

Exhibit N

**UNITED STATES DISTRICT COURT
NORTHERN DISTRICT OF CALIFORNIA
OAKLAND DIVISION**

IN RE: SOCIAL MEDIA ADOLESCENT
ADDICTION/PERSONAL INJURY
PRODUCTS LIABILITY LITIGATION

THIS DOCUMENT RELATES TO:

*People of the State of California, et al. v.
Meta Platforms, Inc., et al.*

MDL No. 3047

Case Nos. 4:22-md-03047-YGR-PHK
4:23-cv-05448-YGR

Judge Yvonne Gonzalez Rogers

EXPERT REBUTTAL REPORT OF NICK FEAMSTER, PH.D.

April 15, 2026

Highly Confidential

I. SUMMARY OF OPINIONS

1. I have been retained by Meta Platforms, Inc. and affiliated defendants (“Meta”), through its counsel, to respond to opinions offered by plaintiffs’ experts Dr. Ryan Sheatsley and Mr. Tim Estes. A summary of my opinions follows below.

A. Opinion 1: Meta’s Defense-in-Depth Approach to Underage Detection Meets or Exceeds Security and Industry Standards

2. Meta’s layered approach to detecting and removing individuals under the age of 13 (“U13s”) meets or exceeds security and industry standards, and neither Dr. Sheatsley nor Mr. Estes shows otherwise. Meta’s systems employ a layered, defense-in-depth approach to underage detection that includes: collection of age data at account signup; automatic checkpoints when users change their age to indicate they are under 13 or over 18; user reporting mechanisms (including both the standard form and SURF); proactive detection via algorithmic signals; human review workflows with instructions to checkpoint when in doubt; hard-linking through Account Center connections; and appeals processes. Neither Dr. Sheatsley nor Mr. Estes acknowledges the full scope of this defense-in-depth approach. This is reasonable system design for platforms of Meta’s scale.

3. Neither Dr. Sheatsley nor Mr. Estes compares Meta’s practices to other social media platforms or offers any opinion that Meta’s practices are deficient compared to other social media platforms. Other major platforms (YouTube, TikTok, Snapchat, Twitter/X) also relied on self-reported age during the relevant period. The age assurance technologies that plaintiffs’ experts cite were not widely deployed by any platform during most of the time period at issue here. Meta was an early adopter of face-based age prediction technologies like Yoti compared to competitors.

B. Opinion 2: Plaintiffs’ Experts Fail to Consider Real-World Factors That Impact the Design and Operation of Age Assurance Systems

4. Both Dr. Sheatsley and Mr. Estes fail to consider critical real-world factors that impact age assurance systems. Neither expert analyzes false positive rates or the precision-recall tradeoffs inherent in any classification system. In none of Dr. Sheatsley’s eight opinions has he done any analysis to determine how many of those users are actually under 13. Checkpointing does not equal confirmation of under-13 status, and even removal does not prove a user was under 13. Neither expert acknowledges that security systems are inherently reactive, operating in an adversarial environment where proactive detection has fundamental limits.

5. Plaintiffs’ experts also fail to consider the real costs of their proposed alternatives beyond monetary implementation costs, including accessibility impacts that would exclude legitimate users, false positive operational burdens that would overwhelm review systems, and privacy tradeoffs. The population of eligible users (anyone 13+) is

Highly Confidential

vastly larger than the population of ineligible users (the subset of U13s who have access to devices and attempt to create accounts), which impacts what would be reasonable and feasible in designing a defense-in-depth approach, and neither Dr. Sheatsley nor Mr. Estes considers that context.

C. Opinion 3: Dr. Sheatsley's Analysis of Meta's Checkpointing Data Is Consistent with a Reasonable and Layered Approach to U13 Detection

6. Dr. Sheatsley's analysis of Meta's checkpointing data demonstrates that Meta successfully identified and disabled a significant volume of potential U13 accounts. He reports that Meta checkpointed and disabled more than one million U.S. users during the time period covered by Meta's produced data, notwithstanding uncertainty as to whether those users were actually U13—for example, where a user fails to respond to a checkpoint, where a user's ID does not match their username, or where accounts are hard-linked in the Account Center (notwithstanding that they may belong to different users), such users' accounts are nevertheless disabled.

7. Dr. Sheatsley's opinion that Meta's use of the SURF form would increase the number of detected U13 accounts is misleading because it focuses on checkpointing rate rather than the ultimate action rate. Meta's own data shows that SURF generated low-quality reports with a 37% false positive rate and had a 25% higher false positive rate and 58% lower successful action rate compared to other reporting mechanisms. The SURF form also generated more noise, not better signal. In any event, some friction in the reporting process serves an important function: it helps discourage abusive and false reports submitted for harassment or other malicious purposes. Sheatsley's opinion that SURF and contact form have comparable checkpointing rates is also not supported by his analysis, which shows lower checkpointing rates for reports submitted through SURF than those submitted via the contact form.

8. Dr. Sheatsley's opinions regarding Meta's use of thresholds and automations fail to recognize that threshold selection is standard machine learning practice, not a design flaw. He provides no alternative threshold and does not analyze the false positive consequences of lower thresholds. With respect to the Adult Classifier Model, he overlooks its actual function in Meta's layered approach: by identifying users over the age of 18, the model preserves and prioritizes finite human resources for reports of potential U13 accounts that would most benefit from human review.

9. Dr. Sheatsley's analysis of review times is misleading. While he opines that many reports are resolved quickly, the average review time is approximately 72 seconds—significantly higher than the median—with substantial long-tail cases requiring extended review time. Thus, while some reports may be resolved quickly, accounts with more complex elements may require more time to assess.

10. Dr. Sheatsley reports that 44,289 checkpoints (3%) took over one year, but he fails to engage with the reasons for this duration. Appeals processes require time

Highly Confidential

(including for users to submit appeals and for appeals to be resolved), distributed systems have technical constraints, and legal obligations for preservation exist, all of which may impact the amount of time before an account is disabled. Moreover, accounts are inaccessible to the user while they remain in the checkpoint. He offers no comparison to any other company achieving better or faster results.

11. Finally, Dr. Sheatsley's opinion that Meta could have removed "many" underage users through hard-linking and soft-matching is unsubstantiated. Soft-matching produces probability scores, not definitive identifications. Automated removal based on probabilistic linkage would necessarily result in many false positives, disabling accounts for users over the age of 13, including siblings or family members accessing accounts from the same IP address or sharing devices. I am not aware of any other social media company that uses soft-matching to automatically disable potential U13 accounts, and Dr. Sheatsley cites no authority demonstrating that other platforms employ this approach.

D. Opinion 4: Mr. Estes's Proposed Alternatives Are Impractical, Would Create Significant Accessibility Barriers, and Are Not Required by Any Applicable Industry Standard

12. Mr. Estes compares Meta to other industries (not other social media companies), implicitly acknowledging that Meta's practices are consistent with what every other social media company does for age assurance. His proposed alternative approaches to age assurance—credit card verification, government ID verification, federated identity, mobile carrier data, and AI-based facial age estimation—would each create significant accessibility barriers, privacy concerns, and operational challenges that he fails to address, and as a result, exclude legitimate users.

13. Mr. Estes's cost-benefit analysis only discusses implementation costs ("dollars and cents") while ignoring accessibility costs, false positive operational costs, privacy tradeoffs, and policy implications of making a public forum inaccessible. His comparison to gaming platforms is inapposite because gaming platforms allow individuals under 13 and use Verifiable Parental Consent (VPC) because they have U13 users, while Facebook and Instagram prohibit U13s under a different regulatory model. Gaming platforms are also built around charging customers and need to collect payment information regardless of age assurance, which is not the case for free platforms like Facebook and Instagram.

II. INTRODUCTION & QUALIFICATIONS

A. Assignment

14. I have been retained by Meta Platforms, Inc., through its outside counsel, to provide expert opinions in this matter—specifically, to review and respond to the March 20, 2026 Trial Report of Dr. Ryan Sheatsley and the March 20, 2026 Trial Report of Mr. Tim Estes. In consideration for my work in this case, I am being compensated in the amount of

Highly Confidential

F. False Positives Impose Substantial Harms

64. Neither Dr. Sheatsley nor Mr. Estes considers the costs of false positives (users over the age of 13 that are incorrectly flagged or removed). Any age assurance system must balance these competing error types, which plaintiffs' experts fail to address at all.

65. False positives impose substantial harm on legitimate users. Legitimate 13+ users incorrectly flagged and removed from platforms lose access to their data, social connections, online communities, educational resources, and in some cases business income (e.g., influencers or small business operators).

66. Mr. Estes discusses face-based age estimation as a solution, but NIST research on facial age estimation shows that accuracy varies significantly by demographic characteristics.¹³ The 2024 NIST Face Analysis Technology Evaluation found that "error rates were almost always higher for female faces than for males," and that accuracy is "strongly influenced by algorithm, sex, image quality, region-of-birth, age itself, and interactions between those factors."¹⁴ NIST concluded that "an algorithm that performs well on certain groups can perform poorly on others," creating discrimination risks if such systems were deployed universally. Single accuracy figures cited from marketing materials are misleading because accuracy varies significantly based on sex, race, lighting conditions, and image quality.

67. Under laws like Illinois's Biometric Information Privacy Act (BIPA), even transient collection of biometric identifiers requires informed consent and creates legal obligations—the collection itself, not just retention, is regulated.¹⁵

collecting personal information from children). Section 230 of the Communications Decency Act, 47 U.S.C. § 230, available at <https://www.law.cornell.edu/uscode/text/47/230>, provides immunity for user-generated content but imposes no affirmative age verification obligations.

¹³ Estes Report ¶¶ 96-101.

¹⁴ NIST IR 8525, Face Analysis Technology Evaluation: Age Estimation and Verification (May 2024), available at <https://nvlpubs.nist.gov/nistpubs/ir/2024/NIST.IR.8525.pdf>; see also <https://www.nist.gov/news-events/news/2024/05/nist-reports-first-results-age-estimation-software-evaluation>.

¹⁵ 740 ILCS 14/15(b), available at <https://www.ilga.gov/legislation/ilcs/ilcs3.asp?ActID=3004&ChapterID=57> ("No private entity may collect, capture, purchase, receive through trade, or otherwise obtain a person's or a customer's biometric identifier or biometric information unless it first informs the subject or the subject's legally authorized representative in writing that a biometric identifier or biometric information is being collected or stored[.] . . . of the

Highly Confidential

68. At Meta's scale of millions of daily signups, biometric data would continuously flow through collection and processing systems, creating ongoing data breach exposure even with rapid deletion.

69. The friction and exclusion problems remain: users without camera-equipped devices, users in low-bandwidth environments, and users uncomfortable providing facial scans would still be excluded from account creation.

70. Timely deletion does not address the underlying accuracy limitations of age estimation technology, particularly for younger users and certain demographic groups. For instance, Meta's ability to develop and deploy machine learning tools to detect U13 individuals is constrained by restrictions on collecting and retaining data from U13 individuals, including under COPPA. This is a challenge for all companies that age-gate U13 individuals.

71. In addition, anonymous speech protections serve important social functions. Requiring real identity verification eliminates anonymity for LGBTQ youth in hostile environments, political dissidents, victims of domestic violence, whistleblowers, and others with legitimate needs for pseudonymous participation. These vulnerable populations benefit from social connection while facing real-world dangers from identification.

G. The Precision-Recall Tradeoff Bounds System Design

72. Every binary classification system faces the precision-recall tradeoff: increasing true positive rate (recall) generally decreases positive predictive value (precision) unless the underlying classifier improves. This is a mathematical relationship, not a design choice.

73. The optimal operating point on the precision-recall curve depends on the relative costs of false positives and false negatives. A system optimized purely to minimize false negatives would flag the vast majority of users for review, creating an unusable platform. A system optimized purely to minimize false positives would rarely flag anyone, potentially allowing widespread underage access. As detailed further below, Meta's approach falls between these extremes and is reasonably tailored to identify and remove U13s, consistent with applicable security standards and rules.

specific purpose and length of term for which a biometric identifier or biometric information is being collected, stored, and used; and receives a written release executed by the subject . . . or the subject's legally authorized representative.").

Highly Confidential

IV. OPINION 1: Meta’s Defense-in-Depth Approach to Underage Detection Meets or Exceeds Security and Industry Standards

74. Meta’s layered approach to detecting and removing U13s meets or exceeds security and industry standards. This approach consists of multiple components that work together to identify and remove underage users from Meta’s platforms. Neither Mr. Estes nor Dr. Sheatsley acknowledges or considers the full scope of Meta’s defense-in-depth approach.

1.1 Meta’s Systems Employ a Layered, Defense-in-Depth Approach

75. Meta’s systems employ a layered, defense-in-depth approach to underage detection that includes: the collection of user age at account signup; automatic checkpoints when users change their age to indicate they are under 13 or over 18; user reporting mechanisms that are available not only to users of the platforms, but to all individuals via a web form; proactive detection via algorithmic signals; human review workflows, with instructions to checkpoint when in doubt; hard-linking (Account Center connections); and appeals processes.

76. Both Facebook and Instagram require individuals to report their age at account registration.¹⁶ Users must enter their date of birth when creating an account. Users who enter a birth date indicating they are under 13 cannot access the platforms, consistent with the terms of service for Instagram and Facebook.¹⁷ The birth date entry screens are neutral—for example, they do not prohibit users from entering birth dates that would render them ineligible to access the platforms, and they do not indicate the

¹⁶ I understand that Instagram has required individuals to report their age at account registration since December 2019, and that in May 2022, Meta “rolled out a process that required all current users on Instagram who did not have an age associated with their account to provide their date of birth.” META3047MDL-012-00000001 at -002. I understand that Facebook has required individuals to report their age at account registration since at least 2009. See META3047MDL-208-00003312.

¹⁷ META3047MDL-091-00080707 (Facebook Statement of Rights and Responsibilities, Nov. 15, 2013, Section 4.5: “You will not use Facebook if you are under 13”); META3047MDL-208-00003308 (Instagram Terms of Use, Jan. 2013, Basic Terms Section 1: “You must be 13 years or older to use the Service”); Meta Defendants’ Responses and Objections to Plaintiffs’ Third Set of Requests for Admission, No. 18 (“Since at least 2012, Facebook’s Terms of Service and Instagram’s Terms of Use have prohibited individuals under the age of 13 from creating a Facebook or Instagram account.”).

Highly Confidential

VI. OPINION 3: Dr. Sheatsley’s Analysis of Meta’s Checkpointing Data Is Consistent with a Reasonable and Layered Approach to U13 Detection

3.1 Meta Successfully Identified and Disabled a Significant Volume of Potential U13 Accounts (Sheatsley Opinion 1)

102. Dr. Sheatsley’s analysis shows that Meta checkpointed and disabled more than one million U.S. accounts during the time period covered by Meta’s productions. This demonstrates that Meta’s layered detection system identifies and removes a significant volume of potential U13 accounts. Meta disabled those accounts notwithstanding uncertainty as to whether the accounts belonged to U13s. For example, accounts are disabled where a user fails to respond to a checkpoint (which does not prove they were actually under 13), where a user’s ID does not match their username (which could occur for legitimate reasons even where the user is over 13), or where accounts are hard-linked in the Account Center (notwithstanding that they may belong to different users).

3.2 SURF Generated Low-Quality Reports with High False Positive Rates (Sheatsley Opinion 2)

103. Dr. Sheatsley opines that expanding Meta’s Shortened Underage Reporting Flow (“SURF”) would increase the number of detected underage users because SURF has “comparable checkpointing rates” to the standard contact form while generating higher reporting volume.³⁸ This opinion is misleading and unsupported for two reasons.

104. First, Dr. Sheatsley’s analysis of SURF checkpointing rates contains a significant methodological flaw. His primary tables purport to show SURF reports across the entire data period, but SURF was not launched until March 22, 2022.³⁹ It is impossible for there to be SURF reports before SURF existed. Dr. Sheatsley offers no explanation for why his primary analysis includes data from this impossible pre-launch period, or why such an analysis would be reliable. Dr. Sheatsley’s primary analysis also doesn’t apply the “Instagram” value to the “app_name” filter.⁴⁰ Given SURF was launched on only Instagram, Dr. Sheatsley’s failure to apply the “app_name” filter may have resulted in the inclusion of Facebook contact forms that may not be comparable to Instagram contact forms and SURFs.

105. When Dr. Sheatsley applies date and platform filters to exclude the pre-launch period and potential Facebook reports, the results undermine his own conclusion. With proper date and platform filtering, SURF checkpointing rates are significantly lower than contact form rates—not “comparable” as he claims. For example, with the filters,

³⁸ Sheatsley Report ¶ 2.

³⁹ Sheatsley Report ¶¶ 87–93.

⁴⁰ Sheatsley Report ¶¶ 87–93

Highly Confidential

contact forms have a pooled checkpointing rate (across reports and form Table 1, Table 1 (Potential US), and Table 2) of approximately 14.6% compared to approximately 4.2% for SURF—more than triple.⁴¹ The only logically valid version of his analysis—the one that excludes the impossible pre-launch data—shows that SURF performed worse than the standard contact form, not comparably.

106. Indeed, Meta’s own data shows that SURF generated low-quality reports with a 37% false positive rate and had a 25% higher false positive rate and 58% lower successful action rate compared to other reporting mechanisms.⁴² SURF generated more noise, not better signal.

107. Second, the fact that an account is checkpointed does not mean that the user is U13. Indeed, many checkpointed accounts clear the checkpoint and are reinstated. Comparing the rate at which accounts reported through the SURF are *disabled* (i.e., fail the checkpoint) to the rate at which accounts reported through the contact form are disabled is the proper analysis to determine whether the SURF form generated more noise than signal. Dr. Sheatsley did not analyze this issue in his report.

108. To the extent that Dr. Sheatsley implies that Meta should have adopted and deployed the SURF form more broadly, I disagree with that conclusion. Some friction in the reporting process serves an important function: it helps discourage abusive and false reports.

109. Evidence shows that a significant number of inauthentic reports are submitted as a form of harassment against high-profile individuals and institutions (e.g., Kim Kardashian and the IDF).⁴³ This problem is well-documented across other platform contexts: research on DMCA takedown notices found surprisingly high percentages of

⁴¹ See my Appendix C-1, where I calculate pooled checkpointing rates based on the numbers Dr. Sheatsley separately reported for Table 1, Table 1 (Potential US), and Table 2 reports. In the same appendix, I also plot the checkpointing rates for SURF and contact form reports over time based on my replication of Dr. Sheatsley’s analysis with the date and platform filters, which shows consistently lower checkpointing rates for SURF compared to contact form reports.

⁴² See META3047MDL-111-00080128 (SURF generated “low quality reports” with 37% false positive rate); META3047MDL-190-00001439 (SURF has 25% higher false positive rate and 58% lower successful action rate).

⁴³ META3047MDL-208-00003249 (internal Meta analysis titled “Exploratory Analysis of Reactive Reporting,” documenting abuse of the U13 reporting system through false harassment reports filed against high-profile accounts including celebrities and public figures).

Highly Confidential

notices of questionable validity generated by automated systems,⁴⁴ and automated enforcement systems have repeatedly generated false reports targeting legitimate content, including incidents where automated systems falsely identified popular software like Skype and Tor as infringing content, or where HBO's own official website was targeted for removal.⁴⁵ If there were no friction whatsoever in the reporting process, platforms would face an overwhelming volume of false reports, creating additional burdens on the review system and making it harder to address actual U13 cases. The reporting form strikes a reasonable balance between accessibility for legitimate reporters and deterrence of abusive reporting.

3.3 Threshold Selection and Automation Are Standard Practice (Sheatsley Opinions 3 and 4)

110. First, Dr. Sheatsley criticizes Meta's automations that "ignore" certain underage reports. He reports that approximately 40,000 reports were ignored because the reported account did not have a bio or photos (the information used by human reviewers to make age determinations), and approximately 13,000 reports were ignored because less than 50% of photos on the user's profile were of children.⁴⁶

⁴⁴ Jennifer M. Urban, Joe Karaganis & Brianna L. Schofield, *Notice and Takedown in Everyday Practice*, UC Berkeley Public Law Research Paper (Mar. 2017), available at <https://ssrn.com/abstract=2755628> (examining over 108 million DMCA requests sent to Google during a six-month period and finding surprisingly high percentages of notices of questionable validity generated by automated systems).

⁴⁵ See Electronic Frontier Foundation, *Absurd Automated Notices Illustrate Abuse of DMCA Takedown Process* (Feb. 12, 2015), available at <https://www.eff.org/deeplinks/2015/02/absurd-automated-notices-illustrate-abuse-dmca-takedown-process> (documenting how Total Wipes' automated system sent Google takedown demands targeting popular software including Skype, Tor, and WhatsApp, attributing errors to "technical servers problems" with its automated system); *System Used By New Six Strikes CAS, Falsely Identifies Game Mods As NBC TV Shows*, Techdirt (Mar. 1, 2013), available at <https://www.techdirt.com/2013/03/01/system-used-new-six-strikes-cas-falsely-identifies-game-mods-as-nbc-tv-shows/> (MarkMonitor's DtecNet division, on behalf of NBC Universal, falsely identified a Guild Wars game modification file as infringing NBC television shows); *How Much Does HBO Pay MarkMonitor To Send DMCA Notices Removing Its Official Content From Google?*, Techdirt (Feb. 5, 2013), available at <https://www.techdirt.com/articles/20130205/03124421884/how-much-does-hbo-pay-markmonitor-to-send-dmca-notices-removing-its-official-content-google.shtml> (MarkMonitor sent DMCA notice demanding removal of links to HBO's own official website for *Eastbound and Down*).

⁴⁶ None of these reports were actually ignored; to the contrary, they were resolved through Meta's multi-layered processes, which include the use of automations to resolve reports,