



NATIONAL ENDOWMENT FOR THE HUMANITIES

OFFICE OF DIGITAL HUMANITIES

Narrative Section of a Successful Application

The attached document contains the grant narrative and selected portions of a previously funded grant application. It is not intended to serve as a model, but to give you a sense of how a successful application may be crafted. Every successful application is different, and each applicant is urged to prepare a proposal that reflects its unique project and aspirations. Program guidelines also change and the samples may not match exactly what is now required. Please use the current set of application instructions to prepare your application.

Prospective applicants should consult the current program application guidelines at <https://www.neh.gov/program/dangers-and-opportunities-technology-perspectives-humanities>

Applicants are also strongly encouraged to consult with the NEH Office of Digital Humanities staff well before a grant deadline.

Note: The attachment only contains the grant narrative and selected portions, not the entire funded application. In addition, certain portions may have been redacted to protect the privacy interests of an individual and/or to protect confidential commercial and financial information and/or to protect copyrighted materials.

Project Title: *Good Decisions: Data Science as a Moral Practice*

Lead Institution: Syracuse University

Project Directors: Johannes Himmelreich

Grant Program: Dangers and Opportunities of Technology: Perspectives from the Humanities (Individual Researcher)

List of personnel

Good Decisions: Data Science as a Moral Practice (single researcher project)

Project director (PD)

- **Johannes Himmelreich**, Assistant Professor of Public Administration and International Affairs, *Syracuse University*

Co-author (CA) (funded independently, not covered by this grant)

- **Sebastian Köhler**, Associate Professor of Philosophy, *Frankfurt School of Finance and Management*, Germany

Good Decisions: Data Science as a Moral Practice

Single researcher project — PI: Johannes Himmelreich, PhD, Syracuse University

Project Summary

This project investigates the technology of data science (a collection of techniques to extract value from data). The project advances the argument that data science is a moral practice characterized by inherent ethical dilemmas. The project makes this argument by bringing normative theories and philosophy of science to bear on the practice of data science. The goal of the project is to offer a systematic analysis of the nature of data science and its inherent ethical dilemmas. The project expands the understanding of a topic in the humanities (values in science) and explores the relationship between technology (data science) and society. Key activities are identifications of ethical dilemmas in each step in the data science work cycle—these steps include data collection, data “cleaning”, data analysis, and communication. The main project outcome is a book manuscript; further outcomes are two peer-reviewed open-access journal articles. Each of the steps in the data science work cycle will be the topic of a chapter and/or article. Work on this project has already begun in an existing co-author collaboration. The co-author is neither eligible nor in need of funding by the NEH. This is thus a single researcher project application.

Significance and Contribution

Data science is the engine of decision-making today. The US Bureau of Labor Statistics estimates the profession will grow 30% over 10 years. In 2019 alone, LinkedIn reports, job ads for data scientists increased by 37%. Several laws require policies and regulations to be grounded in data. This reflects a credo of our times: *Good decisions are data-driven.*

This project advances the argument that this idea of “good decisions” must be understood comprehensively. Good decisions, and good data science, needs to be *good* not only in a factual or epistemic sense of being accurate, but also in a moral sense: Good data science needs to be justifiable and arise from a legitimate process. In a slogan: Data science is a moral practice. To make this argument, the project attends closely to individual steps in the cycle of data science work and identifies ethical dilemmas at each step. The project **contributes to humanities research** by extending both normative theories as well as philosophy of science and bringing them to bear on the practice of data science. The project’s **target audience** are students and scholars in philosophy, as well as students, scholars, and practitioners of data science. The project outcomes will be available open access where possible (see budget justification and below) and adhere to best practices for accessibility (e.g. appropriate digital text file formats).

The project explores the technology of data science. **Data science** combines techniques from statistics, computer science, and management to inform the decision making of public and private organizations. It is used to decide which welfare claims to accept, where to deploy police, which tax returns to investigate, whether a suspect is granted bail, and whether an immigrant is detained. The stakes in such decisions are high and their impacts are profound (e.g. O’Neil, 2016; Eubanks, 2018; Benjamin, 2019; Raji et al., 2020). One example: As the COVID-19 pandemic ravaged prisons, who had to stay in prison and who was released to shelter at home was determined by predictions of inmates’ recidivism risk.

These and other examples of data science raise urgent challenges—for society at large and for data scientists. One is the **challenge of justice**: For example, which idea of a just society, if any, should inform data science recidivism predictions? Another is the **challenge of data dominance**: When predicting fraud, is it true that “we can just let the data speak”, or do values and hidden assumptions enter the analysis? In the public sector in particular, data science raises the **challenge of pluralism**: How should data science reflect the rich range of values in society? How should opposing preferences about police

work inform predictive policing? What makes data-driven policing not just morally justified but politically legitimate? This distinctively political dimension is crucial (Himmelreich, 2020).

This project analyzes data science as a moral practice. It identifies ethical dilemmas inherent in data science. It explains why these dilemmas arise, relates them to the three challenges above, and asks what can be done to address them. The main outcome is a book manuscript that integrates ethical, methodological, and political-theoretical investigations.

The project's **overarching research question** is: What makes for *good* data science and for a good data scientist? The main hypothesis concerns the nature of data science as a moral practice: *Data science acts in the world; it does not merely represent it.* In result, good data science is more than good data.¹ Good data scientists cultivate (epistemic) virtues and consider not only *what* decisions they make but *how* they make them—since decisions need to be justifiable (*what* decisions are made) as well as legitimate (*how* decisions are made). For further research questions, see book outline (below).

The project's **key argument** is that **data scientists act in the world—and data science is a moral practice—because data science raises difficult ethical dilemmas at each step: from project conception, via data “cleaning” and analysis, to product deployment.** For example, data “cleaning”, i.e. identifying outliers and dealing with incomplete records, “is as much a moral as a practical concern” (Barocas & boyd, 2017). Similarly, statistical analyses force a choice between conflicting notions of “fairness” or “equity” and between competing moral interests of different stakeholders (Corbett-Davies et al., 2017). Although many such individual dilemmas are already known, this book provides a synthesizing perspective from the humanities—in particular, from normative theories and philosophy of science.

The project's **intellectual significance** as well as its intended **impact** is a fuller appreciation of why data science is prone to be controversial. In addition to several well-known and important reasons, this project argues that data science is bound to be controversial because of its very nature: As a moral practice, it consists of difficult ethical dilemmas. This view—that data science is controversial because it is ethically difficult—complements views contending that data scientists' biases, their defective intent, or their ignorance about consequences are the reason why data science is prone to controversy. The project argues instead that data science tends to be controversial not only because human and social biases translate into algorithmic bias, and not only because data scientists lack a professional ethos, but because data science is replete with questions that have no uncontroversial right answer. Many technical or methodological questions of data science are also ethical questions.

The project's **research methods and ethical framework** are those of practical and interdisciplinary philosophy. The main method is a theoretical and normative analysis that takes seriously the fact of value pluralism and persistent multifaceted injustice. The project complements literatures that identify injustice and bias insofar as it answers conceptual (*What is data?*) and normative questions (*Should you take data as given?*).

Environmental Scan and Project History

The **field of digital technology ethics** comprises several books for a general-interest audience that highlight the pernicious effects of data science (e.g. O'Neil, 2016; Eubanks, 2018; Noble, 2018). Important contributions examine race and gender (e.g. Benjamin, 2019; Buolamwini & Gebru, 2018; Gebru, 2020).

Science and technology studies (STS) interrogates assumptions, contextualizes technologies, and uncovers their semiotic and material effects in journals such as *Science, Technology & Human Values*. The empirical dimension is fruitfully combined with the feminist tradition (D'Ignazio & Klein, 2020). **Computer scientists and statisticians** have distinguished different definitions of fairness (Chouldechova & Roth, 2018, 2020; Corbett-Davies & Goel, 2018; Kleinberg et al., 2016; Mitchell et al., 2021). The larger research field (around the conference FAccT) has recently been synthesized into a book, which addresses

primarily an engineering audience (Hardt & Recht, 2022). **Human-centered data science** synergizes STS and statistics and advocates for interdisciplinarity qualitative research (Aragon et al., 2022).

The existing **literature of technology ethics within philosophy** does not yet fully appreciate the extent to which *ethics and methodology are intertwined in the data science process*.² This contrasts with philosophical scholarship that engages with other areas of science, which examines how ethics and methodology are intertwined in biology (Leonelli, 2016), emphasizes the role of human agency in inquiry (Reiss, 2015), and admonishes science to be oriented towards democracy (Kitcher, 2011; Longino, 1990) or practicality (Cartwright, 2019; Chang, 2017). The project can thus build on a rich existing literature in philosophy. The project specifically extends existing research on values in science, which was pioneered by feminist scholars (Longino, 1990; Douglas, 2009), by pursuing a pluralistic and political approach (Himmelreich, 2018a, 2020; Schroeder, 2022).

The **project history** originates in courses that the project director (PD) as well as this project's co-author (CA) have been teaching at their respective institutions for four years, which generated the research ideas that are pursued in this project. The project **builds on the PD's on-going research**. The PD has published extensively in applied philosophy and technology: on virtual reality (2018b), responsibility for "killer robots" (2019), digital democracy (2022a, 2022b), the digital economy (forthcoming), self-driving cars (2018a, 2020, 2022c), and structural injustice in artificial intelligence (AI) (Himmelreich & Lim, 2022). Together with an interdisciplinary team, the PD investigates the ethics of AI in government (Young et al., 2022; Young, Himmelreich, Bullock, et al., 2021; Young, Himmelreich, Honcharov, et al., 2021). Joint work with the co-author is on AI and responsibility (Himmelreich & Köhler, 2022).

A pitch for a framing piece for this project, a descendant of which could be the basis for the book introduction, is currently under consideration with an editor of the *Boston Review*.

Activities and Research Team

The project timeline is two years. The **PD's regular teaching load** comprises four weeks of summer teaching (instead of fall teaching) and two courses in the spring semester. The project funding primarily serves to **reduce the PD's teaching load in each year** over the duration of the project.

The PD will pursue complementary opportunities to fund another semester of research leave in the spring of 2025 (with a fellowship of Syracuse University's Humanities Center, one course reduction).

The PD has been granted a research leave for fall 2023, which will be used to finalize the book proposal and two sample chapters. **Year 1** starts with the submission of the proposal to publishers. By the end of the year 2024, four further book chapter drafts will be completed, one will be submitted to a journal (e.g. *Ethical Theory and Moral Practice*). A first milestone in **year 2** is the submission of another chapter as a journal article and revisions of the first journal article. By the end of the year, the manuscript draft and revisions for the second article should be completed.

The **project team** consists of the PD, Dr. Johannes Himmelreich, and a co-author (not covered by the requested funding). The **project director** is responsible for the successful conduct and completion of all project activities. A **co-author** on the book manuscript is Dr. Sebastian Köhler whose expertise on metaethics complements the PD's expertise on applied ethics and the philosophy of social science. Both team members benefit from synergies of their complementary expertise as well as from mutual accountability that arises from undertaking such a collaboration.

Final Products and Dissemination

Final products of this project will be a book manuscript, at least two journal articles, and two domestic trips to disseminate findings at research conferences.

The **book manuscript** is to be drafted by EOY 2025. Research results will be disseminated as **conference presentations** at interdisciplinary academic conferences, such as FAccT or AIES. Some results also will be published as **journal articles**. Journals that offer affordable open access publication will be targeted with priority. The project budget includes funds for **open-access publication**. Additionally, **drafts and accepted publications** will be made available on personal websites and repositories. Special attention will be paid to making the documents **accessible** to people with disabilities (e.g. by making available file formats and formatting suitable for screen readers).

The **final products align with the project's goals** in two ways. First, the book manuscript allows for the intended larger synthesis and the analysis of the entire data science project cycle. Second, open-access journal articles, draft publications, and conference presentations serve important dissemination functions to advance the impact and the contribution of the project to the humanities.

Since no data collection or analysis is undertaken, and given the norms of academic integrity and attribution, as well as legal obligations stipulated in author agreements required by publishers, no significant risks of **privacy, confidentiality, or intellectual property** are anticipated. Research notes and manuscript files are stored on proprietary cloud services (esp. OneDrive, iCloud Drive, Zotero storage), which can be considered reasonably safe.

Book Outline: Topics and Research Questions

The book identifies and examines ethical challenges at each step of the data science work cycle, which consists of (1) project conception, (2) data acquisition, (3) data processing, (4) modelling, (5) evaluation, and (6) deployment.³ Preliminary research topics, questions, and hypotheses have been identified. The three challenges—of justice, data dominance, and pluralism—frame the discussion throughout.

The first step, *project conception*, raises the challenge of **finding the right problem**: To what questions is data science the answer? One chapter of the book demarcates the limits of data science and explains how data scientists are responsible for the consequences of their work. A further chapter discusses **which ethical values, if any, should guide data science**, such as freedom, neutrality, welfare, equity, social justice, or the common good.

On *data acquisition*, one chapter argues that **data acquisition starts with a conception of the social world**. Should, for example, gender be represented as binary? Such representations are subject to ethical evaluation just as actions (Longino, 1995; Basu, 2019; Johnson, forthcoming).⁴ Further chapters investigate the **ethics of data** (of data ownership and data stewardship).

After data are acquired, the next step is *data processing*. Data scientists remove data that are incomplete, invalid, or otherwise erroneous (Ilyas & Chu, 2019). They aim for accuracy: to represent reality correctly (Olson, 2003). But this aim for accuracy is increasingly chimerical as data scientists work with “soft” data (Akerlof, 2020). Soft data—social, emotional, or psychological properties, such as how suspicious a person looks in a video feed or what emotional timbre their voice has in a recording—do not represent reality but interpret it. One chapter pursues the research question of what notion of **data quality beyond accuracy** should guide data processing.

Once processed, data are used in *modelling and evaluation*, that is, in building mathematical representations and using statistical methods to predict, explain, and understand events. Two chapters investigate **how data scientists' modelling choices impact decision-making**. Even technical choices are value judgments, for example, when assessing whether being a woman causes lower wages. When a model separates the broader context of gender socialization (e.g., occupational preferences) from the gender variable, then gender's effect on wages is reduced (Hu & Kohler-Hausmann, 2020).

The final part of the book, on *project deployment*, covers **the ethics of putting data science to use**. The chapters recommend practices that responsible data scientists should engage in.

Workplan

Good Decisions: Data Science as a Moral Practice (single researcher project)

The project’s timeline is two years (January 01, 2024, to December 31, 2025). The project director (PD) is responsible for the successful conduct and completion of all project activities. Conference presentations are not included in the workplan overview.

Overview

	Activities	Output
2023	Fall <ul style="list-style-type: none"> – Research leave, approved (not funded by this application) – Finalize book proposal – Solicit feedback – Writing: two sample chapters 	<ul style="list-style-type: none"> – Book proposal – Two sample chapters – Excerpt of chapter (e.g. introduction or framing piece) as dissemination piece under review at free and popular venue (e.g. <i>Boston Review</i>)
	Spring <ul style="list-style-type: none"> – Submission of proposal and sample chapters to publishers – Writing: one new chapter – Revision: dissemination piece – Negotiation and selection of publisher 	<ul style="list-style-type: none"> – One further chapter draft – Dissemination piece placed – Contract with publisher – Feedback on proposal and existing drafts
2024	Milestone 1: three chapters drafts completed	
	Summer <ul style="list-style-type: none"> – Writing: one chapter 	<ul style="list-style-type: none"> – Drafts of one further chapter
	Fall <ul style="list-style-type: none"> – Writing: two chapters – Revision: existing chapters – Editing: journal article 1 	<ul style="list-style-type: none"> – Drafts of two further chapters – Revision of three chapters – Journal article 1 under review
Milestone 2: six chapters drafts completed, journal article 1 under review		
2025	Spring <ul style="list-style-type: none"> – Writing: one chapter – Revision: two chapters – Editing: journal article 2 	<ul style="list-style-type: none"> – Drafts of one further chapter – Revision of two chapters from spring – Journal article 2 under review
	Milestone 3: seven chapters drafts completed, journal article 2 under review	
	Summer <ul style="list-style-type: none"> – Writing: one chapter – Revision: existing chapters – Revision: journal article 1 	<ul style="list-style-type: none"> – Drafts of one further chapter – Revision of two chapters – Journal article 1 accepted
2025	Fall <ul style="list-style-type: none"> – Writing: final chapter – Revision: existing chapters – Revision: journal article 2 	<ul style="list-style-type: none"> – Drafts of remaining chapters – Revision of two or three existing chapters – Journal article 2 accepted
	Milestone 4: book manuscript complete, journal articles accepted	

Milestones

The project is divided into four milestones. Because the content backbone of the project consists in the research and writing that is undertaken for the book project, and because these activities will take the most time, the milestones are formulated with respect to the progress on the book manuscript.

Milestone 1 (one semester after project start, end of May 2024) is that *three chapters* are drafted. This is enabled by a full research leave in fall 2023 (approved, not funded by this application), which will allow for the finalization of the book proposal and a draft of two sample chapters. The PD also expects to be able to publish a dissemination piece, with material that will explaining the theoretical framing, in an outlet that is freely accessible and has a wide academic readership (such as the *Boston Review*). The PD submitted a first pitch draft in January 2023.

Milestone 2 (end of the year 2024) is that drafts of *six chapters* have been completed and that one of the existing chapters is reworked and submitted to a peer-reviewed journal. This is feasible because the PD has no teaching obligations in the fall of 2024 (because of teaching in the summer 2024).

Milestone 3 (spring 2025) marks the completion of seven chapter drafts; moreover, a second chapter will be reworked and submitted to a peer-reviewed journal for publication. This is feasible because of course reductions funded by this application.

Milestone 4 (end of year 2025) is the completion of the manuscript and the acceptance, or conditional acceptance, of the two journal articles.

Research and writing activities

The NEH DOT project will enable the PD to commit significant time to this project. The project budgets 1.06 academic months each year to the project, which is equivalent to a reduction of the teaching load of one course per year (regular teaching load is 4 courses per year). Syracuse University faculty appointments are for 8.5 months. The PD teaches for four weeks during the summer months.

Risks

The project anticipates one significant risk: slow editorial processes. This risk concerns the publication of the two peer-reviewed articles.

Speed of editorial processes. Since the pandemic, the turnaround time (or time-to-decision) of journals in the humanities has increased significantly, according to anecdotal experience. Often, this is due to the difficulties of soliciting reviewers. Again, anecdotally, this risk is less severe in this the sub-field of philosophy and in applied ethics. However, to mitigate this potential risk, the project workplan prioritizes an early submission of the journal articles.

The project does not consider staff attrition or problems in the co-author collaboration a significant risk. PD and CA have worked together regularly since 2016. Since 2019 they have worked on a joint research project, which has led to this current project and a published co-authored article.