EXHIBIT 15

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.

January 9, 2026

I. SUMMARY OF OPINIONS

1. I have been retained by Meta Platforms, Inc. and affiliated defendants ("Meta"), through its outside counsel, to respond to opinions offered by plaintiffs' expert Dr. Patrick McDaniel. A summary of my opinions follows below.

A. Opinion 1: Meta's Processes for Detecting U13 Individuals Meet or Exceed Security and Industry Standards

2. Meta's layered approach to detecting and removing individuals under the age of 13 ("U13 individuals" or "U13s") meets or exceeds security and industry standards, and Dr. McDaniel fails to show otherwise. Dr. McDaniel's critiques misunderstand the technical and operational constraints of deploying detection systems at platform scale, which requires the use of reasonably tailored filters or other forms of automation. His analysis ignores base-rate effects on classifier accuracy (the mathematical reality that rare events are hard to detect accurately, even with good systems), treats standard machine learning threshold selection—the necessary process of deciding how confident a system must be before flagging an account—as a design flaw, fails to recognize that all security systems are inherently reactive, and critically proposes no viable alternative approach. Fundamentally, Dr. McDaniel does not—and cannot—identify any specific accounts that Meta knew were operated by U13 individuals and failed to disable.

3. Dr. McDaniel cites general security principles but provides no specific industry standard for age assurance on social media platforms at scale. None of the authorities Dr. McDaniel cites required proactive age verification at account creation during the time period in question.

4. Indeed, other major platforms (YouTube, TikTok, Snapchat, Twitter/X) also relied on self-reported age during the relevant period. The age-assurance technologies that Dr. McDaniel cites were not widely deployed by any platform during most of the time period at issue here. In fact, compared to competitors, Meta was an early adopter of face-based age prediction technologies like Yoti.

B. Opinion 2: Meta's Human Review Processes Meet or Exceed Industry Standards

5. Meta's human review process—one of several layers in Meta's U13 detection processes—meets or exceeds security and industry standards. Dr. McDaniel—who simultaneously critiques Meta's failure to use more human review for a greater proportion of cases, the guidance provided to human reviewers, and the time required to complete human review—ignores the infeasibility of comprehensive human review at Meta's scale. Meta operates billions of accounts and has millions of new signups daily. Structured workflows and rapid review of clear cases are necessary features, not flaws. At Meta's scale, efficiency is a form of accuracy—slow review means delayed enforcement. Meta's

human review process is reasonably tailored to balance accuracy with operational efficiency.

6. Dr. McDaniel cites no industry alternative achieving better outcomes. Unstructured discretion, as he suggests, would produce wildly inconsistent outcomes across thousands of reviewers in multiple countries and languages. His own data shows 80% of checkpoints result from users who self-admit their age, which should resolve quickly.

C. Opinion 3: Dr. McDaniel Overstates the Reliability and Viability of Using Soft Matching Data to Identify Additional U13 Accounts

7. Dr. McDaniel overstates soft-matching precision for U13 enforcement. Soft-matching produces probability scores, not definitive identifications of individuals that are necessarily U13. Automated removal of users based on probabilistic linkage (statistical connections between accounts that may suggest, but do not prove, common ownership by the same user) would necessarily result in the disabling of accounts for users over the age of 13, including siblings or parents that access their accounts from the same IP address as, or share devices with U13s.

8. Soft-matching for U13 detection also creates unique technical and legal challenges not present in other enforcement contexts. The false-positive problem is much worse for age verification than for other types of violations because ground truth cannot be easily verified.

9. Dr. McDaniel presents no reliable analysis of how many additional U13 accounts would be caught or what the false-positive rate would be if Meta were to use soft-matching data to disable accounts. His extrapolations from a three-month sample are speculative and assume all soft-matched accounts are actually underage without validation. Dr. McDaniel provides no evidence that any other platform uses probability scores akin to Meta's soft-matching to disable potential U13 accounts, or that any security or industry standard, rule, or best practice required Meta to do so.

D. Opinion 4: Meta Does Not Retain Data For Unreasonably Long Periods After Detection

10. Dr. McDaniel misunderstands industry retention norms. Data retention over extended periods is standard across the industry for compliance documentation under the General Data Protection Regulation ("GDPR") and California Consumer Privacy Act ("CCPA"), litigation hold requirements, and regulatory investigation response—sometimes across dozens, if not hundreds, of different jurisdictions. Appeals processing, re-registration detection, and law enforcement compliance may also require data availability, which Dr. McDaniel ignores.

11. Retention periods are not purely technical decisions—they involve legal, operational, and fairness considerations. For example, users generally have rights to appeal decisions that result in disabling their accounts, which requires that data remain available. None of the authorities that Dr. McDaniel cites, including NIST 800-88 (which relates to media sanitization, not account deletion timelines) and ISO 27001 (which says "no longer than necessary" but does not define "necessary" for this context), support his opinion.

12. Indeed, Dr. McDaniel cites no industry benchmark establishing that Meta's practices are inconsistent with industry standards. Many platforms have 30-day appeal/recovery windows, GDPR allows 30+ days for deletion requests, and Google, Microsoft, and Apple all have multi-stage deletion processes.

E. Opinion 5: Dr. McDaniel's Analysis Fails to Show that Many U13 Reports Were Not the Result of Bot Attacks

13. Meta experienced periods of significantly elevated U13 reporting, characteristic of automated bot attacks. Dr. McDaniel opines that neither the data nor common understandings in the field support that explanation for the elevated volumes of reports. But Dr. McDaniel's own data shows 10-40x burst activity, which is characteristic of bot attacks. Days with 40,000+ reports compared to typical volumes of 1,000-2,000 represent classic bot attack signatures: sudden spikes, short duration, and return to baseline. The SURF form rollback after 10-100x volume increases confirms the vulnerability to automated exploitation.

14. Dr. McDaniel's fan-in/fan-out analysis cannot detect sophisticated single-use bot accounts. Sophisticated bots create new throwaway accounts for each submission precisely to avoid detection through account reuse patterns. His simplistic burst detection method using mean + 2 standard deviations on rolling 7-day windows would miss sophisticated attacks that stay under statistical anomaly thresholds.

15. Dr. McDaniel also did not review data that informed Meta's assessment that bot activity contributed to elevated levels of U13 reporting, including IP address patterns, device fingerprints, and submission timing at millisecond granularity. Meta deployed expensive CAPTCHA systems in response to detected attacks, which Dr. McDaniel's analysis fails to explain if attacks weren't real. His conclusion that the data "do not support" Meta's characterization ignores data he did not review.

II. INTRODUCTION & QUALIFICATIONS

A. Assignment

16. As noted above, I have been retained by Meta through its outside counsel, to provide expert opinions in this matter—specifically, to review and respond to Opinions 1-5 set forth in the November 21, 2025 Trial Report of Patrick McDaniel and the December 15,

37. I have authored papers on topics related to Internet privacy, including some on tracking technologies, especially in the Internet of Things and "smart devices" space. Some of the technologies I have studied extensively involve capabilities and technologies that are analogous to those at issue in this case. These papers include, but are not limited to, those cited here. My CV, attached as Appendix A to this report, contains other such examples.

F. Teaching and Educational Impact

38. I co-founded Georgia Tech's OMSCS (Online Master of Science in Computer Science) program, one of the first and largest online graduate degree programs in computer science, which has educated tens of thousands of students worldwide. I created the first Software-Defined Networking (SDN) Massive Open Online Course (MOOC), which has educated over one million students globally. At the University of Chicago, I teach courses on Machine Learning for Computer Systems, Internet Censorship, and Security/Privacy, where I train the next generation of computer scientists in the principles and practices relevant to the issues in this case.

III. BACKGROUND: Age Gating at Scale

A. Industry-Standard Multi-Layered Detection Architecture

39. All major social media platforms deploy similar multi-layered detection architectures for age gating, reflecting industry consensus on the only viable approach at scale.[2] These systems share common characteristics across YouTube, TikTok, Snapchat, Twitter, and Meta platforms. [3] The architecture generally consists of three primary detection layers operating in parallel: self-reported age verification at signup, automated detection systems using machine learning classifiers, and user-generated reports processed through structured review workflows.

40. The self-reported age layer operates at account creation. Users provide their date of birth, which is stored and used to block account creation for individuals indicating

---

[2] *See* Sarah Forland, Nat Meysenburg & Erika Solis, *Age Verification: The Complicated Effort to Protect Youth Online*, Open Tech. Inst., at 24-25 (April 22, 2024) (discussing the multi-layered approach used by major platforms, including self-reported age at signup, automated detection systems, and user reporting mechanisms); *id.* at 28-29 (explaining that platforms including TikTok, Pinterest, Discord, and Google rely on self-reported age at signup, with age verification required only when users attempt to change their stated age or are flagged as potentially underage through detection systems); Noah Apthorpe, Brett M. Frischmann & Yan Shvartzshnaider, *Online Age Gating: An Interdisciplinary Evaluation*, (Aug. 1, 2024), available at https://ssrn.com/abstract=4937328 (providing interdisciplinary evaluation of online age gating technologies and implementation approaches).

[3] Throughout this report, references to Meta's "platforms" refer specifically to Facebook and Instagram.

an age below 13. However, this layer cannot protect against individuals who misrepresent their age, which represents the primary evasion vector.

41. The automated detection layer employs machine learning classifiers trained on multiple signals to identify potentially underage individuals who evaded the self-reported age check. These classifiers may analyze profile content (photos, biographical information, interests), behavioral patterns (activity times, engagement patterns, content consumption), friend network characteristics (connections to known minors, age distribution of connections), and other signals. These classifiers output probability scores, not binary determinations, requiring threshold selection to convert probabilities into actionable decisions. Notably, U13 detection is significantly more challenging than the U18 age verification employed in other regulated industries (such as online gambling or adult content platforms), as U13 individuals generally lack government-issued identification, there are legal restrictions governing the use and retention of data from U13 individuals (which significantly limits the ability to use data to develop machine learning models capable of accurately predicting whether an individual is under the age of 13), and U13 individuals may exhibit behavioral patterns that are difficult to distinguish from slightly older users.

42. The user-generated report layer processes reports from other users who believe an account belongs to a U13 individual.[4] Reports trigger human review workflows where trained reviewers examine available evidence (which may include profile photos, bio information, stated age if visible, reported age from reporter) and make determinations based on structured decision trees. For large social media companies, this layer handles millions of reports annually and requires careful workflow design to maintain consistency across thousands of reviewers operating in multiple languages and jurisdictions.

B. Automated Detection: Machine Learning Classifiers at Scale

43. Machine learning classifiers for age detection operate as multi-class or binary classifiers outputting probability distributions or scores. A typical age classifier might output a probability distribution across age ranges (e.g., 13-17, 18-24, 25-34, 35+) or a binary probability that a user is under or over a certain age. These probabilities reflect the classifier's confidence based on training data but do not represent certainty about ground truth. Classifier performance is measured using standard metrics: precision (percentage of flagged accounts actually underage), recall (percentage of actual underage accounts successfully flagged), and accuracy (overall percentage of correct classifications). At Meta's scale, even small improvements in these metrics translate to millions of additional

[4] Non-users may also report suspected U13s on the platform through online forms. *See* https://help.instagram.com/contact/723586364339719; https://www.facebook.com/help/contact/209046679279097.

correct classifications, but small error rates also translate to millions of misclassifications.

44. Threshold selection is fundamental to deploying any classifier that outputs probabilities. For example, a classifier might output a 70% probability that an account belongs to a user that is above or below a designated age. These probabilities alone cannot drive enforcement actions—the platform must set a threshold above which it takes action (e.g., flag all accounts scoring above 50%, 75%, or 90%). The choice of this threshold itself involves tradeoffs. Lower thresholds increase recall (catch more underage individuals) but decrease precision (flag more legitimate users). Higher thresholds increase precision but decrease recall. This is the precision-recall tradeoff inherent to all binary classification systems, including age classifiers.

C. Base Rate Effects on Classifier Performance

45. Base rate effects fundamentally constrain classifier performance at Meta's scale. When the true prevalence of the target class is relatively low in the overall population, even classifiers with extremely high accuracy rates will produce substantial false positive rates. This is a mathematical consequence of the base rate, not a flaw in the classifier design.

46. Consider a concrete example with hypothetical numbers. Assume that a platform has 200 million U.S. users, 5% (10 million) are users under 18 who inaccurately reported their age at account signup (i.e., provided a stated age of 18 or above), and a classifier achieves 99% true positive rate and 99% true negative rate—itself an incredibly high standard that even state-of-the-art classifiers would have trouble meeting. The classifier would correctly identify 9.9 million underage users, but it would also incorrectly flag 1.9 million legitimate users (1% of the 190 million legitimate users aged 18+). This results in approximately 84% precision, meaning that even with 99% accuracy on each class, roughly 16% of flagged accounts would be false positives requiring human review to avoid incorrectly restricting legitimate adult users.

47. The only way to improve precision without sacrificing recall is to improve the classifier's true positive and true negative rates, which requires better training data, better features, or better algorithms.

48. The operational consequences of base rate effects are substantial. In the example above, processing 1.9 million false positive cases requires significant human review capacity. Each incorrectly flagged account may trigger an appeal, requiring additional review and verification. Users whose accounts are incorrectly disabled experience immediate harm—loss of access to their social networks, content, and in some cases business operations. The appeals process itself creates friction and delay even when users are ultimately vindicated. Thus, regardless of the level of resources deployed by a company, as detection systems improve and the true prevalence of underage individuals decreases (because more are caught), the base rate problem actually

worsens—lower prevalence means even highly accurate classifiers produce worse precision. This is an inherent issue with binary classifiers and operating age assurance programs at scale.

49. Dr. McDaniel's proposed alternatives—more aggressive use of automated detection, lower thresholds, universal deployment of age estimation—would all worsen the false positive problem without solving the base rate constraint. A system that flags more accounts in an effort to identify more potential individuals under the age of 13 necessarily flags more users above the age of 13 as well, unless the underlying classifier accuracy improves. Dr. McDaniel provides no analysis of how his proposed changes would affect false positive rates or what operational burden those false positives would create.

D. Evasion Techniques and Arms Race Dynamics

50. Age verification systems face determined adversaries who actively seek to evade detection. Individuals may seek to evade age restrictions on the ability to create accounts by: lying about birthdate at signup (undetectable without external verification), creating accounts using parent or sibling information, avoiding posting age-revealing content, or using VPNs or location spoofing to evade location-based restrictions.

51. Even U13 individuals who are successfully detected and removed from a platform may seek to create new accounts immediately after removal using different information, which presents another challenge for platforms seeking to enforce age restrictions. When accounts are removed for age violations, users can immediately create new accounts with different birthdates and email addresses, which can be very difficult to detect. Platforms cannot permanently blacklist minors because minors age into eligibility—i.e., an individual removed from a platform at age 12 should be permitted to create an account at age 13. This creates a fundamental asymmetry: platforms must allow new account creation, but cannot reliably distinguish newly-eligible users from repeat evaders.

52. Dr. McDaniel suggests that Meta could use browser fingerprinting, telemetry-based identification, and device characteristic analysis to identify repeat evaders.[5] However, these techniques face significant limitations for U13 detection. Browser fingerprinting and device telemetry can identify when multiple accounts access a platform from the same device or browser, but they cannot determine which account belongs to the legitimate user. In households with shared devices—common among families with children—multiple users of different ages may use the same computer, tablet, or phone. Blocking all accounts associated with a device fingerprint because one account was flagged as underage would create unacceptable false positive rates, potentially disabling accounts belonging to parents, older siblings, or other household members. Additionally, these techniques become less reliable when users employ privacy-protective browsers,

[5] McDaniel Report ¶ 41.

Dr. McDaniel considers. Using such data for automated decision-making may also trigger additional requirements under GDPR Article 22.[10]

61. No cross-platform database of underage users exists. Privacy regulations generally prohibit sharing personal information across competing companies without user consent. Even if technically feasible, a shared database of "known minors" would face insurmountable obstacles under privacy law[11] and would become outdated as minors age into eligibility.

62. Finally, I am aware of no federal or state law or regulation, or industry rule, standard, or best practice, that required social media platforms to implement proactive age verification at account creation at any point in time, and Dr. McDaniel does not cite to any.[12]

F. Costs of Over-Enforcement: False Positive Harms

63. Dr. McDaniel's analysis focuses exclusively on false negatives (U13 individuals who evade detection) while ignoring 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 Dr. McDaniel fails to address at all. In this respect, Dr. McDaniel's opinions are inconsistent with custom and practice in the security field.

---

[10] *See* Regulation (EU) 2016/679 (General Data Protection Regulation), art. 9(1) (prohibiting processing of "biometric data for the purpose of uniquely identifying a natural person" as special category data requiring heightened protection); *id.* art. 22(1), (4) (restricting automated decision-making based on special categories of data including biometric data, requiring explicit consent or substantial public interest justification); *id.* art. 8 (imposing additional requirements for processing children's data); 740 Ill. Comp. Stat. Ann. 14/15(b) (requiring written informed consent before collecting biometric identifiers or information).

[11] *See* Regulation (EU) 2016/679 (General Data Protection Regulation), art. 6(1) (requiring lawful basis for data processing, typically consent or legitimate interest); *id.* art. 7 (establishing conditions for valid consent); *id.* art. 13-14 (requiring transparency about data sharing with third parties). Sharing personal data between competing platforms would require user consent, appropriate safeguards, and data processing agreements, making a cross-platform database of minors legally and practically infeasible.

[12] COPPA (15 U.S.C. § 6501 *et seq.*; 16 C.F.R. Part 312) prohibits website operators from knowingly collecting personal information from children under 13 without verifiable parental consent but does not require platforms to implement proactive age verification systems at account creation. *See* 16 C.F.R. § 312.5(a)(1) (requiring that operators subject to COPPA obtain parental consent for collection, use, or disclosure of personal information from children); 16 C.F.R. § 312.3 (establishing COPPA's application to operators with "actual knowledge" that they are collecting personal information from children). Section 230 of the Communications Decency Act, 47 U.S.C. § 230, provides immunity for user-generated content but imposes no affirmative age verification obligations.

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

65. Dr. McDaniel discusses face-based age prediction as one recent evolution in age assurance practices. For example, Yoti is a third-party age-checking service that estimates a person's age using an image of their face, without checking physical documents like an identification card and without human intervention. Dr. McDaniel characterizes Yoti as "accurate, privacy-preserving, and designed for youth-protection contexts," and he criticizes Meta for failing to apply Yoti for the purpose of U13 detection.[13]

66. However, NIST research on facial age estimation shows that accuracy varies significantly by demographic characteristics. 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. Dr. McDaniel claims that these solutions "boast accuracies of above 99%",[15] citing Yoti's marketing materials. This single accuracy figure is misleading. As the NIST research demonstrates, accuracy varies significantly based on sex, race, lighting conditions, and image quality. A system might achieve 99% accuracy on one demographic group under optimal conditions while performing substantially worse on others. Even Yoti's own white papers acknowledge these variations. The single "99% accuracy" figure does not reflect performance across diverse populations and real-world conditions.

67. Stricter verification requirements also create privacy invasions at massive scale. Requiring ID verification or biometric data collection from all users would mean collecting and storing sensitive personal information from billions of people, including actual children 13 and over who are legally permitted on platforms. This creates massive data breach risks and privacy harms that must be weighed against potential benefits. While it is theoretically possible that biometric data could be collected at signup and then promptly deleted after age verification, this approach does not resolve the fundamental concerns.

---

[13] McDaniel Report ¶ 323.

[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] McDaniel Report ¶ 47.

68. First, under laws like Illinois's Biometric Information Privacy Act (BIPA),[16] even transient collection of biometric identifiers requires informed consent and creates legal obligations—the collection itself, not just retention, is regulated.

69. Second, at Meta's scale of millions of daily signups, biometric data would continuously flow through collection and processing systems. As with any company at Meta's scale, this creates additional risk associated with potential data breaches, even in the context of well-designed systems with ample protections and even with rapid deletion.

70. Third, 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 creating an account.

71. Fourth, prompt deletion of data collected at signup 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.

72. Finally, 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: Constraints on System Performance

73. 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.

74. 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

---

[16] 740 Ill. Comp. Stat. Ann. 14/15(b) ("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 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 authorize representative."). I understand that COPPA also regulates the collection of data from individuals under the age of 13.

102. Finally, even if Meta were to develop machine learning models based on text and user behavior patterns on its platforms, such models likely could not meaningfully distinguish between 11- and 13-year old individuals. Dr. McDaniel acknowledges that these approaches infer a probable age range rather than verifying age with certainty in paragraph 44 of his report.[60] The reliance on age range undermines these approaches' efficacy for differentiating between 11 and 13 year old individuals because these two age demographics use similar language and have comparable use patterns on social media platforms,[61] making it extremely difficult for automated systems to reliably differentiate between them based on behavioral signals alone. This fundamental challenge further underscores why age detection at scale cannot rely solely on automated approaches and why human review remains necessary for ambiguous cases.

1.5 Security is Inherently Reactive

103. Although, as Dr. McDaniel notes,[62] Meta maintained numerous age-related machine-learning models across different product teams, these models were built for analytics purposes. Analytic models have much lower thresholds for accuracy and operational performance. In addition, analytic models are primarily optimized to identify large user cohorts, not making accurate single-user predictions. Furthermore, at the time that Meta began to develop these tools in 2018, the highly specialized technical approaches, skilled teams, and rich, accessible datasets required to develop these tools into effective U13 detection measures did not exist and take a significant amount of time to build.

104. Dr. McDaniel characterizes Meta's systems as "primarily reactive" as if this were a design failure. All security systems are inherently reactive—intrusion detection, fraud detection, spam filtering all respond to observable behavior. Proactive prevention would require identifying bad actors before they take any observable action, which is impossible.[63]

---

[60] The examples that Dr. McDaniel references in paragraph 44 of his report also all discuss differentiating between adolescents and adults and not distinguishing U13s from people over 13 years old.

[61] *See* Jason M. Nagata et al., *Prevalence and Patterns of Social Media Use in Early Adolescents*, 25 Academic Pediatrics 102784 (2025), https://www.academicpedsjnl.net/article/S1876-2859(25)00009-9/fulltext (finding that 63.8% of participants under 13 years reported social media use despite age restrictions, with early adolescents having similar platform preferences and account numbers across the 11-13 age range, suggesting comparable online behavioral patterns).

[62] McDaniel Report ¶ 142.

[63] *See* Ross Anderson, *Security Engineering: A Guide to Building Dependable Distributed Systems* (John Wiley & Sons, 2d ed. 2008) (explaining that all security systems are fundamentally reactive, responding to observed threats and adversarial behavior rather than preventing attacks before any observable action).

105. Meta engages in proactive detection and enforcement: ML systems continuously scan user content and behavior, age collection was extended to existing users,[64] and cross-platform enforcement acts on linked accounts.[65] Dr. McDaniel himself acknowledges the arms race dynamic: "if users became aware of the cues used to trigger detection, they could likely alter their behavior".[66] This permanent cat-and-mouse dynamic is the nature of adversarial security, not evidence of failure.

106. Proactive age verification at signup would require government ID collection, biometric data collection, or credit card verification. These approaches create privacy concerns, exclude users without readily available IDs, create data breach risk, and have legal restrictions in many jurisdictions.[67] Self-reported age plus reactive enforcement is the industry standard for good reason.

1.6 No Alternative Proposed and No Comparative Analysis

107. Across more than 200 pages, Dr. McDaniel offers no concrete alternatives that, in his view, would suffice to establish the standards against which he purports to assess Meta's platforms. He proposes no system architecture that would work at scale, no specific thresholds, and no cost analysis of alternative approaches. He makes no comparison to Google, TikTok, Snap, or any other platform.

108. If a competitor had solved this problem through different methods, one would expect Dr. McDaniel to cite that example. His silence suggests no such example exists, nor am I aware of any. A critique that establishes only that a system is imperfect—without proposing what the standard should be or how it could be achieved—does not establish that Meta has fallen short.

109. Dr. McDaniel makes no comparison to peer platforms. He provides no data on what YouTube, TikTok, Snapchat, or Twitter did during the same period. He does not analyze whether other platforms achieved better outcomes. The absence of such comparison suggests that other platforms faced similar challenges and used similar approaches.

110. Dr. McDaniel performs no false positive analysis despite extensive data access. He performs no cost-benefit analysis of proposed alternatives. He provides no user impact assessment of stricter enforcement. He does not analyze why certain

---

[64] McDaniel Report ¶¶ 130, 131.

[65] McDaniel Report ¶¶ 128, 129.

[66] McDaniel Report ¶ 140.

[67] *See* Forland, Meysenburg & Solis at 21 ("Youth between the ages of 14 to 16 years old often do not hold any official form of government identification, and those under the age of 18 are unable to hold credit cards. Millions of adults who are 18 years of age and older do not hold a valid government-issued photo identification.").

multi-stage filtering: automated signals → human review → appeals. This is exactly how to handle low-prevalence detection at scale and represents standard industry practice.[86]

2.6 Human Review Evasion Claims Lack Support

128. Dr. McDaniel argues that human review processes are vulnerable to evasion by underage users who can manipulate their profiles to appear older.[87] However, this critique lacks empirical support and mischaracterizes the available evidence.

129. Dr. McDaniel cites a study from Academic Pediatrics (Nagata et al., May-June 2024) as evidence that underage users can easily evade human review. However, this study is a prevalence survey examining rates of social media use among adolescents—it says nothing about evasion methods, human review processes, or how underage users might circumvent age verification systems. The study does not support Dr. McDaniel's characterization that human review is "easily evaded."

130. Moreover, the claim that users can evade detection by altering their profiles applies equally to any detection system, whether automated or manual. Automated classifiers rely on the same profile signals (photos, biographical information) that human reviewers examine. If a user successfully presents as older in their profile, both automated systems and human reviewers face the same challenge. This is not a deficiency specific to human review—it is an inherent limitation of any system that relies on user-provided information, which Dr. McDaniel fails to acknowledge.

131. Dr. McDaniel provides no evidence that human review is more susceptible to evasion than automated detection, nor does he propose an alternative approach that would be less vulnerable to such manipulation. The evasion problem he identifies is fundamental to age detection at scale, not a failure of Meta's human review processes specifically.

VI. OPINION 3: Dr. McDaniel Overstates the Reliability and Viability of Using Soft Matching Data to Identify Additional U13 Accounts

132. Dr. McDaniel argues that Meta should have used soft-matching data more aggressively to identify additional U13 accounts.[88] He claims soft-matching is a "low-cost" approach that could have identified "tens of thousands" of additional underage accounts, and he cites internal documents discussing soft-matching as evidence that Meta ignored an obvious solution. These criticisms overstate the reliability of soft-matching, ignore the

---

[86] *See* Forland, Meysenburg & Solis at 28-29 (describing the multi-stage approach used across platforms: automated detection systems flag potentially underage accounts, which then trigger verification requirements).

[87] McDaniel Report ¶¶ 209, 210, 342.

[88] McDaniel Report ¶¶ 286-289.

false positive consequences of automated enforcement based on probabilistic linkage, and mischaracterize internal engineering discussions. I address each in turn below.

3.1 Probabilistic, Not Deterministic

133. Dr. McDaniel presents soft-matching as if it conclusively identifies accounts controlled by the same user. This is not the case. Soft-matching produces probability scores based on multiple signals, including IP address patterns (whether accounts connect from the same network locations), device fingerprints (browser characteristics, operating system details, screen resolution), session timing patterns, cookie data, and other behavioral signals.[89] However, soft-matching does not produce certain identifications.[90] Shared household devices, dormitory networks, and library computers create false matches between accounts controlled by different people.[91] If Account A is confirmed under-13 and Account B is soft-matched to it, there is no certainty that Account B is also under-13.

134. Another common false positive scenario occurs when a parent creates an account for their child using the parent's own email address, where that email address is also associated with the parent's own account. Soft-matching could flag both accounts as potentially belonging to the same user, and if the child's account were identified as underage and disabled, the soft-matching system might incorrectly suggest the parent's account should also be reviewed or disabled. This illustrates how legitimate, privacy-protective parental oversight practices can trigger false positives in probabilistic matching systems.

135. Moreover, Dr. McDaniel's soft-matching analysis assumes that the original U13 determination for Account A is itself certain. But as discussed above, age predictions are probabilistic. If Account A was incorrectly flagged as under-13 (a false positive), then any accounts soft-matched to it are being linked to an erroneous determination, compounding the false positive problem. This is especially true given that accounts may be

---

[89] *See, e.g.*, META3047MDL-146-00072998 at 3 (matching models or heuristics, account clustering models, unlinking a previously linked account); META3047MDL-146-00072957 at 2 (common keys, same device); *id.* at 3 (previous IP users); *Id.* at 3-5, 7 (cookies); *id.* at 3,6 (same session for pixel only); META3047MDL-044-00112891 (social graph similarity). Common keys include phone number and email, *see* META3047MDL-047-00016158 at 28, same name and birthday, *see* META3047MDL-044-00112891, and common photos, *see* META3047MDL-146-00072998 at 3.

[90] *See* Hartnett Dep. 232:1-23 (describing the "accuracy" of soft-matching as "not pass[ing] what we usually think of as . . . high accuracy" because it looks at accuracy with users in the aggregate, not with individual users); META3047MDL-146-00072994 at 00072995 ("Facebook uses our models . . . to predict the probability of a FB user match[ing] an IG user account." (emphasis added)).

[91] *See* META3047MDL-044-00112891 (Cross-App Reference Types) (stating that soft-matching uses device and IP address information to connect accounts, so those logging in from the same device or location may be inappropriately matched).

removed from the platforms through the checkpointing process where a user fails to provide age documentation during the requisite appeals period, even where that user is in fact over the age of 13.

136. For these reasons, Dr. McDaniel's conclusion that he "found tens of thousands of accounts that Meta had disabled for being underage, which Meta's systems predicted were linked to tens of thousands more accounts that Meta ignored"[92] is unreliable. Dr. McDaniel fails to account for both the probabilistic nature of the confidence score as well as uncertainty regarding whether the original account (to which a different account is soft matched) in fact belonged to a U13 individual.

137. In addition, Dr. McDaniel does not explain how soft-matching data could be used in practice within Meta's human review workflows. If a reviewer sees that Account B has a 65% soft-match probability to a disabled Account A, what action should they take? How would the knowledge that Account B is soft-matched to Account A implicitly impact the reviewer's determination? How does that soft-matching data layer on to other aspects of the human review processes? There is no clear decision rule: the probability score does not tell the reviewer whether the accounts share an owner, whether that owner is underage, or how to weigh this signal against observable profile characteristics. Integrating probabilistic linkage data into human review would require developing entirely new decision frameworks, training thousands of reviewers on statistical interpretation, and accepting increased inconsistency as reviewers weigh ambiguous probability scores differently.

3.2 Over-Enforcement Risk

138. Automatically disabling accounts based on statistical association with a previously disabled or flagged account would harm innocent users. For example, siblings, family members, and roommates sharing devices or networks would face automated enforcement despite being eligible users given the inputs into the soft matching data. I am not aware of any other social media company that uses this type of probabilistic matching to automatically disable potential U13 accounts, and Dr. McDaniel cites no authority demonstrating that other platforms employ this approach. Nor does he point to any rule, regulation, or industry standard that would require it. The false positive consequences would be significant at Meta's scale, which is why this "guilt-by-association" approach has not been widely adopted for age enforcement.

139. The over-enforcement risks associated with using soft-matching for U13 detection also highlight why soft-matching may be an appropriate tool in certain contexts, like using it to target advertisements, but not others. Specifically, the downsides of using Meta's low-precision soft-matching model for advertising, such as showing users less relevant ads and slightly increasing costs for Meta and advertisers, are relatively minor

[92] McDaniel Report ¶ 283.

when compared to the downsides of using a low precision model for U13 enforcement, like erroneously removing account access for users over the age of 13. In other words, soft-matching in the advertising context risks only making Meta's advertising product less efficient whereas soft-matching for U13 detection risks removing access to Meta's product entirely for a potentially substantial number of users over the age of 13.

3.3 Different Use Cases

140. Meta's use of soft matching data in limited contexts does not, as Dr. McDaniel suggests, mean that soft matching data is appropriate to use for U13 identification and removal—let alone that Meta fell short of security standards by failing to disable soft matched accounts. Meta primarily uses soft-matching for platform growth, like facilitating friend and follow recommendations across Facebook and Instagram, and detecting potential off-platform harms, like child safety. Soft-matched inferences facilitate cross-app friend recommendations and allow the company to dedupe users across apps and calculate the number of unique users in its quarterly earnings reports.[93] Meta also employs soft-matching, on a limited basis, to identify repeat violators of its Terms of Service. The company specifically targets the "worst of the worst" on its platforms—clear policy violations with acute dangers, such as human and child exploitation, terrorism, and mass murder, on the linked accounts.[94] Meta's age enforcement is different because it relies on probabilistic inference about identity without a clear violation on the linked account.[95]

3.4 Speculative Extrapolation and Missing Analysis

141. Dr. McDaniel extrapolates from a three-month sample to estimate "tens of thousands" of additional detectable accounts.[96] This assumes constant soft-match rates over time without supporting evidence. As described above, it also assumes all soft-matched accounts are actually underage, which he does not validate. His projections are speculative and he presents no reliable analysis of false positive rates for soft-match-based enforcement.

3.5 Internal Document Context

142. Dr. McDaniel cites internal documents discussing soft-matching as evidence Meta ignored an obvious solution. These documents show Meta actively considered soft-

[93] *See* Hartnett Dep. Ex. 6, at 11.

[94] *See* META3047MDL-146-00072994 at 00072996 (listing the policies soft-matching is employed to target).

[95] *See* META3047MDL-050-00014483 at 00014485 ("[T]he consequences of showing the wrong ad to someone is relatively low compared with the consequence of being checkpointed as an underage account.").

[96] McDaniel Report ¶ 286-289.

matching, not ignoring it.[97] Internal discussions showing debate about trade-offs are evidence of healthy engineering process, not negligence. "Difficult to defend" does not mean "technically feasible and appropriate," nor does it mean required by security or industry practices or standards. Single quotes without full context of meetings and follow-up discussions can mislead.

VII. OPINION 4: Meta Does Not Retain Data For Unreasonably Long Periods After Detection

143. Dr. McDaniel criticizes Meta's data retention practices for detected U13 accounts, arguing that: (1) Meta retains data for unreasonably long periods after detection; (2) this retention exposes children to "unnecessary risk"; and (3) Meta's practices violate standards such as NIST 800-88 and ISO 27001. These criticisms misunderstand and misstate industry retention norms, ignore the legitimate operational and legal reasons for retention, and misapply technical standards designed for different contexts. I address each in turn below.

4.1 Industry Norms

144. Dr. McDaniel implies Meta's retention periods are unusually long. Data retention for extended periods is common across the industry for compliance documentation under GDPR and CCPA, litigation hold requirements, and regulatory investigation response. Dr. McDaniel cites no industry benchmark establishing that Meta's practices are outliers.

145. Most platforms have multi-week appeal/recovery windows. GDPR requires responses to erasure requests within one month, with extensions up to two additional months permitted for complex cases.[98] Google, Microsoft, and Apple all have multi-stage deletion processes. Meta's 90-day backend deletion period represents a reasonable amount of time for distributed systems to propagate deletions across geographically distributed data centers and backup systems.[99]

[97] *See, e.g.*, META3047MDL-050-00014483 at -14485 (discussing whether soft-matching may be a useful age-verification tool and the potential downsides of doing so); *see also* Hartnett Ex. 6 at 13 (noting that work on soft-matching use cases was hindered by a German FCO (Competition Authority) order prohibiting Meta from using soft-matching in certain contexts).

[98] *See* Regulation (EU) 2016/679 (General Data Protection Regulation), art. 17(1), (3) (requiring controllers to erase personal data "without undue delay" upon request, subject to exceptions); *id.* art. 12(3) (requiring responses to erasure requests "without undue delay and in any event within one month of receipt"). GDPR does not mandate immediate deletion but allows reasonable time for technical implementation across distributed systems.

[99] Hartnett Dep. Ex. 6 at 13-15.

4.2 Legitimate Operational Needs

146. Retention serves concrete purposes beyond pure technical requirements. Appeals processing requires that data remain available for users to challenge enforcement. Users deserve time to appeal and gather documents to prove their age. Re-registration detection requires historical records to identify returning users. Law enforcement requests require responsive data to be available. Legal holds for lawsuits and regulatory investigations require data preservation.

147. Data deletion is irreversible—adequate time is needed to ensure accuracy and avoid permanent false positives. Dr. McDaniel does not propose what retention period would be appropriate or analyze the operational consequences of shorter retention.

4.3 Risk Management During Retention

148. Dr. McDaniel claims retained data exposes children to “unnecessary risk.” Accounts are disabled immediately upon age determination. Data is marked for deletion and access is restricted during the retention period. The data is not actively used in production systems. “Retention” does not mean “actively accessible.” Staged deletion reduces risk of accidental deletion and ensures completeness across distributed systems. Dr. McDaniel identifies no specific breaches or incidents involving deletion-flagged data.

149. There is a contradiction in Dr. McDaniel’s reasoning: he characterizes data retention as exposing children to “unnecessary risk,” yet he simultaneously advocates for more aggressive data collection (facial images, government IDs) that would create far greater privacy risks for all users, including children. The risk of retaining disabled account data for a defined period—data that is access-restricted and marked for deletion—is minimal compared to the risks of universal biometric or ID collection that his proposed alternatives would require.

4.4 Standards Misapplication

150. Dr. McDaniel cites NIST 800-88, which is about media sanitization (hard drives, etc.), not account deletion timelines. He cites ISO 27001, which says data should be retained “no longer than necessary,” but it does not define “necessary.” These standards do not specify deletion timelines for this context, do not account for user appeals and legal requirements, and do not address balancing deletion speed against false positive risks.

The undersigned hereby certifies their understanding that they owe a primary and overriding duty of candor and professional integrity to help the Court on matters within their expertise and in all submissions to, or testimony before, the Court. The undersigned further certifies that their report and opinions are not being presented for any improper purpose, such as to harass, cause unnecessary delay, or needlessly increase the cost of litigation.

January 9, 2026

Nick Feamster, Ph.D.