

RAPID ETHICS REVIEW – UK Pandemic Ethics Accelerator

The ethics of data self-reporting: a review of important issues and best practices

Part II: best practice innovations

Niccolò Tempini

University of Exeter & Alan Turing Institute

n.tempini@exeter.ac.uk

This report commissioned by UK Pandemic Ethics Accelerator. Revision submitted 24th June 2022. Please cite as follows:

Tempini, N. 2022. The ethics of data self-reporting: a review of important issues and best practices – Part II: best practice innovations. UK Pandemic Ethics Accelerator (under review). Preprint available at <http://dx.doi.org/10.5281/zenodo.7324079>

Headline points

- The spectrum of self-reporting is extremely diverse and impossible to govern through a one size fits all approach; projects can be scientific or not, can raise important privacy concerns, and can be open to organised manipulation and poor governance of risks and harms; best practices and guidelines must be adapted to the local setting and remain open for improvement
- People take part in data collection for the most heterogeneous of purposes. Closely controlling hopes, aims and beliefs of participants is beyond the reach of any sensible self-reporting exercise; this heterogeneity introduces various biases in the data, and adds unpredictability to the data collection project
- Governance of self-report data collection is a thorny issue; new data governance approaches have been put forth to deal with the heterogeneity of ends that big data collection projects now need to be able to support, that require some experimentation, but the best chance of long-term sustainable governance comes from collective governance frameworks

- Greatest scientific potential of self-reporting is through data linkage with other forms of heterogeneous data, however these forms of data linkage are also the source of the highest risks for privacy and individual harms; flexible and sustained ethical oversight is key
- Participation in self-reporting is not equally distributed across society; when there are benefits associated in participation, it can exacerbate existing inequalities

The ethics of data self-reporting

Part II: best practice innovations

Introduction

Important ethical questions on data self-reporting have arisen both before and during the pandemic. The pandemic saw various kinds of self-reporting apps taken up by extraordinary amounts of contributors, and heightened our sensitivity to the issues that came forth then. But ‘normal times’ had already taught us much on issues of data governance, recognition and valuation that it is important to try to address as many of these issues at once to best prepare for the time after the pandemic. This report discusses some of the key ones, and highlights some of the most promising innovations to tackle them and to help with issues of public trust (Bharti et al. 2021).

The report is divided in two parts. This document, *Part II: best practice innovations*, takes stocks of innovations that can help improve best practices with a view to tackle challenges outlined in Part I. The previous document, *Part I: ethical issues*, outlines the ethical issues that are outstanding the pandemic experience and the decades preceding it in which web technology became mainstream.

1. Best practices

The spectrum of self-reporting is extremely diverse and as such impossible to govern through a one size fits all approach; projects can be scientific and not, can raise important privacy concerns, and can be open to organised manipulation and poor governance of risks and harms; best practices and guidelines must be adapted to the local setting and remain open for improvement. The following are three areas of best practice that have been developed to help manage the ethical issues generated by data self-reporting projects. They are applicable to many different domains and institutional settings. And they should be seen as a rather complementary set –conversely, their power is diminished without one another.

1.1 Data trusts and data cooperatives

Managers and developers of data self-reporting projects should consider what measures they can take to ensure data governance is accountable, inclusive, competent. Projects where data are extracted from a contributor base to then be shifted and shared according to the judgement of a self-selected few are more likely to enter controversy or make poor choices from an ethical point of view. These are the governance approaches privileged by private companies and tech monopolists who intend to control the way in which commercial value can be created by contributed data because they concentrate decision power in the hands of allegiant few.

Arguments in favour of individualised, distributed control of data that resonate with recent visions of decentralised web and organisational governance are problematic because of the burden and risks they put each individual deciding for themselves under. They imagine technological frameworks could make it possible for the data to be locked and unlocked by the individual they refer to, thus giving, in principle, total control over the data back to them. The real world might differ. Few have the expertise and time necessary to scrutinise proposal to use their data that might be put before them. Most are likely to take decisions that are against their best interests or principles because of poor or rushed judgement.

In recent years there has been a great deal of innovation in data governance frameworks that deliver a form of collective control on the projects, uses and aims the data are put to. These are the data trusts (Delacroix and Lawrence 2019) and data cooperative (Ernst Hafen 2019; E. Hafen, Kossmann, and Brand 2014) approaches. They have been applied with particular promise in domains such as health data, where data are sensitive, requiring tight scrutiny, and highly valuable and desired by researchers and business alike.

These approaches reject the assumption that each individual is capable and best positioned to decide what projects are worth giving access to their data. They instead elaborate on forms of delegation whereby a number of research participants are appointed as delegates or trustees with data governance responsibilities, and are charged with day to day decision making and governance. Depending on the setup, members can exercise their own decision over whether to participate in one project or another more or less often. These approaches can be used to take over and collectively govern and mobilise existing datasets that were already generated for other purposes, for instance, receiving healthcare; but can also be used to start and

coordinate the collection of new data from the outset, in a ‘bottom-up’ fashion where the project is managed by participants. They are intended to afford better inclusivity, accountability and competent data management. Different individuals should be able to seek different levels of involvement according to their interest.

Not all self-report collection projects need to be governed through a bottom-up approach such as data cooperatives or trusts and in many cases, data governance might be better managed through review boards composed of independent experts (see for instance, section 1.3) who might or might not be participating in the project themselves. The value of representative bottom-up data governance approaches is still in the potential for the project management to better articulate finalities, values and directions of a project in a way that better reflects the views and aims of its pool of participants, enacting the sort of relational ethics (Birhane 2021) where experiential expertise is valued at the same time as professional expertise

1.2 Solidarity and harm mitigation

Data governance includes the assessment of risks and benefits of ways of working with data to the best of knowledge and assumptions held by those involved in decision making. Assessing risks and benefits of different ways of using data necessarily involves making many assumptions about future developments and events. The exercise is uncertain, and it is possible even for the most suitable experts to make mistakes. This is true even when data are governed through a collective governance framework such as a data trust or cooperative. Harms and benefits will be very likely to distribute unevenly, and collective governance will not in itself make things right or fair if mistaken assessments lead to harm and benefits that are unfairly distributed in ways that could be mitigated.

As Barbara Prainsack argues (Prainsack 2019a; Prainsack 2019b), only when collective governance frameworks are combined with a commitment to the principle of social solidarity they can truly fulfil their revolutionary potential of alternative to the monopolistic extractive models pioneered by tech giants. To achieve this, data self-reporting projects need to accept the possibility of harm and prepare for its mitigation. With her colleagues (Prainsack and Buyx 2017; McMahon, Buyx, and Prainsack 2020) she calls for the creation of harm mitigation bodies that will operate harm mitigation measures. These, they suggest, could include issuance of apologies and stipulation of amendments to avoid future harm recurrence; and financial support for the gravest of cases.

Through a combination of collective governance (see 1.1) and harm mitigation features, self-reporting projects can aspire at the implementation of a *data donation* economy that maximises data's potential to be used for the common good.

Data donation has been used as a buzzword for broadly mimetic purposes by some exploitative projects, but it is not a fuzzy concept. There are a few conditions that must be met for a project to be based on data donation. Barbara Prainsack emphasises that for practices of data reporting and sharing to qualify as based on data donation, they need to exhibit *relationality*, *indirect reciprocity* and *multiplicity*. Relationality requires that the two parties (giver and receiver of data) mutually acknowledge the act of donation; this means that the data receiver (self-report data project manager) honours the work that the data donors (research participants) have done by “*systematically considering the needs and interests of data donors and their significant others*” (Prainsack 2019a:14); this can be reflected in the ways in which a project is organised to register, reflect and enact the views of its contributors. Relational ethics and bottom-up data governance, for instance, can go some way in this direction. But for the data donation economy to ensure that reciprocity is indirect and thus shared across the membership, harm mitigation measures are seen as a necessary instrument. It is through this combination of arrangements that data multiplicity (Prainsack's way of pointing out the ability for data to be reused multiple times) can serve the common good. Data that are collected, for instance, through mainstream, status quo, commercial data governance practices (for instance, data collected through a user-generated data platform and used by a Silicon Valley giant for undisclosed commercial purposes) might still have multiple uses, i.e. multiplicity, but the economy of such a project would not be based on donation, because the data are not received as a *gift* and respected with the ethics gifts command.

1.3 RECs and research ethics governance

A crucially important function that officers of a data trust, data cooperative or data governance organisation form should need to provide is the oversight of ethical risks. At any rate, a data self-reporting project should implement an ethical risk oversight function that is adequate to its size and ambition. A project that aims to create a data resource that can be used time and again and is widely trusted by internal and external stakeholders is highly likely to need a formal oversight structure with clear governance procedures.

The time-tested standard model in this domain are Research Ethics Committees that have been implemented for decades by universities in the UK, US (aka IRB), and the rest of the world. They normally require researchers who intend to embark on new research to submit a research ethics application. This is usually centered on a formal document (e.g. application form) that provides key information about the project aims, methods, resources; and evaluations of ethical risks together with a specification of measures planned to control for them. The application is reviewed by experts serving in the committee who should be able to judge whether the project is responsibly conceived, and the eventual ethical risks are acceptably approached, i.e. control measures are fit for purpose. If the application is rejected, the research cannot start.

Institutions can design, and have designed, research ethics oversight committees and processes differently to adapt to local conditions, but the above tend to be the main features, and this model has been adapted outside academia in the public and private sector as more organisations sought to control risks of their R&D activities (including reputational risk, through the accountability that having a formal process affords). This means that there have been no off-the-shelf approaches to research ethics oversight, though there have been various guidelines.

The need for up to date guidelines and models has recently become apparent as some late developments in research methodology are challenging the research ethics status quo. A report by the Ada Lovelace Institute, the University of Exeter Institute for Data Science and Artificial Intelligence, and the Alan Turing Institute (Petermann et al. 2022) highlights the new challenges and the possible solutions that institutions should consider to keep their research ethics oversight function suitable (also, among others, explored by Ferretti et al. 2021; Jordan; Basl and Sandler 2019; Clark et al. 2019). While for a long time RECs have been focusing on the ethical harms and risks that research could pose for the participants in a study (resulting from the origins of research ethics in medical research and the design of clinical experiments), new kinds of research employing cutting edge data science and artificial intelligence methods are bringing up emerging ethical issues. Here, the new innovations engendered by data self-reporting projects coordinated over the web are part parcel of these changes that need to be confronted.

Of growing concern have been the long term consequences of research that might be developing methods and technology that can be easily re-deployed (recall ‘function creep’) across institutions, jurisdictions, countries, populations, and how they might have disparately unequal social impacts; along with a redefinition of issues of privacy, confidentiality, and consent at a time when researchers can source individual level data from the web or other easily available resources, without asking the individuals these data originated from.

Issues are complex and their governance cannot be a box-ticking exercise. Methods that can be seen as gold standard in tackling some problems still require close and situated scrutiny. For instance, anonymisation can be a strong solution for granting privacy protection to individuals whose data widely available on the web is collected without their explicit informed consent. However, it is widely understood that the strength of anonymisation is always relative to the way in which various other data sources a researcher can have access to can lead to re-identification. Even when anonymisation is considered satisfactory for the purpose of privacy protection, in itself it offers nothing to protect individuals from the eventual unfair impacts of the outcomes of research once they are operationalised and deployed in new systems shaping individual movements, interactions, behaviours and access to services.

In recent years and in response to the emerging ethical challenges posed by the latest innovations in Data Science and Artificial Intelligence, there has been growing demand for the active consideration of broader downstream impacts of such research. Research communities and academic societies have started to require researchers to submit research ethics statements on the broad societal impacts of their endeavours (Petermann et al., 2022). Leading commercial technology organisations have started to experiment with the implementation of internal ethics review processes for their commercial R&D projects, in an effort to improve their reputation and accountability when the public is questioning whether they can be trusted with running the socio-technical infrastructure of much public life. How much can commercial organisations be expected to follow together ethical and business principles, and successfully map, and manage, the broad landscape of social groups, needs and interests, is very much an open question. Criticisms of ‘stakeholder capitalism’ abound. But it can perhaps be argued that when self-report data power so many projects and systems across both the research and commercial worlds, and underpin affairs in the public sphere; and when so much of public life relies on complex infrastructures developed on methods and

techniques of recent invention and never fully understood consequence; then, attempts at translating some of the ethical principles, expectations and methods developed to deal with the arising ethical issues of self-report data might be in order. The implementation of ethical oversight functions, which can take many forms, in new contexts is one such opportunity.

Questions around the re-definition of ethical standards and red lines are tightly linked to practical questions over the right training and composition of experts that REC should seek when evaluating projects with strong interdisciplinary components, whose risks are inherently more difficult to assess (Petermann et al. 2022); and complex projects involving multiple staggered undertakings where a technological solution builds from the previous. Questions that are open to debate and that will require careful consideration when implementing research ethics functions in organisations include also the scope, duration and frequency of interactions between research ethics functions and the researchers leading a project. These questions are not merely operational but really shape the quality of understanding and assessment that a research ethics function is able to deliver.

But at any rate, a form of research ethics oversight is needed in data self-reporting projects. Here again principles of bottom-up participation will be important. The world-leading public health databank SAIL (Ford et al. 2009), which has been managing dozens of highly sensitive, linked health datasets and making them available to hundreds of research projects, involves members of the public, together with domain experts, as reviewing members of research ethics applications.

Bottom-up participation will improve diversity of perspectives and sensitivity to assess bias and unequal impacts, and will be a required condition for a full implementation of collective governance approach involving data management through trust and cooperative; and design and operation of harm mitigation measures. Birhane (Birhane 2021) emphasises, with the concept of *relational ethics*, how those communities and groups that are on the receiving end of a project's impacts should be included in its governance, because they have a key epistemic privilege about its social outcomes – decisions should be taken with and through them, instead of without.

Acknowledgements

I thank Cian O'Donovan, James Wilson and Melanie Smallman for feedback and suggestions.

References

- Basl, John, and Ronald Sandler. 2019. *Building Data and AI Ethics Committees*. Ethics Institute, Northeastern University & Accenture.
- Bharti, Nishtha, Cian O'Donovan, Melanie Smallman, and James Wilson. 2021. 'Public Trust, Deliberative Engagement and Health Data Projects: Beyond Legal Provisions'. *Engaging Science, Technology, and Society* 7 (1): 125–133. doi:10.17351/ests2021.1197.
- Birhane, Abeba. 2021. 'Algorithmic Injustice: A Relational Ethics Approach'. *Patterns* 2 (2): 100205. doi:10.1016/j.patter.2021.100205.
- Clark, Karin, Matt Duckham, Marilys Guillemain, Assunta Hunter, Jodie McVernon, Christine O'Keefe, Cathy Pitkin, et al. 2019. 'Advancing the Ethical Use of Digital Data in Human Research: Challenges and Strategies to Promote Ethical Practice'. *Ethics and Information Technology* 21 (1): 59–73. doi:10.1007/s10676-018-9490-4.
- Delacroix, Sylvie, and Neil D Lawrence. 2019. 'Bottom-up Data Trusts: Disturbing the "One Size Fits All" Approach to Data Governance'. *International Data Privacy Law* 9 (4): 236–252. doi:10.1093/idpl/ipz014.
- Ferretti, Agata, Marcello Ienca, Mark Sheehan, Alessandro Blasimme, Edward S. Dove, Bobbie Farsides, Phoebe Friesen, et al. 2021. 'Ethics Review of Big Data Research: What Should Stay and What Should Be Reformed?' *BMC Medical Ethics* 22 (1): 51. doi:10.1186/s12910-021-00616-4.
- Ford, David V., Kerina H. Jones, Jean-Philippe Verplancke, Ronan A. Lyons, Gareth John, Ginevra Brown, Caroline J. Brooks, et al. 2009. 'The SAIL Databank: Building a National Architecture for e-Health Research and Evaluation'. *BMC Health Services Research* 9 (1): 157. doi:10.1186/1472-6963-9-157.
- Hafen, E., D. Kossmann, and A. Brand. 2014. 'Health Data Cooperatives – Citizen Empowerment'. *Methods of Information in Medicine* 53 (2): 82–86. doi:10.3414/ME13-02-0051.
- Hafen, Ernst. 2019. 'Personal Data Cooperatives – A New Data Governance Framework for Data Donations and Precision Health'. In *The Ethics of Medical Data Donation*, edited by Jenny Krutzinna and Luciano Floridi, 141–149. Philosophical Studies Series. Cham: Springer International Publishing. doi:10.1007/978-3-030-04363-6_9.
- Jordan, Sara R. *Designing an Artificial Intelligence Research Review Committee*. Future of Privacy Forum. Virginia Tech.
- McMahon, Aisling, Alena Buyx, and Barbara Prainsack. 2020. 'Big Data Governance Needs More Collective Responsibility: The Role of Harm Mitigation in the Governance of Data Use in Medicine and Beyond'. *Medical Law Review* 28 (1): 155–182. doi:10.1093/medlaw/fwz016.
- Petermann, Mylene, Andrew Strait, Niccolò Tempini, Ismael Kherroubi-Garcia, and Kirstie Whitaker. 2022. *Ethical Review Processes for AI and Data Science Research*. London and Exeter: Ada Lovelace Institute, University of Exeter Institute for Data Science and Artificial Intelligence, Alan Turing Institute.
- Prainsack, Barbara. 2019a. 'Data Donation: How to Resist the Leviathan'. In *The Ethics of Medical Data Donation*, edited by Jenny Krutzinna and Luciano Floridi, 9–22. Philosophical Studies Series. Cham: Springer International Publishing. doi:10.1007/978-3-030-04363-6_2.

- Prainsack, Barbara. 2019b. 'Logged out: Ownership, Exclusion and Public Value in the Digital Data and Information Commons'. *Big Data & Society* 6 (1): 2053951719829773. doi:10.1177/2053951719829773.
- Prainsack, Barbara, and Alena Buyx. 2017. *Solidarity in Biomedicine and Beyond*. Vol. 33. Cambridge: Cambridge University Press.
http://encore.exeter.ac.uk/iii/encore/record/C__Rb3514718.