Assignment 3: Meta's Social Media Data and Well-Being Research
Law & Algorithms, Spring 2024
due before class on March 28

All of this assignment is based on true information.

1. Meta's Rocky Relationship with Outside Research

Meta, the parent company of Facebook, has had a long and complicated history with academic research. One of its early research partnerships with scholars at Cornell emotionally manipulated users intentionally without their knowledge, raising a great deal of ethical concerns. A study from 2010 about how the platform could alter voting results by selectively encouraging users to vote led to concerns that Facebook could secretly swing an election.

And perhaps most famously, its "Graph API" data portal for on-platform apps and permissive terms of service were leveraged by Cambridge University psychometric researcher Aleksandr Kogan to create a data gathering research application called "This is Your Digital Life." That app was used to harvest profiles of over 87 million Facebook users, most of whom did not use the app but had friends who did. (For certain app developers like Kogan at the time, that was all that was needed to see that secondary user's data through the Graph API.) That data was then licensed to a political consulting firm, Cambridge Analytica, who used the data in service of a series of conservative political campaigns in the US and UK. All of this came to light in 2018 care of a whistleblower inside of the firm, to considerable and sustained public uproar.

That led to the FTC imposing a $5 billion fine — the largest fine in the Commission's history — against Facebook. (It is also believed to be a significant part of why Facebook chose to rebrand as "Meta.") The legal ability for the FTC to seek such a fine was thanks in part to its prior consent decree against Facebook in 2012, where the FTC alleged that its privacy policy was deceptive in part because it failed to address this same type of "friend-of-a-friend" sharing to Facebook apps.

Shortly after the fine, the FTC modified its original 2012 consent decree and ordered the company to place some additional restrictions on "Covered Information," which the consent decree defines as:

> information from or about an individual consumer including, but not limited to: (a) a first or last name; (b) geolocation information sufficient to identify a street name

and name of city or town; (c) an email address or other online contact information, such as an instant messaging User identifier or a screen name; (d) a mobile or other telephone number; (e) photos and videos; (f) Internet Protocol ("IP") address, User ID, or other persistent identifier that can be used to recognize a User over time and across different devices, websites or online services; (g) a Social Security number; (h) a driver's license or other government issued identification number; (i) financial account number; (j) credit or debit information; (k) date of birth; (l) biometric information; (m) any information combined with any of (a) through (l) above; or (n) Nonpublic User Information.[1]

The revised consent decree, among many other things, requires Meta to:

- Ensure that "Covered Information cannot be accessed by any Covered Third Party[2] from servers under Respondent's control" after a user deletes that information or their account, except as otherwise required by law. (2020 Consent Decree, Part III.)

- Implement a comprehensive privacy program to govern Covered Information, including safeguards to ensure that every Covered Third Party that received Covered Information self-certifies annually that they are in compliance with platform terms and for what purpose they will use Covered Information. Facebook must terminate access for those out of compliance, monitor Covered Third Parties for abuse, and enforce its terms and related laws against Covered Third Parties "based solely on the severity, nature, and impact of the violation." (2020 Consent Decree, Part VII(E).)

- Notify the FTC when they have a data breach. More specifically, the decree defines a "Covered Incident" as "any instance in which Respondent has verified or otherwise confirmed that the Covered Information of 500 or more Users was or was likely to have been accessed, collected, used, or shared by a Covered Third Party in violation of Respondent's Platform Terms." (2020 Consent Decree, Part IX.) Notably,

[1] "Nonpublic Under Information," in turn, is defined as "any User profile information (i.e., information that a User adds to or is listed on a User's Facebook profile), or User-generated content (e.g., status updates, photos), that is restricted by one or more Privacy Setting(s)". "Privacy Setting(s)" are defined as "any control or setting provided by Respondent that allows a User to restrict which individuals or entities can access or view Covered Information."

[2] A "Covered Third Party" is "any individual or entity that uses or receives Covered Information obtained by or on behalf of Respondent outside of a User-initiated transfer of Covered Information as part of a data portability protocol or standard," with exceptions for Facebook's service providers (provided they only use the information as Facebook directs and for no other purpose); and as others need to use that information to comply with the law, enforce Facebook's terms of use, or mitigate fraud or security vulnerabilities on the platform.

> Facebook's current terms widely prohibit users from "collect[ing] data from our Products using automated means."

Months later, Meta found itself in even more researcher-related trouble. This time, the issue was the NYU Ad Observatory, a research project based at NYU's Cybersecurity for Democracy group that focuses on how political ads were being shown across Facebook. This library is built in part on data contributed by research subjects who volunteer to install the group's "AdObserver" plugin. The plugin detects political ads the user sees on their Facebook feed, and sends copies of those ads to the Ad Observatory. To help build trust in the platform the group posted the source code for the plugin on GitHub. (Subsequent to some of what follows below, it also commissioned Mozilla to conduct an independent privacy assessment of the tool.) Given what had happened with Cambridge Analytica in the 2016 presidential election, there was considerable public concern about similar influence in the 2020 election. Facebook separately ran its own political ad library for researchers, but it was widely panned as being buggy and ineffective.

On October 16, 2020 — after presidential voting was underway in several states — Facebook sent NYU a cease and desist letter ordering it to disable AdObserver and stop collecting data on the platform. The letter (as reported by the *Wall Street Journal*) alleged that NYU was impermissibly scraping its platform without consent. This led to outcry from numerous nationally-leading advocacy and technology civil liberties organizations, and one BU Law professor.

The researchers, represented by the Knight First Amendment Institute, negotiated for a better part of a year trying to resolve this tension between privacy and accountability, but in August 2021 Meta suspended the accounts of the individual researchers and the pages related to the research project. In a blog post, Meta said that "[w]e took these actions to stop unauthorized scraping and protect people's privacy in line with our privacy program under the FTC Order." The post noted that NYU's extension "also collected data about Facebook users who did not install it or consent to the collection," but it appears that the "users" referenced here are the political advertisers, not personal accounts.

While the FTC Order was only mentioned obliquely, it was enough to attract the ire of the Commission. Two days after that blog post, the FTC Acting Director for Consumer Protection Samuel Levine sent an open letter to Meta CEO Mark Zuckerberg that excoriated the company for suggesting their actions against NYU were required by their consent decree. "Had you honored your commitment to contact us in advance, we would have pointed out that the consent decree does not bar Facebook from creating exceptions for good-faith research in the public interest." But as some commentators noted both as this dispute began in 2020, and as Facebook escalated its actions in 2021, it was not obvious

that Meta's reading of its consent decree was incorrect, or that there wasn't a possible risk of user harm from NYU's project.

2. Meta's Current Transparency Efforts

In an effort to provide insight into its platform while also balancing the many privacy concerns that have arisen from research-related incidents like Cambridge Analytica, Meta has now deployed a variety of privacy-protective tools, largely organized through Facebook Open Research & Transparency (FORT). Some of its more recent projects have included:

- Their 2020 release of the Facebook URLs Dataset in partnership with Social Science One, an initiative from Harvard's Institute for Quantitative Social Science seeking to build greater academic–industry research partnerships. The URL dataset uses noising techniques that ensure a degree of differential privacy, though Social Science One objected to this technique: "We think of differential privacy as a technological solution to a political problem, just as the organizational structure we proposed for this project is an innovation in constitutional design that solved a different political problem." Some research has been directed at the question of whether this use of differential privacy still allows for valid analysis.

- A new Meta Content Library and API for information related to non-individual accounts on Facebook and Instagram, including Facebook Pages, groups, and business accounts. Access is facilitated through the Inter-university Consortium for Political and Social Research (ICPSR) at the University of Michigan, and researcher access is just beginning to be processed as of the time of this writing.

- Meta's "Data For Good" project, which produces datasets for public interest organizations in a variety of contexts. It says that it processes "privacy-preserving data," but is not specific as to what this means.

On January 29, 2024,[3] Meta announced its most recent project in this area, a study on "topics related to well-being" done in partnership with the Center for Open Science (COS), a nonprofit organization dedicated to openness, integrity, and reproducibility of research funded by the Arnold Foundation. The press release notes that any data that will be shared using privacy-preserving computational techniques, though it does not specify how it will do this. The COS plans to invite researchers to submit proposals for research plans that discuss (1) what well-being related topics they plan to study, and (2) what specific social

[3] The very same day that President Freeman announced the AI Administrative Process Task Force you discussed in Assignment 2!

media data they seek from Meta. The Center, the researcher, and Meta will then "collaborate to implement privacy protective measures" that allow studies to be published publicly and reproducible by others. While this is not made explicit, it appears that Meta intends to publish new datasets with some privacy-protecting qualities in response to these prompts that can be used by the general public, much as it did for the Facebook URLs Dataset—though probably not the same size as that dataset, as that has over 10 trillion cell values!

While the call is not specific, one assumes that this attention to "well-being" is oriented toward concerns about harms that arise from social media use, including concerns about social media addiction, how content algorithms may expose users to harmful or abusive content, or how these platforms can be used to facilitate other harmful behavior and activity. The timing of the announcement is likely probative as to its intended scope; it came two days before Mark Zuckerberg testified at a U.S. Senate hearing about child safety online. This research might of course reveal some deeply incriminating things about how Meta has conducted its platforms or may instead show that harms here are overstated or attributable to other factors.

3. Assignment

For this assignment, we would like to hear your group's thoughts on how Meta, researchers, and the Center for Open Science should weigh the competing demands between (1) transparency and accountability of social media platforms for well-being harms, and (2) the privacy of users on the platform. We would like to hear your proposals for how the Center for Open Science should handle requests for user information in light of:

1. Facebook's obligations under the revised consent decree, its terms of service, and its privacy policy. You can assume that Facebook is allowed to make an exception to its terms of service for anything that it prohibits for its own interests, but not in a way that would violate either the consent decree or the privacy policy. The consent decree should be read in light of the FTC's caution about pretextual assertion to deter scrutiny of its platform. You can also assume that any downstream researcher would be a "Covered Third Party" per the terms of the consent decree. Whether the information disclosed is "Covered Information" will likely depend upon the privacy techniques employed. It is likely not possible for Meta to delete any information after it is included in a dataset released through this partnership, given how often those sets will be copied by others.

2. The needs of academic researchers, who enter 2024 with renewed concern about the role of social media as vector for both personal harm and potential election

interference. Professor Ethan Zuckerman has identified a "constellation of factors" that make scrutiny of online platforms especially challenging at this present moment, right as the United States (alongside the UK, Mexico, Taiwan, Pakistan, India, and several other countries) face leadership-changing elections. Tools like platform APIs that were used to study social media have either closed down or become prohibitively expensive for researchers. Twitter has gone as far as to sue platform researchers seeking to expose the rise of hate speech on the platform. "Trust and safety" teams have borne the brunt of the layoffs happening at most major platforms. And some existing online scrutiny tools — like Facebook's CrowdTangle – have been intentionally dismantled. Data sharing like this, especially if it can be expeditiously reviewed and prepared for dissemination, could have an important role in shaping public understanding at this critical moment.

3. The privacy expectations of users, who have suffered harms from research projects in the past, as well as from pretextual data sharing done in the guise of academic research. Your analysis and description of these harms should be informed by the readings we have had about how to define and think of the concept of "privacy."

You can focus your attention on data that may be shared from the Facebook platform specifically, as opposed to other Meta platforms (WhatsApp, Instagram, etc.). For a sense of what data Facebook has on users that could be put to use, review some of the research above, alongside this recent review from *The Markup*. If you are a Facebook user, you can also request a copy of the data Facebook has about you.

Your proposal should assume the structure and format of the data sharing will be as it is set forth in the Meta/COS press release, and should instead speak to the substance of how the anticipated review should be conducted. Specifically, we would like to hear from you about:

1. How COS should select the researchers who will be invited to submit Registered Reports proposals, what the proposals should say about what data they intend to use and how, and how COS should evaluate those proposals.

2. What privacy-preserving computational techniques you believe may be helpful in the course of this analysis, and how those should be developed. You do not have to provide specific computations, but you should identify which general categories of techniques should be employed, and the affordances they provide that are useful here.

3. How your proposal squares with Meta's legal obligations under the revised consent decree, and the values and concerns of the other stakeholders.

Describe how you reached these recommendations in light of the stakeholder concerns discussed above, and use the readings from this part of the class to defend your choices. Your proposal can be addressed to the Center for Open Science, and should be between 2000–2500 words, citations excluded.

Please email your team's submission as a .PDF file to the course instructors (sellars@bu.edu and varia@bu.edu) before our class (i.e., before 2:10 pm ET) on March 28.

4. Rubric

Please review section 10 of the syllabus for our expectations for team collaboration for this and other assignments. Our grading rubric for this assignment will assess:

- How effectively you identify and address the competing values and policy considerations that surround this potential data sharing project.
- How the decisions made in your proposal relate to those values, and how you defend the inherent choices on which values to prioritize and why.
- How the privacy-preserving computational techniques you identified would work to balance these policy considerations in ways that can expand the possible data use in a responsible manner while simultaneously protecting the legal and ethical privacy considerations.
- Your engagement with the assigned readings from Classes 7–9, as well as the classroom discussion.