

FILED UNDER SEAL

EXHIBIT 21

[A/C Priv] Under-13 Age Gating + Enforcement

Cross-Family Review (3/3/20)

Goal

Decide on our approach to collecting age at registration and for enforcing our minimum-age policy after registration.

Agenda

1. Recap the issue of doing more on u13s
2. Make decisions on:
 - a. Workstream 1: Improving Age Collection at Reg flows globally
 - i. **Question 1**: Should we change the current default age on the birthday screen from 25 years old to an age below the minimum required age (i.e. 0 or 1)?
 1. (Previously decided yes, but looking for final confirmation given results)
 - b. Workstream 3: Improving u13 Age Enforcement
 - i. **Question 2**: Should we invest in a u13 classifier to be used for app level enforcement? *If yes, classifier would be used to queue accounts for our existing CO underage review; no proposal to checkpoint accounts without human review.*
 - a. **OPTION 1**: Invest in u13 classifier, commit to using it for app level enforcement; discuss direction externally
 - b. **OPTION 2**: Invest in u13 classifier, reassess fit (for app level, feature level or ad/content enforcement) once there's data on quality; discuss direction externally but minimize expectations
 - c. **OPTION 3**: Do not invest in u13 classifier; define a milestone to revisit this decision
3. Identify other approaches to age enforcement we should explore

Next steps: If we fail to reach a decision, the issue will go to Mark, potentially on 3/5

Out of scope for today: Teens

1. Recap the issue of doing more on u13s

WHY WE'RE TALKING ABOUT THIS

We're seeing increased external pressure (here)– predominantly from policy and sales audiences – around our approach to managing age on the platform. While the specific consequences of this pressure are varied, there are two areas of focus:

- **Under 13s**. We've set a minimum age for FB and IG, but it's perceived that we are not taking sufficient steps to enforce these rules or meet obligations to keep under 13s off FB more generally, especially given the general understanding that an unknown number of u13s are on FB and especially IG, and lying about their age.
- **Teens**. Once people are old enough to be on FB/IG, we're not viewed as taking sufficient steps to create a protective experience for them. We've particularly faced pressure around ads and other content restrictions and setting appropriate default privacy settings.

These concerns became more acute after we announced we were moving to end-to-end encryption in messaging. Satisfying critics that we have methods of protecting u18s even after messages encrypted is critical to landing MIP/e2ee.

The goal of today's discussion is to make decisions about our approach to Under 13s, but we will need to evaluate Teens as a follow on to this discussion.

OUR CURRENT APPROACH

We have a policy that says you have to be 13 to use our services,* recognizing that we haven't built FB/IG to be suitable for young users (and that US and other privacy law restricts our ability to collect data about u13s without parental permission).

- We ask people to state their age at registration and block attempts to create accounts by people who say they're too young. (But the default age is 25 and many people accept that.)
- People can report accounts as under13 ("u13"), which leads to a manual review by CO to see if the person appears to be underage. If so, we require them to produce ID to validate that they are over 13 before they can continue to use their account. We don't otherwise proactively enforce our minimum age policy.

* except WA, which requires age 16 in Europe, and Messenger Kids, which has no age requirement

THE "PEOPLE PROBLEM" WE'RE TRYING TO SOLVE

This policy creates problems for two groups of people:

Group 1: Age liars. People under 13 may want to use FB despite our policy.

- It's relatively easy to get around our age gate at registration, either by accepting the default (25) or simply choosing an age that's over 13.
- If a minor takes no action -- i.e., accepts the default -- then they will pass the age gate, and currently we aren't doing additional verification for people in this situation.
- If there is a catastrophic event involving someone in this group, it's likely that our current approach won't be viewed as sufficient.

Group 2: Adults with low digital literacy. People over 13 may want to (and be permitted to) use FB but may have trouble creating accounts because they're having trouble with our age collection interface.

To address the concerns of Group 1, we've been testing improvements to collecting age at registration that would collect age in a neutral way (i.e., one that doesn't have a "passing" age as the default and that doesn't clearly signal what the "right answer" to the age question is).

- This approach would be consistent with guidance from the FTC and children's privacy and safety experts, which focus on using what they call a "neutral age gate" at registration.
- In testing this approach, we found that a neutral age gate had a negative impact of ~12.5M MAP/year. We don't know how many of those are people in Group 2 (i.e., people who are allowed to be on FB).
- We need to decide whether we should follow this guidance, take an approach that will be considered substandard, or if there is a different approach that could help us meet the needs of Group 1.

THE POLICY PROBLEM WE'RE TRYING TO SOLVE

While neutral age collection at reg historically was viewed as sufficient, this is no longer true as policy makers are beginning to adopt new requirements around age and some companies are beginning to take more proactive steps, like the use of AI, to identify people in Group 1 (who are lying about their age).

- Over 2020/2021, we expect a number of jurisdictions – most significantly including Europe, the UK, and India – to adopt new regulatory obligations around the concern that platforms like ours are not doing enough to enforce age proactively. [More detail on the external context is [here](#).]
- These proposals will likely include requirements to do things like:
 - Verify age, potentially including a requirement to collect or verify identity before account creation
 - Set all minors to the most conservative privacy settings possible, or restrict all users to those conservative privacy settings if we don't verify age in an "appropriate" way
 - Get parental consent before collecting data from teenagers, potentially increasing the age of consent to 16 or 18
 - Restrict the ads and content that we show to people under 18
 - Get explicit consent before collecting certain kinds of data deemed "sensitive" from under-18s
- These measures are being finalized now, but because we're viewed as not doing enough about the Group 1 problem we don't have sufficient credibility to influence them. Being able to say that proposals policy makers are putting on the table aren't workable, but that we are working on other measures that might be, will win us a seat at the table to influence these outcomes.
- If we do nothing, it will be harder to influence the rules that will be adopted. But it's likely we'll still be among the first companies that face enforcement when these obligations come into effect. We also may be held to a different standard than domestic businesses in certain countries.

PAST DECISIONS

- We should change default age at registration from "today's date - 25" to "<13", so people have to specifically state an age to pass the age gate (10/22/19 Cross Family, 2/28/20 Central Privacy)
- While we generally should be consistent in how we check for u13s globally, we can test alternative "fallback" interfaces if people repeatedly provide age <5.
 - We should not test non-neutral fallbacks that clearly state or imply the "correct answer" (2/28 Central Privacy)
 - Revisit this decision if we can show that (a) actual minors are generally getting blocked correctly, and (b) low digital literacy is the main cause of people failing the age gate

EXISTING WORKSTREAMS

Since early 2019, we have been working on improving age collection and age modeling and made progress on both. We've also explored building a classifier for u13 age enforcement but been blocked on whether to proceed.


- Workstream 1: Improving Age Collection at Reg flows globally
 - Started collecting age on IG; agreed on general FB approach, but escalating decision on certain options for low digital literacy populations
- Workstream 2: Consolidating Age Data Sets across the company and ensuring Proper Use
 - Consolidated age models, implemented new governance frameworks, improved FB age models and are actively working on IG. Operational adult classifier in a few locales.
- Workstream 3: Improving u13 Age Enforcement
 - Developed a policy proposal, but haven't made progress on executing due to disagreements on approach and priority.
 - Decision from this group needed to move forward.

INDUSTRY BENCHMARKING

Most industry peers, competitors, and challengers focus on collecting age at reg, but we are seeing a growing trend toward exploring the use of AI for more robust age verification.

NOT AICE

Most peers focus on collecting age at reg



https://www.youtube.com/watch?v=... as of 4/14/20

2. Make decisions

a. Improving age collection at reg flows globally

Question 1: Should we change the current default age on the birthday screen from 25 years old to an age below the minimum required age (i.e. 0 or 1)? (Previously decided yes, but looking for final confirmation given results)

- Materials here

b. Exploring enhancements to age enforcement

Question 2: Should we invest in u13 classifier to be used for app level enforcement?

	A	B	C	D
1		Option 1	Option 2	Option 3
2	Description	Invest in u13 classifier, commit to using it for app level enforcement; discuss direction externally	Invest in u13 classifier, reassess fit once there's data on quality; discuss direction externally but minimize expectations	Do not invest in u13 classifier; define a milestone to revisit this decision
3	1. Safety	Highly Positive	Positive (if classifier used)	Negative
4	2. Privacy and Policy Makers	Highly positive	Positive (if classifier used)	Negative
5	3. Legal Profile	PRIV		
6	4. Reputation	Positive	Positive (if classifier used)	Negative
7	5. Advertisers	Highly positive	Positive (if classifier used)	Neutral
8	6. Business	Unknown MAP impact	Neutral	Neutral
9	7. Incremental investment	Highest	Medium	Not required
10	8. User experience	Negative	Depends on use cases	Neutral
11	9. Industry "first mover"	No alignment on what's best for the company		
12	Policy recommendation with explanation	Preferred approach. Enables policy to engage quickly to shape the many existing age proposals into something we can live with. Reduces point of policy criticism. Improves positioning for e2ee meaningfully.	Policy acceptable if we ultimately use the classifier. Being more tentative will limit our ability to shape the regulations being actively debated, but saying we're working on this gives us a seat at the table and puts us in a stronger position to push our preferred outcome once we reach alignment.	Not recommended. We'll remain at risk of criticism that we aren't taking u13s seriously enough, which are intensifying because of our e2ee announcement. We would have limited ability to influence debates about how to protect minors but would almost certainly face new compliance obligations in the next 1-3 years. High likelihood we're an early target of enforcement and may be held to higher standard than other companies.

Notes

- under Options 1 and 2, classifier would be used to queue accounts for our existing CO underage review; details here
- **PRIV**
- we'll have to architect/build to adhere to our public statements about the use of the classifier, e.g., if we commit to only using it for enforcement, we won't use it for Analytics, Data Science, Investor Relations, other Integrity, Growth

Criteria details

1. Safety – does it help with keeping u13s safe?
2. Privacy and Policy Makers – how much does it help us meet Privacy expectations and shape regulatory debate?
3. Legal Profile

PRIV
4. Reputation – how does it impact our exposure to high risk low probability event?
5. Advertisers – how does it help in addressing select advertiser concerns?
6. Business – do we commit to unknown future MAP hit?
7. Incremental investment – how high is the incremental investment required?
8. User experience – what impact does it have on user experience, e.g., through enforcement?
9. Industry “first mover” – do we want to lead the industry on at-scale u13 enforcement?
(Note – people haven't explicitly aligned on what the “better” answer is here)

3. Identify other approaches to age enforcement we should explore

Discussion question: Independent of whether we build a u13 classifier, should we focus on other alternatives, either in internal work or our external advocacy? If so, what?

Some ideas:

1. **Browser/OS/ISP signal.** Put the burden on browsers, operating systems, and ISPs, rather than FB/IG, to signal when a user is u13.
 - a. **PROS:** Imposes same level of obligation on everyone, not just us; shifts burden to actors that have more data/ability to address it at scale across services; more consistent experience for people across services
 - b. **CONS:** Gives Google/Apple significant leverage over us; we lose the ability to define level of confidence or mitigation steps (Google/Apple will likely require specific restrictions if signal is sent); implications for Oculus
2. **Identify verification.** Rather than inferring age, support identity verification either at reg or for a subset of users (ex., people with stated age 13-18 in US/UK/India)
 - a. **PROS:** Higher confidence that we're blocking people who are actually underage; relying on a third-party addresses data minimization concerns (but increases pressure to oversee data practices of third party); aligns with expectations from child advocates
 - b. **CONS:** Significant growth cost if obligation is imposed at registration; digital divide issues (teens and people in developing countries are less likely than adults in developed countries to be able to supply identity documents); difficult to limit spread (i.e., if we propose ID collection for subset, a regulator could require it at reg); data minimization/advocacy concerns
3. **Parental consent.** Allow parents to verify/authorize u18 accounts – and potentially manage them.

- a. **PROS:** Builds on approach we use in Messenger Kids and GDPR; leverages existing expectations around parental consent from COPPA in the US.
 - b. **CONS:** Parental consent mechanism would need to be developed; doesn't address underlying concerns about verification (since if this is all we do we won't have verified the adult's age); reinforces efforts in US/BR/IN to raise "age of consent."
 4. **Build more u18 and u13 services.** Rather than focusing on getting kids off FB/IG, build a distinct u18/u13 experience.
 - a. **PROS:** Greatest ability to build a product that aligns with needs of u18s/u13s; possibility to build in partnership with (and get support from) outside experts, similar to Messenger Kids.
 - b. **CONS:** does not fully address underlying concerns; costly.
 5. **Other?**
-

Reference Materials

- Policy: Age Enforcement Proposal [\[here\]](#)
 - External Landscape Overview [\[here\]](#)
 - Age Management Post Mortem [\[here\]](#)
 - Age Collection (Fallback Screen) at Reg Test Proposal [\[here\]](#) and Escalation [\[here\]](#)
-

Appendix

- Summary of Current State
- [WIP] Timeline and Scoping

Summary of Current State

Our current Age Management approach consists of:

- Requiring users to be 13+ to use all services (and 16+ to use WA in EU) and 18+ to access select products (eg dating tab) and ads (eg alcohol)
- Requesting age at reg on 1) Facebook and IG (DoB) 2) WhatsApp in EU under GDPR (16+ checkbox)
- Not proactively enforcing, except for when underage users are community reported

Age Management was presented as a top Policy priority in early 2019 following concerns from external stakeholders who believed there're many underage users on our services, our responsibility to keep kids safe was growing, and we were not doing enough about it (highest pressure in UK, US, India, Brazil). The goal was to get ahead of growing regulatory pressure.

Externally, most companies collect age at registration (except for Twitter and Reddit) and don't do age verification and age appropriate experiences based on predicted age for 13-18s once they are on the platform. There are a few notable exceptions: Snapchat is using inferred age for restricting vice ads to minors, YouTube is using machine learning to detect child directed content and adjusting the in-app experience appropriately, and TikTok is developing enforcement based on predicted age of the people in videos.

Since then, we have been working on 3 workstreams:

- Workstream 1: Improving Age Collection at Reg flows globally
- Workstream 2: Consolidating Age Data Sets across the company and ensuring Proper Use
- Workstream 3: Exploring enhancements to Age Enforcement

Workstream 1: Improving Age Collection at Reg flows globally

- We started collecting self reported age on IG globally in Dec 2019
- We are testing new reg flows on FB/MSGR globally (<100% neutral age gate on age fallback screen at Reg), but haven't aligned on the best approach for low digital literacy markets - as is defaults to today's date minus 25 years
- There is no on-going discussion on WA age collection at reg

Workstream 2: Consolidating Age Data Sets across the company and ensuring Proper Use

- We have consolidated u18 and 18+ age models from across the company
- We have built new sources for age predictions
 - Binary classifiers of Adults & Minors (13-17, 18+)
 - Cohort predictions, e.g., adults 18-22
 - Pinpoint age predictions per user
- We have designed a governance framework for using age models on a per use case/product team basis
- FB age models can now be used for analytics globally and passive gating (e.g., dating tab) in US/Can - more countries to be added in H1'20 (details here). IG models work is in progress and high pri for H1'20
- There's currently no u13 model at the company; we can't accurately estimate precision/recall given the COPPA standard for actual knowledge towards the training set and the classifier output

Workstream 3: Exploring enhancements to Age Enforcement

We have developed a Policy proposal for enforcement based on inferred age (details here). However, making further progress on using age enforcement presents several challenges.

The arguments supporting age enforcement are:

- **Safety:** increasing our ability to keep u13 off decreases their potential exposure to harm while using our services
- **Reputation:** being ahead of the industry on the topic; being prepared for a low prevalence/high risk event; minimizing external criticism if we were to only be using age prediction for business (e.g. DOB for payments or predicting age on unmatched account for Ads targeting)
- **Regulatory:** waiting till pushed or regulated wouldn't let us control the narrative
- **Privacy control:** Privacy advocates argue for inferred age enforcement to reduce the risk of us inadvertently possessing u13 data
- **Product:** improving our u13 enforcement will help us better land e2ee and sensitive products like Dating and Messenger Kids

The risks are:

- Legal:

PRIV

- **Privacy data minimization:** age enforcement would go against the principle (not collecting or developing more data than needed to operate our service)
- **Business:** it's unclear what MAP we are losing until we commit to losing it - as we will only be able to have clearer estimates of u13s on our services after we have the classifier (for compliance reasons)
- **Precedent:** execution in even one app/market will set precedent, putting us in strategic industry disadvantage, as well as opening the door for more pressure
- **Lack of alignment on "leading the industry":** there's no peer company doing u13 app level enforcement based on inferred age at scale, yet. People disagree on leading vs following . (TikTok is developing tech to do this and Snap is using inferred age for ads)
- **Lack of capacity and alignment on priority:** this work will have to be staffed and prioritized for execution across Product, Policy, Legal, Ops

Policy thinks Workstream 3 should be a high priority within the Account/Actor Problems section of the Integrity priorities: ■ Integrity 2019-H1/H2 Review

Age Models as of 11/26/2019

Model		Output	Maturity	Use Case
Known External Usage				
FB	FB Adult Classifier	Adult or Minor	Operational	Product - Dating (US/CA)
FB	"Single Factor"	Adult or Minor	Basic	Product - Dating (ROW)
FB	FB Age Affinity	Single-year	Basic	Integrity - PYMK, IIC
IG	IG Age Affinity	Adult or Minor	Basic	Integrity - IIC
IG	IG Ads: Age	Single-year	Basic	Ads (unmatched users)
Known Internal Usage				
FB	Age Affinity	Single-year	Basic	Analytics
FB	Adult Classifier	Adult or Minor	Basic	Analytics
FB	Lifestage	Cohort	Basic	Analytics
IG	Age Affinity	Single-year	Basic	Analytics
IG	Adult Classifier	Adult or Minor	Basic	Analytics

[WIP] Timeline and Scoping

TIMELINE

- **Dec 2019** - Build Age Enforcement Policy and Proposed Approach (draft), tee up decisions
- **Feb 2020** - Align on Age Enforcement approach and 2020 plan across the company
- **Mar - May 2020**
 - Align on the data set approach
 - Brief regulators globally (FTC, IDPC)
 - Secure advocates
 - Finalize Age Enforcement Policy
 - Staff the teams (Product, Ops, Policy, Legal)
- **Jun - Aug 2020** - Identify the data set, align on precision/recall, gather the data

- Sep - Nov 2020
 - Start training the u13 classifier, improve precision/recall to ready to launch level
 - Start building the checkpoint flow
- Dec 2020 - Run tests, get approvals to launch
- Jan 2021
 - Launch a u13 classifier
 - Launch the checkpoint flow (assuming we can leverage our existing ID review based flow; a parental consent flow would take longer unless we build in parallel)

BUILDING A U13 CLASSIFIER

Total resources: 2 Eng, 1 DS, 1 PM, 1 Legal, 1.5 Policy, 1 CDS, 0.5 DE

Time: 9-12 months

1. Identifying a legal-approved dataset, align on precision/recall, gather the data: 10-20 weeks
2. Start training a classifier, have early estimates of precision/recall: 6-8 weeks
3. Get more data, improve precision/recall to ready to launch level: 4 weeks
 - a. Get approvals to test: 2-4 weeks
4. Run enforcement test, check precision/recall: 4 weeks
5. Get approval to launch: 2-8 weeks

Key considerations and decisions:

- Execution risk, irrespective of investment - depends on quality and quantity of ground truth
- Aligning on the data set approach (internal, external tbd)
- Precision/recall threshold for enforcement given legal and business considerations needs to be defined

BUILDING THE ENFORCEMENT/CHECKPOINT EXPERIENCE (ASSUMING WE USE AN ID REVIEW BASED FLOW)

Total resources: 1 PM (Age verification and recovery), 10 Eng (User experience and Back end), 1+ Legal (depending on number of countries), TBD Ops

Time: 6+ months (if we can reuse ID verification flows; TBD if we need to build new flows involving parental consent, etc)

1. Investigate and align on the best flow option (2-3 weeks)
2. Develop the early flow, test, get early estimates for recovery rate (4-6 weeks)
3. Iterate, refine, test (6-8 weeks)
4. Finalize, get approvals to launch (2-4 weeks)

Key considerations and decisions:

- Recovery rate needs to be factored in the decision - likely <10% to start
- Time recoverable once enrolled needs to be defined

TEAM AND STAKEHOLDERS

Policy: Nick Clegg, Joel Kaplan, Rob Sherman (Lead), [REDACTED], Antigone Davis, [REDACTED]

Legal: [REDACTED], [REDACTED], [REDACTED], [REDACTED], [REDACTED]

Product Integrity: Guy Rosen, [REDACTED]

Product Growth: [REDACTED], Andrew Bocking

Product CP: [REDACTED], [REDACTED] [REDACTED]
Product DS: Alex Schultz
Privacy Program: [REDACTED]

For alignment: CP Leads, App Leads, Central Privacy Leads

Related documents

Cross-Family Review (Oct'19) [here](#)
Age Management: PM Ownership (May'19) [here](#)
Global Policy Priorities (Apr'19) [here](#)
Age Strategic Plan (Policy) (Mar'19) [here](#)
Other related documents and announcements [here](#)

Notes

Risk Assessment Buckets:

- Bucket #1: Bucket of things we agree we should change
- Bucket #2: Bucket of things we would change under pressure
- Bucket #3: Bucket of things we would defend vehemently

POLICY PROPOSAL SUMMARY

Key Principles

1. We should have protections for u18 and u13 users, and those protections should not be easy to circumvent
2. If a signal is used to determine age for business purposes, it should be used to enforce on age too
3. Our enforcement should generally be consistent across (a) all apps for a single, known person, and (b) globally

Proposed protocol for u13 and 13-17 enforcement

- For FB/MSGR and IG accounts, enforce on age based on the following signals:
 - Community reporting (*current state*)
 - Profile age change (*new, launched in Feb'20 for FB, proposed for IG*)
 - Age modeling (*new, proposed*)
 - Hard-matched to an IG/FB account we enforced on (*new, proposed*)
- If at least one of the signals above indicates the user might be u13, validate through Profile human review, and if the suspicion is confirmed, checkpoint
 - The checkpoint and appeal is 28 days long requiring an ID to pass, with the account suspended until resolution; if no resolution or user confirmed to be u13, delete data in 30 days max after account is disabled
- If at least one signal indicates 13-17, feature-level enforce by gating adult-only products and features, (e.g., user does not have access to Dating tab)
 - Where users will be age-gated from features they currently have access to, provide notice and point them to appeals process (to be developed by product)

Note: There are two future decisions that will need to be made in the next year on soft-matched accounts and user transparency of these predictions. These are not for discussion today.

Global roll out recommendation

- Start rolling out u13 age enforcement proposal on FB and IG in a few selected countries in 2020, and expand globally in 2021
- The proactive feature-level enforcement for 13-17 should roll out globally in 2020, with risk-based prioritization for use of predicted age to age-gate specific features
- We will strive to stay consistent globally where possible, but also explore whether variations are justified based on user need in a specific country