whether You contend that Uber did not disclose the date on which the following month or year's Uber One subscription would be charged. If You so contend, Identify all facts in support of your response. If You do not so contend, so state." Following some objections, "Alabama answers that its allegations do not revolve around Uber's failure to disclose a billing date. Rather, Alabama alleges that Uber charged consumers up to forty-eight hours before the stated billing date, thereby charging consumers without their express informed consent and causing them to pay additional fees despite the consumers having cancelled their subscriptions before the stated billing date."

The Court agrees with Uber in part and with the States in part. The Court gathers that the true answer to rog 12 is "yes, we contend that," but the States are concerned that answer could be misinterpreted and then seemingly disproved by Uber showing the trier of fact many examples

where it told consumers what the billing date is. Accordingly, the Court thinks the States mean to say something like: "We do not contend that you did not disclose *a* date on which the following month or year's Uber One subscription would be charged. However, you actually charged consumers up to 48 hours *before* the date you 'disclosed.'"

The problem with the States' response as currently drafted is that sentence #1 that says what the States don't contend uses elliptical and indirect language ("its allegations do not revolve around"), while sentence #2 is clear on what the States do contend. Uber is entitled to a more direct sentence #1.

However, the Court does not agree with Uber that "stated billing date" is vague. An ordinary reader of the English language would understand this to mean the date that Uber stated it would charge the following month or year's Uber One subscription. Uber's rog does not draw a distinction between the stated date and the actual date, referring generally to the disclosure of "the date" on which the following month or year's Uber One subscription would be charged. But the States are free to say that there are two dates – the stated one, and the real one.

Accordingly, Uber's motion to compel as to rog 12 is **GRANTED IN PART** and **DENIED IN PART**.

I. Next Discovery Hearing

The Court **SCHEDULES** a further discovery status conference for July 24, 2026, at 10:00 a.m. The parties shall file a discovery status report no later than July 17, 2026.

IT IS SO ORDERED.

Dated: July 2, 2026

THOMAS S. HIXSON
United States Magistrate Judge