


EXHIBIT 6



I. Collection, Use, and Disclosure of User Data (Topics 2(a), 2(b))

A. SME: [REDACTED] (Advertising)

- Privacy Policy: how data collected is used for advertising purposes
 - Privacy policy does not make a distinction between teens and adults on this point
 - Meta provides clarification on additional limitations for teen users in the privacy center that says that the only data we use for teens in ads are age and location. This is not included in the policy.
 - Categories of data that are addressed in the privacy policy as data we use to show ads
 - Profile information
 - Activity off our products
 - Content that you create across the app
 - Inferences
 - Meta uses stated age to target ads to users
 - Information we share with 3rd parties is aggregate, anonymized reporting data eg: total impressions an ad got from 13-17 year olds
- Data use and targeting restrictions on Youth Ads
 - Delivery means what we do to show an ad - most of the time this term refers to Ranking
 - Targeting refers to an advertiser selecting their audience
 - For BAY, options that advertisers have for targeting are minimal. In BAY, you can only choose age and city-level or higher location.
 - Pre-2021, Youth Ads used all data sources but had content restrictions for youth (eg: alcohol, gambling). All users including teens had the ability to opt out of the use of 3P data.
 - 2020-2021 time period, Meta began to look at a different approach for Youth Ads targeting. With GDPR, we introduced low-data use ads types in Europe for certain teens. In 2021, other laws passed that required that we did not create a "profile" of a teen (ie, the company may not use data to create a prediction of a teen's interests for advertising purposes) and Meta therefore decided to launch the low-data version globally to U18.
 - Slimmed Down Services - SDS. Teens were defaulted into the setting where 3P data was off for ranking, targeting and delivery use cases. Targeting was only on age, gender, city-level or higher location. Launched in Sep 2021.
 - Basic Ads Youth Tier: Announced in Jan 2023. Began to make advertiser side changes. Started rolling out in March 2023 to users. Launch

happened in May-Aug 2023 for the US. The only data Meta now uses to show a teen an ad is age and city-level or higher location.

- Content-based restrictions for Youth Ads
 - Ads policies on content have always been applied across all versions.
 - Content gating is enforced on predicted age
 - Advertising standards (externally published): There are a series of policies that outline the content we show to Youth in ads.
 - Age gating:
 - Alcohol
 - Online Gambling
 - Tobacco
 - Drugs (OTC and pharmaceuticals)
 - Weight loss & cosmetic products
 - Adult products
 - Adult nudity
 - Dating ads
 - Financial and insurance products
 - Weapons accessories
 - Prohibited for everyone: eg: gun content, overlap with community standards
 - How enforced:
 - Advertiser sets up targeting criteria
 - If advertiser for example says that they want to show alcohol ads to teens then ad is blocked and is not posted
 - There is also reactive enforcement where we monitor all ads to ensure they are not violating age gating or prohibited topics

B. SME: [REDACTED] (Logged-Out Data)

Scope:

Facebook: 2014 to present

Instagram: 2019 to present

- Data collected from logged out individuals
 - FB and IG follow similar principles for the implementation of logged out data collection
 - We do collect data before a user logs in to the website or the app. We collect raw data on IP address, device ID, cookie data, button clicked, page visited. We do not collect any personal data. Eg: user ID, password, age, etc. The only exception is if you are in a logged out research session where we explicitly got user consent.

- If the user gets to a signup flow and types in data and then they cancel, they may have some of the data they typed in logged, but they agree to TOS right up front.
 - We use this logged out data for security. The data is therefore very locked down - only limited part of the team can access.
 - Other use case is growth analysis. We want to build our logged out pages in a way that helps the product grow.
 - Before 2020: it was possible for other teams in the company to use the logged out data but there were not any known use cases for advertising. After 2020, we locked down the data behind an access control accessible only by Security and Growth teams.
- Data collected and shared for marketing
 - There is an effort in Meta to advertise our own apps on other properties. For that effort, we do some targeting. Eg: if the user is already an active FB user, we do not need to show them ads. To execute on this, we have been communicating with other ad networks (eg: Google) using pixel data (industry standard). Any data used in this communication is totally anonymized. We do not rely on any logged out meta data to target users - we only use data that Google provides.
- Retention
 - All of the data we collect for logged out is only retained for 35 days post-2024.
 - Prior to 2024, there were some cases where logged out data was retained longer than this - 90 days, then 45 days.

C. Documents

- Privacy Policy (TAB 1)

- If yes to any questions except non human entity then account enters UFAC
- Meta's Oversight board maintains a list of people that are journalists, political figures, major celebrities. If an account is flagged that matches this list of people, it is intercepted on way to UFAC and reviewed by FTE against the policy
- Once in UFAC, you have 30 days to appeal
- If do not appeal in 30 day window, begins the clock for deletion
- If a user appeals and hard fails [explanation in other notes], clock for deletion begins immediately.
- Deletion process for online data (data in prod)
 - Deletion process begins on day 31 and takes 90 days to complete
- Abuse of the U13 reporting flow
 - One of most abused reporting flows at Meta
 - Well known way to get accounts taken down because it ends in account suspension / checkpoint
 - Top reported account in 2024 was the IDF by an extremely large margin
 - Other reported accounts, several times a day, Kim Kardashian, @creators handle, Logitech the mouse company was reported 200 times over the course of a year
 - **Stats:** in the US, ~14% of human reviewed reports are checkpointed
 - This varies by market. In Brazil, 60% of human reviewed reports are checkpointed. In India, only 1% are checkpointed. In India the abuse vector is primarily men harassing women. In the US the abuse vector is public figures or teen bullying.
- Cross-app enforcement
 - When an account has been triggered for a checkpoint, this checkpoint propagates to any hard-linked accounts, meaning accounts the user has explicitly linked
 - The user is only able to appeal from the original seed account that was checkpointed
 - If the user fails or does not complete appeal, all accounts are deleted

B. SME: [REDACTED] (cross-app enforcement)

- Underage enforcement in mid 2021
 - Using UFAC on Facebook and an old checkpoint experience on Instagram
 - At some point previously, cross-app propagation had been set up. This worked for hard linked accounts. It was set up at some point prior to 2020. At the point Rohan inherited it, it was broken.

- Late 2021, fixed cross-app propagation. At some point in the past, the way the triggers worked had broken.
- Cross-app propagation, 2021 to present
 - If a report comes in, we assess and then checkpoint them
 - They are put in a checkpoint or appeal experience for 30 days where they can submit an ID
 - At the point of checkpoint, we propagate the checkpoint to the hard linked account.
 - You can only create an appeal from the source account.
 - (Eg) If you are checkpointed on IG, your FB account is disabled. You will be asked to log in to Instagram to look at the enforcement. This works in both directions.
 - Cross-app propagation infrastructure supports multiple types of enforcements in addition to underage
 - Only works for hard link propagation. Looked at possibility of doing soft link propagation.
- Abuse of the reporting form
 - Most of our reporting processes for other issues are in app
 - In the case of underage, the reporting form was at the time an off platform / logged out contact form. We needed that because we wanted people without accounts to be able to report underage (eg: teachers).
 - As a result of it being off platform / logged out contact form, we experience a lot of abuse in the form.
 - Manual reviewer is told if unsure the account is underage, take it down. So it was therefore an attack vector where people would mass report accounts.
 - Introduced IP-based rate limiting and CAPTCHA. Also global-based rate limiting.
 - Never got to the bottom of the motivation of the attackers.
 - Scripting attack examples:
 - Delta checkpoint experience - multiple checks across devices that you have logged in. "Did you log in on this device?"
 - A bunch of malicious users got access to accounts via eg hacking passwords and they then got the Delta experience. They would get blocked by 2FA.
 - They figured out that if you reported the account as underage and it was a young-looking account, the underage checkpoint was higher priority than Delta checkpoint.

- Underage checkpoint was logged in, if they were logged in, they could submit an ID and when the appeal is successful they were given logged in Delta experience.
- This was just on IG.
- Reports of 600K users over the course of a few hours from this type of attack.
- There were 2 or 3 large scripting attacks like this that happened before we introduced friction.
- The scripting attacks were fixed in part by moving to IG UFAC so it was a logged in experience (2022).
- Scripting attacks were not the only type of form abuse:
 - Eg: celebrities / influencers frequent target of underage reporting
 - Legacy checkpoint experience on IG. In order to appeal, you were provided a contact form that you had to fill out. Before the migration to UFAC happened, that form was also subject to abuse as well. Attacker would submit an appeal with an ID that was underage on purpose to make sure the account was down forever.

C. SME: [REDACTED] (2021 Backlog)

- Background
 - Jan 2021: Increase of FB and IG jobs because of more accounts being created overall and raters being sent home due to Covid
 - Flows that go to human review:
 - Edit DOB: When a user requests to change Edit DOB.
 - There was discrepancy between IG and FB:
 - On IG, you could not select a birthdate under 13 because that does not meet TOS
 - On FB, when you edit DOB you can put your age as U13 so Edit DOB queue needs U13.
 - Experience was standardized to include ability to put a U13 birthdate
 - U13 reporting: User account is placed in underage human review queue. Account if believed to be U13 is checkpointed.
 - U13 checkpoint appeal queue is a separate queue
 - Volume of backlog:
 - 2.3M reports across all three queues Jan 2021. Peaked at 2.7M in April 2021.
 - Backlog was outpacing Meta's capacity to review because there continued to be incoming demand.
- Meta's efforts to reduce backlog

- With prior improvements, we can keep up with incoming demand but were not working on backlog at a fast enough pace
 - Requested additional capacity as well as a reprioritization of reviewers
 - \$2.5M increase in budget for these review flows to hire and train new reps
 - Brought in other reviewers from related queues at Meta to bring in to U13 review. You cannot always do this due to reviewer training, but if the reviewer has the right training.
- Results
 - May 17 2021, all of the levers were launched but there was still a 1.2M backlog
 - 453K in U13 enforcement queue
 - 63K in underage appeals queue
 - 733K in EditDOB queue
 - To address remaining backlog we implemented automations for the U13 backlog only as a one time measure and were not applied to new jobs coming in:
 - Auto ignore jobs with college or university degree signal. Account was not checkpointed. If account got another report, the account would be reviewed again and would not have used that automation.
 - Source of data: user-supplied. Or affiliation to a university or college group.
 - Auto ignore accounts with work experience signal.
 - Together these automations addressed 33% of the remaining backlog.
 - Edit DOB backlog:
 - 733K total - 718K from FB, 15K from IG
 - Auto reject entire remaining backlog of EditDOB requests. We believe it was less risky to require the users to not get age update granted. If they want to they needed to resubmit. This is lower risk than auto-accepting.
 - With these measures, by Aug 2021 the backlog was close to 0 on FB and IG. Existing supply of human reviewers then was able to manage incoming capacity.

D. SME: [REDACTED] (soft matching background)

- Background

- Meta had a centralized team who owned soft matching
- Team used phone number and email to match across and within FB and IG
- Two kinds of emails and phone numbers:
 - Email and phone numbers that Meta collected directly - account creation, security (2FAC). Meta had phone number or email for all users because required during registration
 - Email and phone numbers where user did not provide directly but Meta inferred that it belonged to the user. For example, address book contacts could be imported by a user and Meta would then match a contact from her address book with her friend with the same name
 - FTC order 2019/2020: deprecated use of soft matching for ads purposes and more detailed disclosures. Highlighted uses like growth and integrity in the disclosures.
- Use cases
 - 1. Growth
 - Friend recommendations w/in and across apps
 - 2. Measurement
 - Meta began to report unique users in quarterly earnings. Using soft matching to dedupe users across apps.
 - 3. Ads
 - Matching contact points against advertiser consumer info - eg: Macy uploads a list of consumers and Meta targeted those ads. This is now deprecated.
 - 4. Integrity
 - Identifying people who were repeat offenders of Meta's TOS
- Accuracy
 - Meta measures accuracy of the inference. Different thresholds for different use cases - product team decides. If consequences were higher for inaccurate soft matching (eg: integrity use case) use higher precision. Not every use case used same accuracy.

E. SME: [REDACTED] (soft matching background + analytics)

- Beginning of family metric reporting
 - Launched financial statements with family metric reporting in Q4 2019
 - In 2018, we were stating MAU and DAU metrics for FB alone in our financial reports. We wanted to be able to share a more expansive metric that demonstrated the whole family.

- This started using soft matching to begin reporting family metrics. In its disclosures, Meta now seeks to report unique users using at least one of Meta's products.
- Family metrics represents an estimate. When we just had FB, we could get to a distinct count that was detailed. Family metrics are a model-based metric and there are large errors bars that we also report in our financial statements. That is why we do not report the metric at a more detailed level than what we report in financial statements today.
- Prior to this metric being in the financial statements, this process was highly experimental. It was sparsely used, were not sure it would work.
- Method for calculating family metrics
 - Meta uses models to seek to count the individual people behind user accounts
 - This includes matching accounts to an individual user within and across products
 - This requires significant judgement and models are tuned on a quarterly basis. They are subject to data limitations and variances that come along with any model based metric.
 - Family metrics can be noisy. From day to day the number can fluctuate due to the model shifting. Its very hard to tease out what is model vs. user fluctuation. We do not analyze trends at a granular level because we know a lot of the variance is the model.
 - We do not report family metrics on a daily basis due to variance.
 - Publicly report family metrics on a global rounded basis only.
- Ground truth
 - Ground truth is survey-based. This requires statistical analysis to extrapolate the survey to the population from a sample. On a quarterly basis, we will tune modeled metric to the survey metric. Surveys would size the proportion of users on one or many apps. The survey is daily on the other apps. Survey asks "How many times have you used Instagram in the last 30 days? Do you have one or more accounts?" and asked about other apps
 - Survey mean always has error bars.
 - This ground truth is not at a user level and there is no human labeling. We attempted the labeling but there was no consistency with the survey. The survey had a better premise for accuracy vs. human labelers which we did not find to produce high quality data. Abandoned labeled ground truth.
 - We cannot correctly predict a single user from this methodology and that is not our aim
- Difficulties in attributing multiple user accounts to a single person

- Did not attempt to get accuracy at a single user level
- Focused on accuracy at the level of hundreds of millions of users
- Highly inaccurate to identify whether a specific account is definitively part of an individual cluster
- We do not have confidence that the cluster data is accurate. We therefore only use soft matching for very limited use cases.

F. SME: [REDACTED] (soft matching for integrity)

- Impetus for changes
 - Work on softmatching came out of order from German FCO (Competition Authority) that Meta could no longer soft match accounts
 - There are soft matching use cases that are at the core of our work, which includes integrity
 - The FCO order allows certain exceptions including integrity purpose soft matching
- Integrity purpose reviews
 - Product teams use this forum / process to request to use soft matching for an integrity purpose
 - The IPR only reviews: does this use case serve an integrity purpose? The implementation of using soft-matching was handled by a separate XFN
 - Office hours: teams can come in to explain soft matching use cases. IPR template was filled out. It was a necessary step for every single review. Template documents our review process and provide evidence for our ultimate conclusion as to whether use case serves an integrity use case.
 - Information included: what team are you mapped to? What is your mission statement? What are the details of the proposed use case?
 - Teams needed to identify at least one Meta policy that the use case would further.
 - Product team was responsible for making the case to IPR. We approved way more IPRs than rejected in part because it was such a big effort. Product teams had good understanding of what an integrity use case and tended to bring in projects that were indeed integrity purposes.

G. SME: [REDACTED]

- Background:
 - Three states: checkpoint, disable, delete
 - Once we start the deletion of an account, we treat it the same way as any other account that is going to be deleted

IV. Age Collection / Age Verification (Topic 4(g))

A. SME: [REDACTED] (general background)

- Detection and enforcement of U13 during Age Collection at sign up
 - 5-12 experience
 - If gives this age, user continues account creation flow to the end, then given a neutral error
 - They can try again, if they again enter 5-12, they again finish the flow, get a neutral error and are unable to try again for 12 hours [see blocking notes below]
 - U5 experience
 - Assume it is an error when a user gives a DOB U5, give a neutral error, unlimited retries
- Edit DOB flow
 - U18 aging up to O18
 - User is not checkpointed but does have to provide age verification
 - Any user who edits DOB to U13
 - If they do this, they are enrolled in a checkpoint. In order to get out of the checkpoint, they need to submit an ID.
 - Checkpoint stops the user from accessing their account at all and makes their account not appear on the platform
 - We have recently found that 70% of accounts who edit DOB to U13 seem to be involved in fraud. They seem to be running ads without paying for them or they are running ads that are in violation of our policies. They are purposefully reporting their accounts as U13 so that their account gets deleted. Length of time this has been going on is unclear but it has been going on since before April 2024.
- History of age collection
 - FB
 - FB has collected age at signup since before 2012
 - Since before 2012, FB's TOS have required users to be at least 13
 - In current age collection, FB signup screens populate today's date.
 - IG
 - Mandatory age collection on IG started in Dec 2019. It was mandatory only for new signups.
 - IG TOS have required users be at least 13 since before 2012
 - Instagram's age collection has always had DOB pre-populated with today's date
 - Instagram implemented optional age collection for users who did not already have an age in Jan 22 through May 22.

- In initial testing period, Yoti adult unblock rate was 54% for Yoti compared to 78% for IDV and 48% for social vouching.
- Yoti published data on “mean average error” Yoti provides estimates of how off their age predictions are for certain users. For 6-12 years, they report 1.56 mean average error. Meta does not affirmatively leverage Yoti’s technology to identify U13s nor can any checkpointed user identified as U13 use Yoti FBAP to get back on to the platform.
- However, Meta checkpoints any users predicted by Yoti to be U13. They then need to use IDV to regain access.
- For 13-19, mean average error is 1.52 years.
- Post rollout of Yoti:
 - Since 2022, we have continually monitored block and unblock rates to ensure consistency.
 - To do this we leverage labeled birthday posts for ground truth to evaluate rate at which Yoti blocks users who are actually U18 or unblocks who are actually O18.
 - Meta does not internally test Yoti’s accuracy or effectiveness in identifying U13s as that is not the use case we leverage it for.
- Broader age verification usage and use cases
 - Beginning in Dec 2022, Meta began rolling out the age verification options (IDV, FBAP) to verify users in some situations where users were predicted to be U18 but Meta wanted to gate the product to O18: Dec 2022 for FB Gating, FB Marketplace Oct 2024.

C. SME: [REDACTED] (Age Gate; Block Enforcement)

- History of age gating at signup
 - FB implemented age gating at signup at least prior to 2009. We maintained 10 different versions of the registration flow throughout the 2010s depending on which interface the user was employing. Age collection was different on each flow.
 - In 2019, we began to explore not defaulting to a certain age on FB. The pattern we had landed on prior was today’s date minus 25 years - this was a 2017 launch. This format was helping users on bad phones who could not use a scroll.

- Changing this was a huge change to registration. We spent almost a year testing different ways to collect age in a different way. Where we landed at today's experience. IG also began to collect age with this method at the time.
- We default to today's date. Depending on interface, picker is different.
- If you select age 0-4, we feel comfortable that was not a real entry and then we give you one more chance with original picker and then on third try we do age fallback of "what is your age?" We then give an interstitial that says "We are setting your birthday as X date based on today's date + age they picked.
- 5-12: we assume it is more likely to be genuine. We give one retry and then you are blocked from account creation for 12 hours.
- We unified the age collection experience across IG and FB by 2021 to the experience defined above.
- Prior to 2021, there were varying practices exactly how the age collection experience was implemented based on the interface the user was using. Eg: FB Lite is for phones with bad internet and so there were technical limitations to how things could work.
- Details on 12 hour block period
 - 12 hr block is implemented with cookie on web and app id on app.
 - Post 2021, 12 hour block consistently in place.
- Data collection during registration
 - Before you register, we have a pre-TOS data collection policy. We do not send any data to the server - only client side data until you click submit. We only retain minimum data necessary to achieve the specific purpose of registering. The only exception is that we have to retain the block. We must store that we are blocking you.

D. Documents

- Presentation re Registration Flows (Tab 8)

V. Age Modeling (Topic 6)

A. SME: [REDACTED] (Historical Background)

- Pre-2018 Age Modeling
 - All age models were distributed across multiple product teams working on models that were relevant to their own product teams
 - Eg: Integrity would build an adult classifier for use case around harmful behaviors
 - Eg: Age Affinity team built model to help guide product decisions
 - Eg: FB built a life stage model to understand users of different ages
 - Very little documentation, many of these models were not actually used
- 2018
 - Survey and consolidation
 - Core Dimensions was formed. Collected all of the Age Models across the company and took ownership of them.
 - Cleaning up pipelines - looked into how the models were built
 - A key part of building models is establishing “ground truth.” The ground truth is established if 2 labelers agreed on the persons age.
 - This “ground truth” is not real truth. This involves a person looking at a protocol and a profile and trying to guess age but we do not know if it is accurate.
 - The team determined age model performance was not very good and decided to just prioritize investment in the 18 boundary
 - U13 model
 - Decided not to do U13 model for reason above
 - 2019
 - Made architectural decisions to ensure consistency across models
 - Eg: Age affinity model is nested within the adult classifier
 - Age affinity follows within each branch of the adult classifier
 - Also deprecated age models including Age Signals on IG, Lifestage consolidated into Age Affinity. Age Affinity / Lifestage is only for analytics.
 - By 2019 the only models were Adult Classifier and Age Affinity
 - Dating was the first product use case in 2019. Integrity interventions had leverage models prior to this.
 - 2020
 - Still focusing on improving the quality of the adult classifier
 - 2021

also saw many shared accounts, parent-managed accounts, pet or business accounts are false positives.

- Possible that users can learn and game the heuristic over time.
- If the person is a legitimate user above 13 but does not have an ID, they have no way of appealing.
- We ran into concerns that we do not have sound measurement on phase 1 heuristics so we prioritized the measurement work and clarifying the policy.
- Established a policy where we enforce on a user with high confidence signals that they are representing their age as U13. Team has built out measurement of this with SME program.

D. Documents

- Core Dimensions Presentation (Tab 9)