Data&Society

Letter on Proposed Changes to the Common Rule

Re: Notice of Proposed Rulemaking in the Federal Register (Vol 80, No. 173)

Council for Big Data, Ethics and Society

December 29, 2015

Dear Colleagues,

This letter from members of the Council for Big Data, Ethics, and Society[1], a National Science Foundation-sponsored project, regards the proposed revisions to the Common Rule as discussed in the Notice of Proposed Rulemaking (NPRM) entered into the Federal Register (Vol 80, No. 173)[2] on September 8, 2015. We are collectively writing this public letter to express our concerns regarding the consequences of the proposed rules for the emerging field of data science. The ethics of data science research and practice have risen to public notice during the time that proposed revisions to the Common Rule have been discussed and drafted.[3] Research methods in a number of fields are rapidly changing in response to the capacity to analyze ever-larger datasets networked with other large datasets collected from markedly different contexts. Not surprisingly, researchers and practitioners are increasingly finding that these new methods of knowledge production raise ethical challenges that do

[1] http://bdes.datasociety.net/. This letter represents the views only of those members who have signed it.

[2] Department of Health and Human Services (2015) *Notice of Proposed Rulemaking: Federal Policy for the Protection of Human Subjects. Federal Register.* Available from: http://www.gpo.gov/fdsys/pkg/FR-2015-09-08/pdf/2015-21756.pdf (accessed 21 October 2015).

[3] Executive Office of the President (2014) *Big Data: Seizing Opportunities, Preserving Values.* The White House. Available from: https://www.whitehouse.gov/sites/default/files/docs/big_data_privacy_report_may_1_2014.pdf (accessed 10 November 2015).

not easily translate into the regulatory frameworks developed over the last several decades.[4,5]

We wish to express our view that any rules which include or exclude data science from federal ethics regulations should be based on sound research and reasoning about risks to human subjects and preservation of social justice, and achieve clarity about when and how ethics regulations should apply. The proposed revisions in the NPRM fall short of this in several regards.

Many of the regulatory changes in the NPRM touch on areas that are relevant to data science, such as informed consent rules for re-use of biospecimens, the concern about re-identification of whole genome sequences, and the standardization of data privacy protections. However, we anticipate that many other parties will comment on these specific proposals and thus have chosen to limit our recommendations to the areas of the NPRM most relevant to the practices of data science that are primarily outside of biomedicine: the exclusion of research using non-research and public datasets, and the exemption of research using private, de-identified data.

The Excluded category

In creating the new "excluded" category (§___.101(b)), the Department of Health and Human Services rightly sought to clarify and standardize what types of research pose such minor risk to subjects that they should not fall under the scope of IRB review. Researchers outside of the biomedical sciences—including those from social, legal, and computational fields on this Council—have struggled with inconsistent application of IRB regulations primarily calibrated to the methods and risks of biomedical research.[6] The excluded category will have a meaningful impact in relieving that burden, where appropriate.

However, we believe that the proposed regulations contain a fundamental oversight about increasingly common data-intensive research methods. Section §___.101(b)(2) defines research methods that are to be *excluded* from IRB oversight because "they are considered to be ***low-risk*** human subjects research", and therefore the marginal protections offered to research subjects are not worth the administrative burden involved in regulating them. In particular, we express concern about section §__.101(b)(2)(ii),

[4] Zwitter A (2014) Big Data ethics. *Big Data & Society* 1(2): 2053951714559253.
[5] boyd d and Crawford K (2012) Critical Questions for Big Data. *Information, Communication & Society* 15(5): 662–679.
[6] Committee on Revisions to the Common Rule for the Protection of, Board on Behavioral, Cognitive, and Sensory Sciences, Committee on National Statistics, et al. (2014) *Proposed Revisions to the Common Rule for the Protection of Human Subjects in the Behavioral and Social Sciences*. Available from: http://www.nap.edu/read/18614/chapter/1 (accessed 21 October 2015).

which excludes "Research Involving the Collection or Study of Information that has been or will be Collected". The NPRM's extended discussion of this section states that the exclusion covers:

- research involving the collection or study of information that has been or will be acquired solely for non-research activities, **or**
- was acquired for research studies other than the proposed research study when the sources are publicly available, **or**
- the information is recorded by the investigator in such a manner that human subjects cannot be identified, directly or through identifiers linked to the subjects, the investigator does not contact the subjects, and the investigator will not re-identify subjects or otherwise conduct an analysis that could lead to creating individually identifiable private information.

Despite being grouped together, these types of research pose very different risk profiles in the context of big data research methods. Additionally, different big data research projects using the same methods and meeting these qualifications for exclusion can have zero risk or enormous risk for the subjects of that research. Notably, these criteria for exclusion focus on the *status* of the dataset (e.g., is it public? does it already exist?), not the *content* of the dataset nor *what will be done* with the dataset, which are more accurate criteria for determining the risk profile of the proposed research.[7]

At a prior point in history, it was reasonable to assume that *publicness* of an existing dataset was an adequate proxy for informational risk. Because disclosure of the data itself was the relevant harm, it was safe to assume that if data was already publicly available then by definition any harm was already done. However, the power and peril of big data research methods is that large datasets can theoretically be correlated with any other large datasets to algorithmically discover patterns that can re-identify and/or impact individuals in ways that were previously impossible. The technological and mathematical innovations of big data have greatly increased the utility of non-contextual data—datasets now travel in unanticipated ways, and can be used for unpredictable purposes.[8] Combined with the fact that vastly more sensitive information about human lives is recorded and stored in datasets for non-research purposes that is publicly available (including for purchase), correlational research using non-contextual data can expose subjects to as-of-yet-unknowable informational risk within a largely unregulated field.[9]

Researchers and practitioners have found any number of surprising correlations that can disclose

[7] We note that there is relatively little empirical research quantifying such risk for any individual user.

[8] Polonetsky J, Tene O and Jerome J (2015) Beyond the Common Rule: Ethical Structures for Data Research in Non-Academic Settings. *Colorado Technology Law Journal* 13.

[9] Ohm P (2014) What Do the Rules Say About Data Analysis? In: Julia Lane, Victoria Stodden, Stefan Bender, et al. (eds), *Privacy, Big Data, and the Public Good: Frameworks for Engagement*, New York, NY: Cambridge University Press.

sensitive information about persons in public datasets.[10] For example, recently a Freedom of Information Act (FOIA) request allowed a programmer to receive the entire, anonymized dataset of the New York City Taxi Commission's trip records.[11] Such a dataset could be used for any number of highly useful research projects about civic planning and transportation infrastructure. However, from this dataset, adept data analysts were also able to determine the likely religion of certain cab drivers[12], determine which rides were taken by celebrities and how much they tipped[13], determine the likely identity of individuals frequenting strip clubs[14], and de-anonymize the names of drivers based on medallion numbers[15], which could then be correlated with other private details such as religion and income. Although this particular dataset turned out to be poorly anonymized, it demonstrated a well-established pattern that anonymization is not meaningfully protective or technically viable in common circumstances.[16]

Many publicly available datasets can include or be used to infer a person's geolocation history, health status, financial well-being, political orientation, sexual status, etc. Private data brokers can collect and sell a tremendous breadth of data that could or should be considered private in any colloquial sense. Yet the *a priori* exclusion of research using those and other public datasets as proposed by the NPRM has the result of eliminating any point in the regulatory process requiring assessment of the risk posed by the content or uses of those datasets[17].

It is no longer reasonable to claim that either the prior existence or the publicness of a dataset is a reasonable proxy for minimal informational risk posed by the data contained therein. We note that

[10] Such concerns have been discussed in Nissenbaum, H. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*, (Cambridge: MIT Press), and even prior to that in Nissenbaum H, "Protecting Privacy in an Information Age: The Problem of Privacy in Public," *Law and Philosophy*, 17: 559-596, 1998.

[11] http://www.andresmh.com/nyctaxitrips/

[12] Franceschi-Bicchierai L (2015) Finding Muslim NYC Cabbies in Trip Data. *Mashable*. Available from: http://mashable.com/2015/01/28/redditor-muslim-cab-drivers/#0_uMsT8dnPqP (accessed 6 November 2015).

[13] Tockar A. Riding with the Stars: Passenger Privacy in the NYC Taxicab Dataset. *Neustar Research*. Available from: http://research.neustar.biz/2014/09/15/riding-with-the-stars-passenger-privacy-in-the-nyc-taxicab-dataset/ (accessed 6 November 2015).

[14] *ibid.*

[15] Pandurangan V (2014) On Taxis and Rainbows : Lessons from NYC's improperly anonymized taxi logs. *Medium*. Available from: https://medium.com/@vijayp/of-taxis-and-rainbows-f6bc289679a1 (accessed 10 November 2015). NB: this particular breach of privacy is not strictly a matter of correlational research using auxiliary data, but rather was a matter of a poorly hashed dataset.

[16] Ohm P (2009) *Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization*. SSRN Scholarly Paper, Rochester, NY: Social Science Research Network. Available from: http://papers.ssrn.com/abstract=1450006 (accessed 13 November 2015); and Narayanan A and Shmatikov V (2009) De-anonymizing Social Networks. In: *2009 30th IEEE Symposium on Security and Privacy*, pp. 173–187; and Narayanan A and Shmatikov V (2008) Robust De-anonymization of Large Sparse Datasets. In: *IEEE Symposium on Security and Privacy*, 2008. SP 2008, pp. 111–125; and de Montjoye Y-A, Hidalgo CA, Verleysen M, et al. (2013) Unique in the Crowd: The privacy bounds of human mobility. *Scientific Reports* 3. Available from: http://www.nature.com/articles/srep01376 (accessed 13 November 2015); and Montjoye Y-A de, Radaelli L, Singh VK, et al. (2015) Unique in the shopping mall: On the reidentifiability of credit card metadata. *Science* 347(6221): 536–539.

[17] For discussion of these issues as related to court records, a major source of information for data brokers, see, Conley A, Datta S, Nissenbaum H, and Sharma D, "Sustaining Privacy and Open Justice in the Transition to Online Court Records: A Multidisciplinary Inquiry," *Maryland Law Review*, 71:3 (Summer 2012).

within the NPRM, *public* and *private* are used in a manner that leaves this regulatory gap open. The NPRM uses "public" to modify "datasets"—*public* describes the type of access to or availability of a dataset. In contrast, the NPRM uses "private" to modify "information" or "data"—*private* describes the expectation that a reasonable subject has about the relative availability of sensitive information. *Publicly available* datasets containing *private data* describes many of the sources most interesting to data researchers and practitioners, and are arguably most risky for subjects. Yet research using such datasets remain unaddressed by the proposed regulations because public datasets would be *a priori excluded* from the Common Rule based on the assumption of minimal risk. Should this be intentional, we ask that the final rules make it clear that these exclusions are not happening due to limited potential of harm.

The Exempted category

A similar dynamic informs our concern about aspects of the *exempted* category, which in the NPRM would involve even less oversight than is currently applied to exempt research. Specifically, we have reservations about section §__.104(e)(2), which is tailored to facilitate the secondary use of identifiable private information collected for non-research purposes, a practice which the NPRM recognizes as central to big data research.

> (e) The following categories of exempt human subjects research allow for the collection of sensitive information about human subjects, must not involve biospecimens, must be recorded as required in paragraph (c) of this section, and require application of standards for information and biospecimen protection provided in §__.105:
>
> > (2) Secondary research use of identifiable private information that has been or will be acquired for non-research purposes, if the following criteria are met:
> >
> > > (i) Prior notice has been given to the individuals to whom the identifiable private information pertains that such information may be used in research; and
> > > (ii) The identifiable private information is used only for purposes of the specific research for which the investigator or recipient entity requested access to the information.

The NPRM discussion states that IRBs often waive consent for such research under the current rules and claims that the prior notice requirement would ultimately provide greater (or at least more consistent) respect for persons. Additionally, the (as of yet unwritten) privacy rules at §__.105 should provide standardized privacy practices.[18]

[18] For challenges to the gatekeeping capacities of of informed consent due to big data practices, see. Barocas S and Nissenbaum H (2014) "Big Data's End Run Around Consent and Anonymity," In *Privacy, Big Data and the Public Good* (Eds.) Eds. Lane J, Stodden V, Bender S, and Nissenbaum H, Cambridge: Cambridge University Press.

As with our concerns about §__.101(b)(2)(ii) above, we feel that this exemption lumps together an overly broad range of risks based on assumptions about the status of a dataset, rather than its intended uses or contents. Identifiable private data collected via Internet services and made available for secondary research can represent a tremendous breadth of "prior notice" practices and be collected for purposes of the recipient entity (i.e., the service provider) that are opaque or invisible to the user of that service.[19,20,21] Indeed, these issues were central to the highly contentious case of the "Facebook emotional contagion study", the most high-profile data research ethics controversy to date.[22] Experts disagreed widely as to whether Facebook's prior notice met standards for respect for persons, particularly given widespread misunderstanding on the part of users about how Facebook algorithmically filters the content shown to users.[23,24,25,26] Additionally, it appeared that even Facebook was internally unclear about what prior notice was provided or required for them to publish experimental results from manipulating users' News Feeds, leading to an editorial expression of concern from the publishers.[27]

We suggest that "prior notice" alone is too vague a standard to meet the requirements of respect for persons. Instead, either in the text of the Common Rule or in policy advice issued separately, the Federal government should identify best practices for prior notice and hold researchers accountable for using such practices. We also wish to note that empirical and legal scholarship has identified mere mandated disclosure as a weak model for maintaining respect for persons in a research context.[28] The final rules should therefore encourage or accommodate notice practices that go beyond simple compliance with mandated disclosures.

Requiring that researchers restrict themselves to only the "specific research for which the investigator or recipient entity requested access to the information" is a potentially useful check against correlational research that risks sensitive disclosures. It is the nature of large datasets, particularly those

[19] Luger E and Rodden T (2013) Terms of Agreement: Rethinking Consent for Pervasive Computing. *Interacting with Computers* 25(3): 229–241.
[20] Chee FM, Taylor NT and de Castell S (2012) Re-Mediating Research Ethics End-User License Agreements in Online Games. *Bulletin of Science, Technology & Society* 32(6): 497–506.
[21] Ohm (2014)
[22] Kramer ADI, Guillory JE, and Hancock JT (2014) Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of Sciences* 111(24): 8788–8790.
[23] Meyer MN (2015) *Two Cheers for Corporate Experimentation: The A/B Illusion and the Virtues of Data-Driven Innovation*. SSRN Scholarly Paper, Rochester, NY: Social Science Research Network. Available from: http://papers.ssrn.com/abstract=2605132 (accessed 19 October 2015).
[24] boyd danah ((in press)) Untangling Research and Practice: What Facebook's 'Emotional Contagion' Study Teaches Us. *Research Ethics*.
[25] Crawford K (2014) The Test We Can—and Should—Run on Facebook. *The Atlantic*. Available from: http://www.theatlantic.com/technology/archive/2014/07/the-test-we-canand-shouldrun-on-facebook/373819/ (accessed 21 January 2015).
[26] Grimmelmann J (2015) *The Law and Ethics of Experiments on Social Media Users.* SSRN Scholarly Paper, Rochester, NY: Social Science Research Network. Available from: http://papers.ssrn.com/abstract=2604168 (accessed 19 October 2015).
[27] Verma I (2014) Editorial Expression of Concern and Correction. *Proceedings of the National Academy of Sciences* 111(29): 10779.
[28] Ben-Shahar O and Schneider CE (2014) *More Than You Wanted to Know: The Failure of Mandated Disclosure*. Princeton University Press.

generated by Internet services, to have many purposes, especially in relation to other datasets. Respect for persons may require that the use of data in research not just reflect the scope indicated by the data recipient, but also a reasonable accounting for the expectations of the subjects whose data is being shared. This caveat is significantly more important regarding data collected for non-research purposes, in which case we may reasonably infer that the research uses of the data are less obvious to the users of the Internet service.

Responses to requests for input

In question #17 the NPRM asks whether including research using public datasets under the Common Rule would meaningfully add to the protections for human subjects, whether the exclusion should be narrowed, and whether these activities should instead be classified as an exemption. Our members hold a range of opinions on these specific matters. However, we wish to express consensus that Excluded or Exempt status for public data sets *should not* be justified in terms of inherently low risk to human subjects. As discussed above, the publicness of a dataset is not a reasonable proxy for level of risk to human subjects, and the HHS will generate confusion for human subjects and researchers alike if the final rules assert this assumption as a sound justification for exclusion. Declaring that public datasets pose no risk *a priori* will subvert efforts to build robust models of accountability that are properly scaled to data science, whatever those models may ultimately look like.

That being said, some of our members believe that non-biomedical data science should largely fall outside the purview of the Common Rule due to a poor fit between the epistemic conditions of data science and assumptions about scientific practice built into the Common Rule. This is a notably different position than contending that data science should be Excluded because it poses inherently low risk, as expressed by the NPRM; rather, these members desire that the Exclusion be justified clearly in terms of the proper purview of the Common Rule. Other members believe that an Exempt status would be more appropriate for public datasets. At a minimum, this would allow the standardized Exemption-determination tool to inquire about risks imposed by correlative research and create some potential for IRBs to identify research that has empirically demonstrable higher risk for human subjects.

In question #52, the NPRM requests input on the importance of prior notice when determining the exempt status of research that uses identifiable private data procured for non-research purposes. We advise that the quality of prior notice is perhaps the more relevant criteria than the mere fact of it. “Prior notice” is a rather minimal bar if it includes notice hidden inside of End-User License Agreements (EULA), especially if the service provides no option to opt-out or audit the uses of one’s own data. Notice that only involves consent to an EULA is not meaningful protection. Requiring a higher quality of notice practices may provide meaningful additional protections without creating unfair burdens to researchers.

Conclusion

Although members of this Council have some differences of opinion about the best route to promoting ethical research practices in data science—including a range of views about whether the Common Rule ought to apply to data science at all—we wish to reiterate that it is critically important that any regulations be developed around sound empirical research and reasoning about what constitutes risk for human-subjects in the age of big data. That is a daunting challenge as the research literature is still sparse on that matter and so many key ethical infrastructures and core concepts are simultaneously called into question by the methods of big data. Nonetheless, the NPRM's proposed revisions addressed herein contain a mismatch between the criteria that signal risk to IRBs and the factors that can reasonably be assumed to modulate the actual risk to human subjects.

We submit these recommendations in hope that the language contained in the final rules can achieve more clarity on these matters and will successfully accommodate the epistemological and ethical challenges wrought by big data research and practice.

Sincerely,

The Council for Big Data, Ethics, and Society

Signed by:

Solon Barocas, Princeton University

Geoffrey C. Bowker, Evoke Laboratory, University of California, Irvine

danah boyd, Microsoft Research and Data & Society Research Institute

Kate Crawford, Microsoft Research, MIT Center for Civic Media, and NYU Information Law Institute

Alyssa A. Goodman, Harvard University

Rachelle Hollander, National Academy of Engineering

Barbara Koenig, University of California, San Francisco

Jacob Metcalf, Data & Society Research Institute

Arvind Narayanan, Princeton University

Helen Nissenbaum, New York University

Frank Pasquale, University of Maryland

Latanya Sweeney, Harvard University Data Privacy Lab

Matthew Zook, University of Kentucky

The Data & Society Research Institute Program on Ethics has investigated the potential benefits and challenges put forward in this letter. Through partnerships, collaboration, original research, and technology development, the program seeks cooperation across sectors to innovate and implement thoughtful, balanced, and evidence-based responses to our current and future data-centered issues.

Data & Society is a research institute in New York City that is focused on social, cultural, and ethical issues arising from data-centric technological development. To provide frameworks that can help address emergent tensions, D&S is committed to identifying issues at the intersection of technology and society, providing research that can ground public debates, and building a network of researchers and practitioners that can offer insight and direction. To advance public understanding of the issues, D&S brings together diverse constituencies, hosts events, does directed research, creates policy frameworks, and builds demonstration projects that grapple with the challenges and opportunities of a data-saturated world.

Contact

Jacob Metcalf
jake.metcalf@datasociety.net
Data & Society Research Institute
36 West 20th Street, 11th Floor New York, NY 10011
Tel. 646-832-2038
datasociety.net