

The Belmont Report in the Age of Big Data:

Ethics at the Intersection of Psychological Science and Data Science

To appear in S. E. Woo, L. Tay, & R. Proctor (Eds.), Big data methods for psychological research: New horizons and challenges. (American Psychological Association; anticipated publication date: 2020.) Published version may differ slightly from this accepted version.

Alexandra Paxton^{1,2}

¹Department of Psychological Sciences, University of Connecticut

²Center for the Ecological Study of Perception and Action, University of Connecticut

Corresponding author:

Alexandra Paxton
Department of Psychological Sciences
406 Babbidge Road, Unit 1020
Storrs, CT 06269
alexandra.paxton@uconn.edu

Acknowledgements

My thanks to Tom Griffiths for conversations about related issues during our work together on Data on the Mind; to Julia Blau for invaluable feedback on this chapter; to Aaron Culich for thoughtful discussions about securing computational pipelines; to audiences at U. Cincinnati's Center for Cognition, Action, & Perception Colloquium and U. Connecticut's Perceiving-Acting Workshop for insightful questions and comments during presentations of earlier versions of this chapter; to R. Stuart Geiger for sharing his thoughts on this work from a data ethics perspective; and to attendees and organizers of the 2018 Moore-Sloan Data Science and Data Science Leadership Summits (Park City, UT) for discussions about algorithmic justice.

The Belmont Report in the Age of Big Data:

Ethics at the Intersection of Psychological Science and Data Science

Forty years before the publication of this chapter, the U.S. National Institutes of Health released the Belmont Report (1979) to establish ethical guidelines for researchers working with human subjects. Since then, the Belmont Report has not only guided ethical principles but has also shaped federal policy for biomedical and psychological research (45 CFR 46, 2018). In many ways, psychological science today still strongly resembles psychological science from 40 years ago. Researchers are still captivated by understanding many of the same affective, behavioral, and cognitive phenomena. Participants still largely consist of undergraduate student volunteers. Research methods still include self-report surveys and painstaking observation, along with ever-improving dynamics-focused equipment like eye-trackers (Cornsweet & Crane, 1973) and accelerometers (Morris, 1973).

However, technological innovations over the intervening decades have opened doors that the authors of the Belmont Report likely never imagined. Today, humans generate quintillions of gigabytes of data every day (James, 2018). These digital traces of human activity are incredibly useful to private corporations and to government agencies, but they also hold immense promise for understanding psychological phenomena outside of the laboratory. This promise has drawn in pioneering researchers from psychology (e.g., Goldstone & Lupyan, 2016) to network science (e.g., Vespignani, 2009) in the hopes of tapping these data to reconstruct and predict the human behavioral, affective, and cognitive processes that generated them.

The increasing popularity of this approach—along with the increasing richness of the underlying data—have prompted increasingly pressing questions about ethics. While this new

frontier of data¹ presents unprecedented challenges to human-subjects ethics, I argue that the core principles of the Belmont Report are broad enough to encompass any medium of human-subjects research, whether in the lab or in the wild. After situating ethics of large-scale human-derived data use in a historical context, I will discuss how the fundamental principles of the Belmont Report can be expanded to address the emerging research landscape. This chapter then ends with a consideration of open questions that pose some of the biggest concerns for ensuring continuing protection of human subjects.

At the outset of this chapter, it is important to stress that the concerns noted in this chapter are not limited to any particular type of data. While the majority of examples given here will focus on social media or user behavior data, this focus is a natural byproduct of the kinds of data that have been available for study to date. However, as society's online interactions become more complex—and as it becomes cheaper to store and share the increasingly complex data that result from those interactions—it is important for psychological scientists to apply these principles to *all* forms of human data and to carefully consider what new privacy and security challenges richer data may pose (e.g., video data; cf. Bertino, *this volume*).

A Brief History of the Ethical Landscape for Psychological Science

To understand the challenges facing our field, we should first examine why our ethical and legal frameworks for human-subjects research ethics exist and how they manifest themselves today.

The Belmont Report and the Common Rule

Egregious violations of human rights in the mid-20th century led the U.S. Congress to enact legislation that was pivotal in creating the current U.S. system of human-subjects ethics. A

¹ While this chapter is most directly interested in exploring large-scale data use, many researchers who use smaller-scale online data may also find these questions useful to consider in their work as well.

comprehensive recounting of the emergence of the Belmont Report is outside of the scope of the current chapter, but a brief sketch of what ethical historians consider to be the three most influential experiments will be helpful for framing this discussion. (For more on the historical, ethical, and philosophical contexts of these events—including other, less well-known horrors from the biomedical and behavioral sciences—see Baker, 2001; Beauchamp, 2011; Drewry, 2004; and Rice, 2008.)

The first two experiments were biomedical atrocities. First, the Nazi human experiments on unwilling prisoners in the 1940s—exposed to the world during the Nuremberg trials—catalyzed the development of worldwide ethical principles for human biomedical research (see Annas & Grodin, 1992). Second, the Tuskegee Study of Untreated Syphilis tracked the progression of untreated syphilis from 1932 to 1972 in hundreds of poor African-American men who were unaware of the experiment and uninformed of eventual treatment options (see Farmer, 2003; Reverby, 2009).

The third experiment was by no means equivalent in the magnitude of harm caused by the first two experiments, but it nevertheless demonstrated the potential risks posed to participants by behavioral research. U.S. psychologist Stanley Milgram (1963)—directly inspired by the Nuremberg Trials—deceived and coerced participants into delivering what they believed would be painful electric shocks to another individual. The study’s methods raised ethical questions for social and behavioral research, especially for the use of deception (Baumrind, 1964, 1979; Englehardt & Englehardt, 2013; Schlenkerand & Forsyth, 1977).

Although the Nazi and Tuskegee experiments were incomparably different from Milgram’s (1963) experiment in the type, duration, and level of harm that they caused, these (and other) patently immoral and unethical studies sparked efforts to create legal and moral

frameworks for human-subjects research around the world. In 1947, the Nuremberg Code emerged as a result of the Nuremberg trials (reprinted in Annas & Grodin, 1992) and laid down 10 principles that eventually formed the basis for the medical research ethics outlined in the Declaration of Helsinki nearly two decades later (World Medical Association, 1964). At the time, the United States signed onto the Declaration of Helsinki and ostensibly adopted its standards for biomedical research. However, ten years later, public outcry at the Tuskegee syphilis experiment—along with increasing questions about the potential dangers of behavioral research (cf. Baumrind, 1964)—led Congress to create the National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research to explore options for improving human-subjects research safety (Public Law 93-348, 1974).

Five years later, the committee's work culminated in the publication of the Belmont Report (National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research, 1979). The Belmont Report was intended to be a non-legislative statement of core values for human-subjects research. It laid out three foundational principles with clear ties to specific requirements at various stages of the research process:

1. *Respect for persons* upholds the dignity and autonomy of all human research subjects.

From it, we have requirements for informed consent, for additional constraints for researchers who intend to recruit participants from protected populations, for maximizing voluntariness, and for strict guidelines on any research involving deception.

2. *Beneficence* is a conceptual extension of the “do not harm” principle. It explicitly mandates that researchers maximize potential benefits *and* minimize potential harm to individual research subjects. From this principle, we have the obligations

to balance the ratio of individual risks to potential social benefits, to assess the severity and probability of individual risks, and to more carefully weigh risks to protected populations.

3. *Justice* calls for the equal distribution of potential benefits and potential risks across all groups who could potentially benefit from the research. From it, we have the duty to equitably select research subjects from the broader population by minimizing individual or systemic biases that would shift potential risks onto a subset of the population (especially members of protected populations or underrepresented groups) while allowing society at large to benefit.

The U.S. Department of Health and Human Services formally incorporated the Belmont Report's guidelines into binding policies under the Common Rule in 1981 (revised again in 1991 and 2018; 45 CFR 46, 2018). Today, the Common Rule applies to human-subjects research that falls under the purview of 16 U.S. federal agencies and departments, from the Department of Agriculture to the Social Security Administration. Perhaps the most visible contribution of the Common Rule for most researchers is the creation of Institutional Review Boards (IRBs); these ethical bodies are responsible for overseeing human-subjects research that receives direct or indirect funding from the U.S. government.

Current Ethical Oversight for Federally Funded Research

Crucially, activity must meet two very specific requirements to be subject to IRB review: It must be (1) *research* involving (2) *human subjects*. *Research* is defined as “systematic investigation, including research development, testing, and evaluation, designed to develop or contribute to generalizable knowledge” (Section 46.102(1); 45 CFR 46, 2018). In its most recent revision, the Common Rule has been explicitly updated to *exclude* certain categories of activities

that could have been construed as research—specifically, “[s]cholarly and journalistic activities (e.g., oral history, journalism, biography, literary criticism, legal research, and historical scholarship),” “[p]ublic health surveillance activities,” “[c]ollection and analysis of information [...] for a criminal justice agency,” and “[a]uthorized operational activities [...] in support of [...] national security” (Section 46.102(1)(1-4) 45 CFR 46, 2018). A *human subject* is defined as:

[...] a living individual about whom an investigator (whether professional or student) conducting research:

- (i) Obtains information or biospecimens through intervention or interaction with the individual, and uses, studies, or analyzes the information or biospecimens; or
- (ii) Obtains, uses, studies, analyzes, or generates identifiable private information or identifiable biospecimens. (Section 46.102(e); 45 CFR 46, 2018)

Under these definitions, most psychological scientists in academia have engaged with IRBs through their work collecting and analyzing laboratory-based scientific data. For these scientists, it would seem only natural that the collection of new data after directly recruiting participants—whether online or in person—would require them to submit their *protocols* (i.e., formal research plans) for review and approval by their IRBs prior to beginning their research. Many researchers who are first considering working with non-laboratory datasets may not think to seek approval from their IRBs, but federal guidelines do require some oversight for certain kinds of such datasets.

Currently, IRBs can make one of three determinations on research projects using big data and naturally occurring datasets. First, the research could be considered *not human-subjects research*, meaning that the IRB does not need to review it. Second, it could be ruled as falling under category 4 (“reanalysis of existing data”) of the *exempt* IRB classification—somewhat of a

misnomer given that it still falls under a lighter form of IRB review. Finally, it could fall under *expedited* or *full-board* classifications, both of which require a higher level of review.

Taking a very simplified view of the regulations, we can essentially classify the review of existing datasets by answering 4 questions: (Q1) whether the dataset is *entirely* available to the public (without any restrictions whatsoever, including registering for free accounts);² (Q2) whether the dataset contains “private data” (like medical or school data)³; (Q3) whether the data includes identifiers or possible identifiers; and (Q4) whether the data were received first by the researcher as identifiable. (See Figure 1 for flow chart.)

 Insert Figure 1 About Here

A non-human-subjects-research determination can be made either [1] when the dataset (Q1) is publicly available and (Q2) contains no private data or [2] when the dataset (Q1) is publicly available, (Q2) contains private data, (Q3) *currently* contains no participant identifiers, and (Q4) was never sent to the current researchers with any identifiers. This is possible because of the definition of a “human subject” in 45 CFR 46 (e.g., University of California Berkeley Committee for the Protection of Human Subjects, 2016; University of Chicago Social and Behavioral Sciences IRB, 2014). However, individual universities may decide to systematically

² However, requiring payment is generally considered permissible, so long as there are no restrictions designating eligible purchasers.

³ According to the Common Rule, the question of whether data are “private” essentially refers to whether there could be a reasonable expectation of privacy around the data. Simply including personally identifiable information is not sufficient to be considered private. For example, a photograph is inherently personally identifiable information, but a photograph that is shared publicly on a social media website would not *necessarily* be considered private data. Issues of privacy are discussed more in the Open Questions section at the end of this chapter.

limit non-human-subjects-research determinations (Inter-university Consortium for Political and Social Research, n.d.).

A determination of exempt category 4 can be made when datasets—either [1] (Q1) publicly available datasets (Q2) with private information or [2] (Q1) non-publicly available datasets—have (Q3) no identifiers because (Q4) the identifying data were removed from the dataset by the current researchers. Interestingly, the most recent update to the Common Rule has grown to include *prospective* data acquisition under exempt category 4, whereas the pre-2018 Common Rule required the data *already exist* prior to the current researcher’s involvement (45 CFR 46, 2018). Generally, this means that datasets with (Q2) private *and* (Q3) identifiable data will be subject to expedited or full-board review.

The determination of whether a project falls under “human-subjects research” (or any other IRB classification) may *only* be made by an IRB; no researcher can make this determination for themselves. While this may be natural to researchers in psychology, it is important to note that some academic researchers are engaged in IRB-eligible activity without being aware of it (e.g., Dittrich & Kenneally, 2012). This is especially likely to occur in computer science, mathematics, statistics, and other fields that have not traditionally conducted human-subjects research but are now interested in big data or data science (e.g., Metcalf & Crawford, 2016). Accordingly, all researchers—especially those conducting federally funded research or who work at public U.S. institutions—should consult their IRB prior to beginning work on any human-derived data.

Belmont Principles in the 21st Century

Keeping our field’s legal and ethical framework (and its history) in mind, let’s move on to consider how our current challenges can fit within our existing framework.

Ethics Lessons from Recent Studies of “Wild” Data

Big data or naturally occurring datasets (BONDS; Paxton & Griffiths, 2017) afford psychological scientists the opportunity to test, expand, and refine theories by analyzing human behavior in the real world. BONDS are typically not created for general scientific purposes but can, with a bit of careful thinking and the right computational tools, provide crucial insights into psychological science and complement rigorous lab-based experimental inquiry. Keeping in mind the proper awareness of limitations, messiness, and potential biases of these data (e.g., Ioannidis, 2013; Lazer, Kennedy, King, & Vespignani, 2014), real-world data—especially from social media or other social platforms—have been increasingly seen as another valuable tool for psychological scientists to add to their research toolkits (e.g., Goldstone & Lupyan, 2016; Jones, 2016; Lazer et al., 2009). To be clear, BONDS research should not be seen as rebuking or replacing traditional experimental psychological methods: Rather, the clearest value of BONDS to psychological science lies in their ability to *complement* these traditional methods, creating a “virtuous cycle of scientific discovery” (Paxton & Griffiths, 2017, p. 1631).

Along with the promising theoretical and empirical contributions of BONDS research, however, some scientific⁴ BONDS research has raised ethical concerns. In one example, academic researchers scraped over 700,000 profiles from a dating website and then published the entire dataset—including highly identifiable information like usernames—in an open-access repository (Kirkegaard & Bjerrekaer, 2016). The resulting public outcry over the breach in participant privacy without participant consent or IRB oversight eventually caused the repository to remove both the data and manuscript preprint (see Zimmer, 2018).

⁴ There are, of course, equally or more problematic *non*-scientific uses of BONDS data (e.g., the Cambridge Analytica scandal; Granville, 2018; Laterza, 2018). To the extent that these uses intersect with scientific concerns, these are discussed later in the chapter; otherwise, an in-depth discussion of them are outside the scope of the current chapter.

In a second example, researchers scraped over 35,000 profile pictures and corresponding sexual orientation data⁵ from the publicly available online dating profiles of nearly 15,000 users (Wang & Kosinski, 2018). These data were used to create various classifiers that could identify a person's sexual orientation from a photograph with 71-91% accuracy. In addition to scientific concerns about the study's methodology and interpretation (e.g., that the classifier likely picked up on other markers of self-presented sexual orientation and gender identity, like hair styles or photo filters; Gelman, Mattson, & Simpson, 2018), the public and LGBTQ+ advocates raised concerns about potential implications for the physical and social well-being for gay men and women (e.g., GLAAD & Human Rights Campaign, 2017).

While the first two examples used existing data, the final example discussed here was an experimental study conducted on social networks (Kramer, Guillory, & Hancock, 2014). Specifically, by manipulating over 689,000 users' news feeds to show either more positive-emotion or more negative-emotion posts from their Facebook friends, these researchers were the first to demonstrate that real-time in-person communication was not necessary for *emotional contagion* (i.e., becoming more similar to someone else's displayed emotion) to occur. Despite this novel scientific finding, general public alarm at the study—especially because users were unaware of their participation in the study, were not given the opportunity to provide informed consent, and could not withdraw their participation—led Facebook to establish new internal principles for research review (see Jackman & Kanerva, 2016). However, public and scientific scrutiny over the process raised significant questions about how Cornell University (the collaborating academic institution) justified its decision to claim that its researchers were not involved in human-subjects research (e.g., Meyer, 2014), despite vague reports from Cornell that

⁵ By comparing the user's gender with the gender of their desired partners.

the researchers had been involved in “initial discussions” about the study (cf. quoted Cornell press release in Sullivan, 2014; original press release removed from Cornell University’s website: <http://news.cornell.edu/media-relations/2014/06/30/media-statement-on-cornell-universitys-role-in-facebook-emotional-contagion-research/>).

Each of these scientific projects presents a mix of unique and overlapping ethical concerns. The new opportunities from new data sources are incredibly exciting, but in an era without clear ethical guidelines for today’s changing data landscape, even the most well-intentioned researchers can make mistakes. Although the researchers on these projects may not have been fully aware of the potential ethical implications of their work, lessons from these projects can inform not only the ethical considerations for BONDS research in psychology but also ongoing conversations about BONDS ethics more broadly. Critically, because of psychology’s longstanding concerns with human-subjects ethics, psychological scientists are well-poised to influence BONDS ethics policies even outside of human-subjects research.

Belmont Today

Although the writers of the Belmont Report could not have anticipated the new challenges brought by 21st-century technology and data, the principles underlying the Belmont Report are broad enough to apply to research in any century. For example, the Menlo Report (Bailey, Dittrich, & Kenneally, 2013; Dittrich & Kenneally, 2012) sought to apply Belmont principles to “information and communication technology” researchers and gave some prescient insights into problems facing such researchers at the dawn of BONDS research. Building off of the Belmont and Menlo Reports, I extend the three core Belmont principles to apply to human-focused research (especially in academia) today.

Many of these issues target data *collection*—a particular challenge that many researchers who use existing datasets may not directly face. However, as ethics-bound researchers, we should consider how data were collected prior to using it: The *fact* of data availability does not immediately confer the *ethicality* of analysis, and researchers should seriously consider refraining from using unethically collected or unethically shared datasets. On the other hand, some researchers may be in a position to effect change through direct collaboration with the entities creating BONDS datasets. Whether or not they were engaged in the data collection, each researcher should take it upon themselves to safeguard public trust and individual rights. Public trust in science is a precious shared resource, and taking advantage of that resource to improve personal fame, notoriety, or career advancement will lead to its marked depletion to the detriment both of the remainder of the scientific community and society at large (cf. tragedy of the commons; Lloyd, 1883).

In this section, I outline some of the major hurdles to Belmont-compatible behavior in using BONDS research, along with some ways that psychological scientists could begin to address them.

Respect for Persons. Under this principle, individuals' autonomy should be upheld throughout the research process by giving them with enough information to make an informed decision about participation and allowing them to withdraw that consent at any time. Today, BONDS-based researchers are increasingly questioning how compatible the principle of informed consent is with opaque, overly broad terms of service and with rampant dragnet data collection (e.g., Flick, 2016; Zimmer, 2018). An increasing proportion of daily activity occurs online on forums with mandatory terms of service—long, convoluted, and dense documents that

a miniscule percentage of people read (e.g., approximately 2% of users; Obar & Oeldorf-Hirsch, 2018). How voluntary and informed, then, is consent under these conditions?

To more fully embrace the respect for persons principle in BONDS research, researchers should move to opt-in models of research participation, outlining explicitly how and to what end their data will be used if they choose to participate. Alternatively, in situations where opt-in models are judged infeasible or impossible by their ethics board, researchers should work with their ethics board to create an opt-out model with clear and frequent reminders to participants of their rights to drop out of the study at any time. All information shared with participants should be easily obtainable and accessibly written, especially as it pertains to the technical details surrounding the storing and use of their data.

Beneficence. While individual participants may not benefit directly from their participation in BONDS research (as is true with most in-lab psychological science research), researchers must protect their participants from modern-day financial, reputational, and other harms. Today, beneficence requires keen awareness of issues around data privacy and data breaches. Even dealing strictly with public-domain data, massive data-gathering and data-sharing efforts allow researchers to gather dangerous quantities of data about specific individuals (e.g., Metcalf & Crawford, 2016; Rivers & Lewis, 2014; Zimmer, 2018) and raise serious questions about whether de-identified data can ever be truly free from the possibility of re-identification (cf. Netflix challenge data re-identification in Narayanan & Shmatikov, 2008; see also Ohm, 2010). Moreover, some of the most valuable data to companies and researchers today—namely, video and audio data—are inherently and inextricably identifiable (see, e.g., Aghajanzadeh, Li, Jebb, Lu, & Thiruvathukal, *this volume*; Blake, Lee, De La Rosa, & Sherman, *this volume*). Given the potential risks that even public data pose, the potential risk posed by

breaches to sensitive data is nearly unimaginable even if the data do not include financial records or government identification (e.g., ransomed data from a social media site facilitating extramarital affairs; Mansfield-Devine, 2015). How protected, then, are individuals from potential harms of not only participation itself but from the potential misuse or mismanagement of the associated data?

To more fully embrace the beneficence principle in BONDS research, researchers should minimize harm to individual participants by exercising extreme caution when deciding what data to collect. No data are free from the *possibility* of a hack or other form of breach, no matter how secure or well-designed data protections are. Therefore, researchers should minimize the type and amount of data collected from individual participants, with extra consideration given to any data that could be considered potentially sensitive or private. For large-scale data collection, the best protection for participants' long-term data security is, quite simply, not collecting the data in the first place. In addition to enduring commitments to data security, researchers reconsider their motivations in data collection and be judicious in what data are gathered: Just because the data *can* be collected, does not mean that they *should* be collected.

Complementarily, to maximize the benefits to individual participants, researchers should more freely share data and deliverables with participants. Questions of data ownership have been on the rise in industry and medicine for some time (e.g., Cios & Moore, 2002), but members of the public today are expressing increasing interest in having access to their own data (e.g., Ritzer & Jurgenson, 2010; Tene & Polonetsky, 2013). One tangible benefit to participation in BONDS research could be in helping to satisfy participants' curiosity about themselves and their digital lives. Given how much easier open science tools have made sharing data and deliverables (e.g.,

apps, data exploration tools) among researchers and the public, the additional effort to provide insight would be relatively minimal relative to the potential value to participants.

Justice. The principle of justice calls for an equal distribution of benefits and risks across the population. Today, psychological scientists must ensure that the risks of BONDS data do not fall on the limited and non-representative subset of the population who engage with entities and services collecting the data (e.g., Vitak, Shilton, & Ashktorab, 2016) and that groups of people—especially vulnerable groups—are not further disadvantaged by data nor by products created with those data (e.g., O’Neil, 2017). While issues around algorithmic justice for women, communities of color, and other underrepresented groups have been gaining some attention over the past few years (e.g., Hamilton, 2018; O’Neil, 2017), concerns for justice are especially sharp as they relate to data from and about children—a population that is not considered capable of giving consent themselves for traditional laboratory experiments (45 CFR 46, 2018) but that often has data shared about them (e.g., on social media) without their knowledge or assent by caregivers and others (Berman & Albright, 2017). How representative or biased, then, are our datasets and results?

To more fully embrace the principle of justice in BONDS research, researchers should be as concerned with questions of representation among their BONDS participants as they would be for traditional laboratory paradigms. Explicit demographics information may not be available (or advisable to always collect, as discussed in the “Beneficence” subsection above), but when possible, efforts should be made to ensure that data are as representative as possible. Issues of self-selection and online access may make this particularly difficult to achieve with certain groups (e.g., rural communities, lower socioeconomic status, older individuals). Researchers using existing data could remedy this by subsampling their dataset to statistically over-represent

underrepresented groups and under-represent overrepresented groups in the final sample; researchers collecting new data could additionally improve the representativeness of their sample by improving recruitment to over-represent underrepresented groups. Subsampling data to reduce the impact of participants from overrepresented groups on the results could result in lower statistical power, but it could also increase the applicability and validity of findings. This is especially true when researchers are working with data that are linked with and potentially contain human biases and structural inequalities (e.g., arrest records, hiring records) to create deliverables that could exacerbate those inequalities (e.g., Hamilton, 2018; O’Neil, 2017).

Questions of fairness and equity should also significantly influence researchers’ selection of data and research questions. Again, the simple fact that a dataset is available does not mean that a researcher is ethically legitimized to use it, and taking advantage of existing (even public) data could cause certain groups of people to bear an unnecessary share of risk. For example, when creating their classifier of gay and straight faces, Wang and Kosinski’s (2018) stated desire to show the limitations of human perception and the superiority of computer algorithms put greater share of the risk of participation on gay men and women who could face economic, social, and even physical harm by being identified against their will. This risk falls not only on the gay men and women who were included in the sample but potentially *all* gay men and women worldwide; because there is no obvious additional benefit to gay men and women for having participated, this is a significant cause for concern for the principle of justice. To improve this, researchers should be acutely aware of the real-world implications to the groups included in and potentially affected by their research, and researchers who are unsure of the potential implications could solicit advocacy groups for feedback and input on how using BONDS data from their groups could help or harm them.

Beyond the Belmont Report: Open Questions for Today's Psychological Scientists

Today's technological, legal, and social milieu present new challenges for psychological scientists using BONDS. Here, I outline some new considerations that have not yet been neatly addressed by our ethical framework. While this section provides some suggestions on how to address some of them, careful thought and collaboration among all stakeholders will be essential to tackling the new ethical challenges of 21st-century human-subjects research.

Balancing Open Science with Participant Rights

Psychological science—like other scientific areas—has been increasingly focused on adopting open-science practices, including open data (Kidwell et al., 2016). Intriguingly, the Menlo Report (Dittrich & Kenneally, 2012) advocated explicitly for forward-thinking improved transparency as part of its principle of respect for law and public interest. Transparency—not only in information and communication technologies research but also in all scientific work—helps improve public trust in science by facilitating access to the scientific process and products.

However, while some principles of open science can be readily embraced by BONDS researchers (e.g., open materials, open access), the principle of open data should be carefully considered, given the potential privacy ramifications for the participants included. As mentioned under the Respect for Persons subsection above, serious concerns exist about the limits of de-identification and the real risks of re-identification (even of sparse datasets) in a cultural moment when so much data are shared (Narayanan & Shmatikov, 2008; Ohm, 2010). This risk of re-identification becomes more concerning to individual participants' rights as researchers use potentially sensitive data and as researchers derive potentially impactful deliverables from those data. This is especially true for BONDS researchers who are working with datasets that are not

entirely public (although even entirely public datasets should not be used without consideration, as discussed below).

The tools for sharing data from laboratory experiments could provide a model for BONDS data, too. For example, the ICPSR (<https://www.icpsr.umich.edu/>) has a special class of “gated” repositories for sensitive data: These *restricted-use repositories* are freely available to any ICPSR-verified researcher, blending a commitment to open data with dedication to participant rights. While appropriate de-identification should take place to the maximal extent possible, such a vehicle for sensitive data-sharing could help minimize the impact of open data on individual participants.

In any case, specific questions of open data must be evaluated according to the unique needs and risks of each dataset. The value of open BONDS data could also be weighed according to whether we see open data as *public goods* or *common goods*, a debate that can be informed by similar debates within genomics (e.g., Bialobrzeski, Ried, & Dabrock, 2012). Briefly, a *public good* would be like clean air: something that everyone should be entitled to use by fiat, that no one should be barred from using, and that cannot be used up in zero-sum way. By contrast, a *common good* is like a public park or universal healthcare—something that should be shared by all because it can help improve the general well-being of others but that is subject to terms set by group deliberation. A model of participant data as a public good would mandate data access to every (valid) researcher but would also require that every researcher contribute all data. By contrast, a model of participant data as a common good allows for more flexibility for participants and researchers in what, how, and with whom data are shared, with compliance emerging from a common conviction of the utility and power of the data shared.

In light of the respect for persons principle, it seems most reasonable to treat BONDS research in psychological science as a common good. Such a perspective affords flexibility in data-sharing while fostering a community commitment to the value of open data. For example, Bialobrzeski and colleagues (2012) proposed that genetics repositories could work with patients to provide them with options for how they would share their data and for how long it could be kept on file. A similar model could be considered in BONDS research, allowing participants to decide how much (if any) of their data they want to have publicly shared rather than simply requiring participants to share their data as a condition of their participation.

Contextualizing Human Data Use Outside Academia

The Common Rule only applies to research conducted using federal funds, but companies and private organizations are also exploring human behavior and cognition through their users' data. Most of this would likely not fall under the definition of human-subjects research, since it often doesn't seek to find "generalizable knowledge" (45 CFR 46, 2018) and focuses instead on corporate goals like improving user experience. However, even projects that are intended to contribute generalizable findings would not be bound by these requirements because they are not federally funded.

This clear gap between the expectations for human data use inside versus outside of academia is an area of concern for ethicists and data experts broadly. Although peer-reviewed scientific journals increasingly require authors confirm that their protocol was approved by an appropriate ethics board in the manuscript (e.g., Graf et al., 2007), a plurality of industry-focused research is never ultimately published. To that end, Calo (2013) proposed an analogue to IRBs for industry—*consumer subject review boards* (CSRBs). Although CSRBs have not been widely adopted, interest in these and similar entities have increased as the line blurs between research

and corporate activity (e.g., Jackman & Kanerva, 2016; Polonetsky, Tene, & Jerome, 2015; Tene & Polonetsky, 2016). With the volume and heterogeneity of companies' user data collection on the rise, data scientists and big data researchers using human-derived data are grappling with whether and how to provide oversight to corporate research.

Combining BONDS and Institutional Ethics

The current provisions in the Common Rule (45 CFR 46, 2018) were originally developed to handle existing or public datasets that were dramatically different from the options today. The Belmont Report and original Common Rule emerged from a time when large-scale data collection was only possible through massive organizational initiatives (e.g., government records, academic research) and a handful of corporate entities (e.g., insurance companies). Now, hundreds of thousands of organizations collect and share data on hundreds of millions of individuals worldwide.

Despite the significant concerns about the gaps in ethical oversight at the intersection of academic and industry research in that first came to light following the Facebook affective contagion study (e.g., Flick, 2016; Sullivan, 2014), the 2018 revisions to the Common Rule make it clear that there is still relatively little understanding of the unique dangers posed by BONDS, especially when combining multiple public datasets (cf. Metcalf & Crawford, 2016). Industry researchers are grappling with their own questions of oversight (as noted earlier), but academic researchers should think carefully about the risks of BONDS collaborations in light of their potential impact.

As academic researchers collaborate more with companies on interesting and complex datasets, the ethics boards governing academic research must begin to raise real questions about the risks posed even by data collected by businesses as a matter of course or through internal

experimental work. Under the current terms of IRB review, academic researchers could simply claim no ethical responsibility for the data collection and proceed to work with those data, but in the face of current skepticism over science and scientific practices, it seems short-sighted to divorce considerations of ethical data collection from ethical data use.

Researchers should have a duty—not only to specific participants in a study but also to the scientific community and the public—to conduct ethical research: Using ethically obtained data (no matter *who* obtained it) must remain a pillar of that ethical obligation. However, at present, individual researchers—not an ethics board—would make such a decision about whether a non-human-subjects dataset *should* be ethically used. Moving forward, academic researchers and ethical bodies should consider whether it would be more prudent to formalize the process to provide additional oversight.

Securing Computational Pipelines

The increasing complexity of data collection and analysis necessitates increasing scale in our computational and storage capacity, and the ease of cloud-based options is attractive to many scientists working on collaborative projects. The necessity of these resources have grown to the point where the National Science Foundation has funded national cyberinfrastructure for researchers in the U.S. (e.g., XSEDE; Towns et al., 2014). Because these computational pipelines are growing faster than most psychological scientists can learn them, systems administrators and other cyberinfrastructure personnel are essential to ensuring the smooth (and safe) functioning of the hardware and software.

Despite the heavy reliance on the computational and human systems enabling discovery using BONDS, relatively less attention has been paid to the ethical oversight of these systems. For example, let's say that a researcher is working with a cloud computing platform to analyze

sensitive large-scale data—perhaps running computer vision algorithms over videos or natural language processing over patient medical files. While the latter is much better protected by HIPAA regulations, much less consensus exists around how to secure computational pipelines for inherently identifiable data that are not HIPAA-protected, like video. It is important to note that this extends beyond simply hardware and software: Questions of what ethical regulations, training, or approvals might apply to systems administrators are similarly important and as yet unsettled.

Redefining “Minimal Risk” and “Private” Data

Given the power of connecting multiple datasets—even those that are entirely publicly available—data scientists themselves are increasingly arguing for oversight of big data (Metcalf & Crawford, 2016). A requirement of expedited and exempt IRB categories is that projects under these designations pose no more than *minimal risk*, defined as research in which “the probability and magnitude of harm or discomfort anticipated in the research are not greater in and of themselves than those ordinarily encountered in daily life or during the performance of routine physical or psychological examinations or tests” (Section 46.102(j); 45 CFR 46, 2018). The paramount importance of minimizing risk to participants is, at first glance, consistent with why entirely public data is not considered human-subjects research: Because the data are already public, there is an assumption that no *additional* risk would be incurred if the data were used for scientific purposes.

However, the vast quantities of rich data available can—when woven together—produce insights that could have the power to harm specific individuals. For example, Hauge, Stevenson, Rossmo, and Le Comber (2016) used only freely available data in an attempt to uncover the real identity of the artist Banksy, but the project still uncovered a massive amount of information

about the person singled out (for more, see Metcalf & Crawford, 2016). This is also related to concerns around re-identification of private data discussed earlier (e.g., Narayanan & Shmatikov, 2008; Ohm, 2010)—for example, if researchers leverage open datasets to re-identify de-identified datasets with private information. Understanding the true potential for harm in these data—especially when using open data to conduct research on underrepresented or potentially vulnerable groups (e.g., gay men and women; Wang & Kosinski, 2018)—should give researchers and ethical bodies pause when considering whether research activities using open data truly pose “minimal risk” simply by virtue of their openness.

Dovetailing with concerns about what should count as “minimal risk” are questions about what data should count as “private.” According to current federal regulations,

[p]rivate information includes information about behavior that occurs in a context in which an individual can reasonably expect that no observation or recording is taking place, and information that has been provided for specific purposes by an individual and that the individual can reasonably expect will not be made public (e.g., a medical record).

(italics in original; Section 46.102(e)(4); 45 CFR 46, 2018)

An essential part of this definition is the concept of whether the individual has a *reasonable expectation* that their activity will not be recorded or observed. This is, for example, one reason why research that relies on observation of public behavior is considered minimal risk and falls under an exempt category. Again, the assumption is that—because the behavior itself was executed in public—there would be no *additional* risk to participants if data on the public behavior were used for scientific purposes.

Therefore, a crucial question is whether individuals acting online expect that they are acting publicly or privately. The majority of people do not read privacy policies or understand

the legality of broad tracking initiatives online (e.g., Hoofnagle, Urban, & Li, 2012; Martin, 2015)—an unfortunate reality that could explain the “privacy paradox” (i.e., the widespread prevalence of sharing data despite widespread stated concerns about privacy; Smith, Dinev, & Xu, 2011) and that presents concerns for researchers using online data. By contrast, ethics boards would be appalled if the majority of participants in a lab-based experiment failed to understand what data they were giving to researchers.

Even in the cases of outright sharing (e.g., on social media), many IRB professionals express extreme reservations with considering such data as public, citing concerns about whether the individuals truly understood the impact of their sharing (Woelfel, 2016). Others have suggested that social media sites present a sort of public-private hybrid that has no real face-to-face or in-person analogue (e.g., Strauß & Nentwich, 2013). Put simply, if the majority of people are not aware that their behavior *could be* (and likely *is*) tracked extensively on a single online platform or across the internet—regardless of whether a researcher or ethics board finds that lack of awareness to be reasonable—we should be hypervigilant about perceptions of violations of privacy. Moreover, although traditional types of private data included medical and educational records, this lapse in general understanding of tracking suggests that we might move to align our concepts of “private” data to better conform to the general public’s understanding of what data could reasonably be considered private. This is especially true in an age when data shared online are likely to exist in perpetuity.

Despite surface parallels with in-person observation, researchers should grapple with questions of scope in online settings. In real-life settings, a crowd can provide a form of anonymity through obscurity by providing too many targets for a single person to reasonably track; however, in online arenas, both the virtual crowd and its constituent members can be

simultaneously tracked with high fidelity in real time. Online data collection affords passive, dragnet data collection at a scale and precision that would be unimaginable to attain using human observers. Observation through BONDS data collection—especially when combining multiple datasets (e.g., Metcalf & Crawford, 2016)—is so vast as to feel *qualitatively* different from observations performed by note-taking humans or even by passive video cameras. This difference in perception should drive us to reevaluate whether our consideration of behavior in digital “public spaces” is truly equivalent to behavior in real-life public spaces. However, as observation of private and public spaces through large-scale video analysis becomes more prevalent and more computationally tractable, similar questions may come to be asked even of real-life behavior (cf. Aghajanzadeh et al., *this volume*).

Limitations

The present work has been intended to situate new questions of ethics in BONDS research within the existing ethical frame for psychological scientists. This chapter—like all scientific works—has its own set of limitations, including noting that several interesting and important questions fall outside of its scope.

First, this chapter has focused on concerns both for researchers involved in new BONDS data collection and for researchers (re-)analyzing existing BONDS. This, of course, does not completely address the unsettling collection and use of data by companies in the first place—a problem that has been increasingly recognized in the U.S. and around the world. However, as psychological scientists, we often have less direct control over that problem. Instead, our consideration of ethical questions for datasets can guide our choices of which datasets to use, which companies to collaborate with (or work for), what curricula to teach, and what ethical and legal structures to advocate.

Second, legal and ethical questions about what companies can and should be doing with users' data are being raised worldwide as the public becomes increasingly aware of companies' collection, tracking, and use of user data. The lines between *scientific research* (for identifying generalizable knowledge about human behavior) and *company testing* (for improving a specific business's or industry's performance) are legally distinct in the U.S.—even if many users (and scientists) might see them as nearly identical. Large-scale collection by companies, of course, is not unique to this time—for example, actuarial research by insurance companies have long aggregated data as a core part of their business—but it now occurs at an unprecedented granularity and pace. Addressing such questions would require interrelated legal and ethical frameworks; however, such proposals are outside the scope of the current chapter.

Third, this chapter largely centered on the U.S. legal and ethical framework, but care for human subjects has always been of international concern. The Nuremberg Code (see reprint in Annas & Grodin, 1992) and the Declaration of Helsinki (World Medical Association, 1964) both originated outside of the U.S., for example. Recently, the European Union General Data Protection Regulation (GDPR; European Parliament, 2016) enacted sweeping reforms to data collection and use in the E.U., prompting some ancillary changes in the U.S. as international entities shifted their business practices. Among other things, the GDPR reinforced the “right to be forgotten” (Ausloos, 2012)—which itself could present new challenges to researchers aggregating and storing naturally occurring data—and mandated that all user-focused research be explicitly opt-in. These reforms address some of the concerns outlined in this chapter, and similar reforms should be seriously considered (and supported) by U.S.-based researchers.

Finally, as of the time of writing this chapter, the 2018 revision to the Common Rule still contains several gaps that the U.S. Department of Health and Human Services has yet to fill,

some of which will be relevant to BONDS researchers. (One of the most notable is a flow chart to determine whether a project will require IRB oversight.) However, additional guidance and documentation may emerge that could alter the landscape of ethical review for BONDS researchers. As these and other changes take effect, BONDS researchers should continue to educate themselves about their ethical responsibilities—and to call for stronger legal and ethical frameworks to protect human subjects, science, and society.

Conclusion

Psychological scientists today have an unprecedented opportunity to expand our field of study into more natural arenas through capitalizing on big data and naturally occurring datasets. By adopting the tools of data science and staying grounded in rich theoretical and experimental traditions, we can use these data as a window into real-world behavior, cognition, and emotion to help us test, expand, and refine psychological theory. Despite these promising avenues, this paradigm presents new ethical challenges to individuals and to society. However, our core ethical principles—the Belmont principles of respect for persons, beneficence, and justice—can be expanded to address the risks and benefits of today’s data, not only protecting the rights and dignity of our individual participants but also preserving the public’s faith and trust in psychological science.

References

- 45 CFR 46 (2018). Department of Health and Human Services. Retrieved from <https://www.hhs.gov/ohrp/regulations-and-policy/regulations/45-cfr-46/index.html>
- Annas, G. J., & Grodin, M. A. (1992). *The Nazi doctors and the Nuremberg Code: Human rights in human experimentation*. New York: Oxford University Press.
- Ausloos, J. (2012). The “right to be forgotten” - Worth remembering? *Computer Law & Security Review*, 28, 143–152. <https://doi.org/10.1016/j.clsr.2012.01.006>
- Bailey, M., Dittrich, D., & Kenneally, E. (2013). *Applying ethical principles to information and communication technology research*. Retrieved from <http://www.dhs.gov/csd-resources>
- Baker, R. (2001). Bioethics and human rights: A historical perspective. *Cambridge Quarterly of Healthcare Ethics*, 10, 241–252.
- Baumrind, D. (1964). Some thought on ethics of research: After reading Milgram’s “Behavioral Study of Obedience.” *American Psychologist*, 19(6), 421–423. <https://doi.org/10.1037/h0040128>
- Baumrind, D. (1979). IRBs and social science research: The costs of deception. *IRB: Ethics & Human Research*, 1(6), 1–4. Retrieved from <https://www.jstor.org/stable/pdf/3564243.pdf?refreqid=excelsior%3A078c3d2d14862a94a133eba8e32313f1>
- Beauchamp, T. L. (2011). The Belmont Report. In E. J. Emanuel, C. Grady, R. A. Crouch, R. K. Lie, F. G. Miller, & D. Wendler (Eds.), *The Oxford textbook of clinical research ethics*. New York: Oxford University Press.
- Berman, G., & Albright, K. (2017). *Children and the data cycle: Rights and ethics in a big data world*. Retrieved from <https://www.unicef-irc.org/publications/907/>

- Bialobrzeski, A., Ried, J., & Dabrock, P. (2012). Differentiating and evaluating common good and public good: Making implicit assumptions explicit in the contexts of consent and duty to participate. *Public Health Genomics*, *15*, 285–292. <https://doi.org/10.1159/000336861>
- Calo, R. (2013). Consumer subject review boards: A thought experiment. *Stanford Law Review Online*, *66*(97), 97–102.
- Cios, K. J., & Moore, G. W. (2002). Uniqueness of medical data mining. *Artificial Intelligence in Medicine*, *26*, 1–24.
- Cornsweet, T. N., & Crane, H. D. (1973). Accurate two-dimensional eye tracker using first and fourth Purkinje images. *Journal of the Optical Society of America*, *63*(8), 921–928.
- Dittrich, D., & Kenneally, E. (2012). *The Menlo Report: Ethical principles guiding information and communication technology research*. Retrieved from <http://www.dhs.gov/csd-resources>
- Drewry, S. (2004). The ethics of human subjects protection in research. *The Journal of Baccalaureate Social Work*, *10*(1), 105–117. <https://doi.org/10.18084/1084-7219.10.1.105>
- Englehardt, E. E., & Englehardt, R. K. (2013). The Belmont Commission and a progression of research ethics. *Ethics in Biology, Engineering & Medicine-An International Journal*, *4*(4), 315–326.
- European Parliament. General Data Protection Regulation (2016). European Parliament. Retrieved from <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32016R0679>
- Farmer, P. (2003). *Pathologies of power: Health, human rights, and the new war on the poor*. Berkeley: University of California Press.
- Flick, C. (2016). Informed consent and the Facebook emotional manipulation study. *Research Ethics*, *12*(1), 14–28. <https://doi.org/10.1177/1747016115599568>

- Gelman, A., Mattson, G., & Simpson, D. (2018). Gaydar and the fallacy of decontextualized measurement. *Sociological Science*, 5, 270–280. <https://doi.org/10.15195/v5.a12>
- GLAAD, & Human Rights Campaign. (2017). GLAAD and HRC call on Stanford University & responsible media to debunk dangerous & flawed report claiming to identify LGBTQ people through facial recognition technology. Retrieved from <https://www.glaad.org/blog/glaad-and-hrc-call-stanford-university-responsible-media-debunk-dangerous-flawed-report>
- Goldstone, R. L., & Lupyan, G. (2016). Discovering psychological principles by mining naturally occurring data sets. *Topics in Cognitive Science*, 8(3), 548–568. <https://doi.org/10.1111/tops.12212>
- Graf, C., Wager, E., Bowman, A., Fiack, S., Scott-Lichter, D., & Robinson, A. (2007). Best practice guidelines on publication ethics: a publisher's perspective. *International Journal of Clinical Practice*, 61(Suppl. 152), 1–26. <https://doi.org/10.1111/j.1742-1241.2006.01230.x>
- Granville, K. (2018, March 19). Facebook and Cambridge Analytica: What you need to know as the fallout. *The New York Times*. Retrieved from <https://www.nytimes.com/2018/03/19/technology/facebook-cambridge-analytica-explained.html>
- Hamilton, I. A. (2018, October 10). Amazon built AI to hire people, but it discriminated against women. *Business Insider*. Retrieved from <https://amp.businessinsider.com/amazon-built-ai-to-hire-people-discriminated-against-women-2018-10>
- Hauge, M. V., Stevenson, M. D., Rossmo, D. K., & Le Comber, S. C. (2016). Tagging Banksy: using geographic profiling to investigate a modern art mystery. *Journal of Spatial Science*, 61(1), 185–190. <https://doi.org/10.1080/14498596.2016.1138246>

- Hoofnagle, C. J., Urban, J. M., & Li, S. (2012). Privacy and modern advertising: Most US internet users want “Do Not Track” to stop collection of data about their online activities. In *Amsterdam Privacy Conference*. Retrieved from <http://ssrn.com/abstract=2152135>
- Inter-university Consortium for Policital and Social Research. (n.d.). Institutional Review Boards (IRBs). Retrieved from <https://www.icpsr.umich.edu/icpsrweb/ICPSR/irb/>
- Ioannidis, J. P. A. (2013). Informed consent, big data, and the oxymoron of research that is not research. *The American Journal of Bioethics*, *13*(4), 40–42.
<https://doi.org/10.1080/15265161.2013.768864>
- Jackman, M., & Kanerva, L. (2016). Evolving the IRB: Building robust review for industry research. *Washington and Lee Law Review Online*, *72*(3), 442–457.
- James, J. (2018). Data Never Sleeps 6.0. Retrieved October 9, 2018, from <https://www.domo.com/blog/data-never-sleeps-6/>
- Jones, M. N. (2016). Developing cognitive theory by mining large-scale naturalistic data. In M. N. Jones (Ed.), *Big data in cognitive science* (pp. 1–12). New York, NY: Routledge.
- Kidwell, M. C., Lazarević, L. B., Baranski, E., Hardwicke, T. E., Piechowski, S., Falkenberg, L.-S., ... Nosek, B. A. (2016). Badges to acknowledge open practices: A simple, low-cost, effective method for increasing transparency. *PLoS Biology*, *14*(5), e1002456.
<https://doi.org/10.1371/journal.pbio.1002456>
- Kirkegaard, E. O. W., & Bjerrekaer, J. D. (2016). The OKCupid dataset: A very large public dataset of dating site users. *Open Differential Psychology*, (May).
<https://doi.org/10.26775/ODP.2016.11.03>
- Kramer, A. D. I., Guillory, J. E., & Hancock, J. T. (2014). Experimental evidence of massive-scale emotional contagion through social networks. *Proceedings of the National Academy of*

Sciences, 111(24), 8788–8790. <https://doi.org/10.1073/pnas.1320040111>

Laterza, V. (2018). Cambridge Analytica, independent research and the national interest.

Anthropology Today, 34(3), 1–2.

Lazer, D., Kennedy, R., King, G., & Vespignani, A. (2014). The parable of Google Flu: Traps in big data analysis. *Science*, 343, 1203–1205.

<https://doi.org/10.1109/PASSAT/SocialCom.2011.98>

Lazer, D., Pentland, A., Adamic, L., Aral, S., Barabási, A.-L., Brewer, D., ... Van Alstyne, M. (2009). Computational social science. *Science*, 323, 721–723.

<https://doi.org/10.1145/1134271.1134277>

Lloyd, W. F. (1883). *Two lectures on the checks to population*. England: Oxford University.

Mansfield-Devine, S. (2015). The Ashley Madison affair. *Network Security*, 2015(9), 8–16.

[https://doi.org/10.1016/S1353-4858\(15\)30080-5](https://doi.org/10.1016/S1353-4858(15)30080-5)

Martin, K. (2015). Privacy notices as tabula rasa: An empirical investigation into how complying with a privacy notice is related to meeting privacy expectations online. *Journal of Public Policy & Marketing*, 34(2), 1547–7207. <https://doi.org/10.1509/jppm.14.139>

Metcalf, J., & Crawford, K. (2016). Where are human subjects in Big Data research? The emerging ethics divide. *Big Data & Society*, 3(1), 1–14.

<https://doi.org/10.1177/2053951716650211>

Meyer, R. (2014, June 28). Everything we know about Facebook’s secret mood manipulation study. *The Atlantic*. Retrieved from

<https://www.theatlantic.com/technology/archive/2014/06/everything-we-know-about-facebooks-secret-mood-manipulation-experiment/373648/>

Milgram, S. (1963). Behavioral Study of Obedience. *Journal of Abnormal and Social*

Psychology, 67(4), 371–378. <https://doi.org/10.1037/h0040525>

Morris, J. R. W. (1973). Accelerometry—a technique for the measurement of human body movements. *Journal of Biomechanics*, 6, 729–736. [https://doi.org/10.1016/0021-9290\(73\)90029-8](https://doi.org/10.1016/0021-9290(73)90029-8)

Narayanan, A., & Shmatikov, V. (2008). Robust De-anonymization of Large Sparse Datasets. In *IEEE symposium on security and privacy* (pp. 111–125). Retrieved from https://www.cs.utexas.edu/~shmat/shmat_oak08netflix.pdf

National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research. (1979). *The Belmont report: Ethical principles and guidelines for the protection of human subjects of research*. Bethesda, MD.

O’Neil, C. (2017). *Weapons of math destruction: How big data increases inequality and threatens democracy*. New York, NY: Broadway Books.

Obar, J. A., & Oeldorf-Hirsch, A. (2018). The biggest lie on the Internet: ignoring the privacy policies and terms of service policies of social networking services. *Information, Communication, & Society*. <https://doi.org/10.1080/1369118X.2018.1486870>

Ohm, P. (2010). Broken promises of privacy: Responding to the surprising failure of anonymization. *UCLA Law Review*, 57, 1701–1777.

Paxton, A., & Griffiths, T. L. (2017). Finding the traces of behavioral and cognitive processes in big data and naturally occurring datasets. *Behavior Research Methods*, 49(5), 1630–1638. <https://doi.org/10.3758/s13428-017-0874-x>

Polonetsky, J., Tene, O., & Jerome, J. (2015). Beyond the Common Rule: Ethical structures for data research in non-academic settings. *Colorado Technology Law Journal*, 13(101), 333–337.

- Public Law 93-348. National Research Act (1974). United States Congress.
- Reverby, S. M. (2009). *Examining Tuskegee: The infamous syphilis study and its legacy*. Chapel Hill: University of North Carolina Press.
- Rice, T. W. (2008). The historical, ethical, and legal background of human-subjects research. *Respiratory Care*, 53(10), 1325–1329.
- Ritzer, G., & Jurgenson, N. (2010). Production, consumption, prosumption: The nature of capitalism in the age of the digital “prosumer.” *Journal of Consumer Culture*, 10(1), 13–36.
<https://doi.org/10.1177/1469540509354673>
- Rivers, C. M., & Lewis, B. L. (2014). Ethical research standards in a world of big data. *F1000 Research*, 3(38). <https://doi.org/10.12688/f1000research.3-38.v1>
- Schlenkerand, B. R., & Forsyth, D. R. (1977). On the ethics of psychological research. *Journal of Experimental Social Psychology*, 13, 369–372.
- Smith, H. J., Dinev, T., & Xu, H. (2011). Information privacy research: An interdisciplinary review. *Management Information Systems Quarterly*, 35(4), 989–1015.
- Strauß, S., & Nentwich, M. (2013). Social network sites, privacy and the blurring boundary between public and private spaces. *Science and Public Policy*, 40, 724–732.
<https://doi.org/10.1093/scipol/sct072>
- Sullivan, G. (2014, July 1). Cornell ethics board did not pre-approve Facebook mood manipulation study. *The Washington Post*. Retrieved from <https://www.washingtonpost.com/news/morning-mix/wp/2014/07/01/facebooks-emotional-manipulation-study-was-even-worse-than-you-thought/>
- Tene, O., & Polonetsky, J. (2013). Big data for all: Privacy and user control in the age of analytics. *Northwestern Journal of Technology and Intellectual Property Volume*, 11(5),

240–273. <https://doi.org/10.1177/073490419901700102>

Tene, O., & Polonetsky, J. (2016). Beyond IRBs: Ethical guidelines for data research.

Washington and Lee Law Review Online, 72(3), 458–471. Retrieved from

<https://scholarlycommons.law.wlu.edu/wlulr-online/vol72/iss3/7>

Towns, J., Cockerill, T., Dahan, M., Gaither, K., Grimshaw, A., Hazlewood, V., ... Wilkins-

Diehr, N. (2014). XSEDE: Accelerating scientific discovery. *Computing in Science and Engineering*, 16(5), 62–74.

University of California Berkeley Committee for the Protection of Human Subjects. (2016).

Research involving the secondary use of existing data. Retrieved from

<https://cphs.berkeley.edu/secondarydata.pdf>

University of Chicago Social and Behavioral Sciences IRB. (2014). *Guidance on Secondary*

Analysis of Existing Data Sets. Chicago, IL. Retrieved from

<https://sbsirb.uchicago.edu/page/secondary-data-analysis>

Vespignani, A. (2009). Predicting the behavior of techno-social systems. *Science*, 325, 425–428.

<https://doi.org/10.1016/j.socnet.2009.02.004>

Vitak, J., Shilton, K., & Ashktorab, Z. (2016). Beyond the Belmont principles: Ethical

challenges, practices, and beliefs in the online data research community. In *Computer*

Supported Cooperative Work (pp. 941–953). New York, NY: ACM.

<https://doi.org/10.1145/2818048.2820078>

Wang, Y., & Kosinski, M. (2018). Deep neural networks are more accurate than humans at

detecting sexual orientation from facial images. *Journal of Personality and Social*

Psychology, 114(2), 246–257. <https://doi.org/10.1037/pspa0000098.supp>

Woelfel, T. (2016). *Behind the computer screen: What IRB professionals really think about*

social media research. University of Washington. Retrieved from

https://digital.lib.washington.edu/researchworks/bitstream/handle/1773/36448/Woelfel_washington_02500_16207.pdf

World Medical Association. Declaration of Helsinki - Ethical principles for medical research involving human subjects (1964). Helsinki, Finland.

Zimmer, M. (2018). Addressing conceptual gaps in big data research ethics: An application of contextual integrity. *Social Media and Society*, 4(2), 1–11.

<https://doi.org/10.1177/2056305118768300>

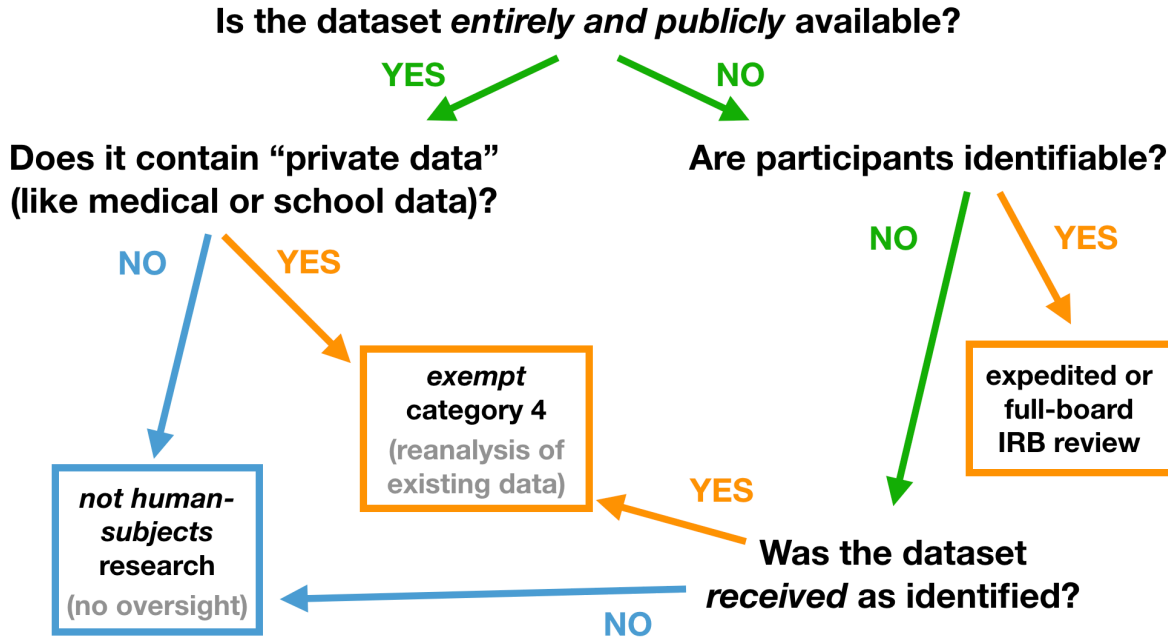


Figure 1. Simplified flow chart of the regulations used to determine the level of oversight required for existing datasets in federally funded research. However, all final decisions about IRB review are made by IRBs, not by individual researchers. (Blue lines lead to a non-human-subjects-research determination. Orange lines lead to a determination requiring IRB oversight.

Green lines indicate a path that could end in either determination.)