

Title: Big Data, Big Responsibilities: The Law and Ethics of Business Analytics

Instructor: [Mihir Kshirsagar](#)

Contact: mkshirsa@wharton.upenn.edu

Office Hours (Zoom): Thursdays 2:00p -3:00p or by appointment

General Description

This course introduces students to the legal, policy, and ethical dimensions of algorithmic decision making. We are in the midst of a profound shift in society where critical decisions in commercial and social contexts are increasingly being run through automated decision-making systems. This increased reliance on data-driven decisions creates opportunities for both financial gain and social good. But it poses challenges to social norms and laws that requires careful thinking about how to make sound, ethical decisions. In particular, because regulations are rapidly evolving in this space, there is an urgent need for forward-thinking businesses to anticipate these novel risks and build sustainable business models. The class will engage in discussions about cutting-edge research and novel challenges that are now frequently in the news.

Learning Objectives

My goal is to help you think about challenging analytical questions through an ethical framework. The course is designed to help you:

1. Understand the legal and ethical implications of business decisions;
2. Exercise thoughtful judgment and communicate your values;
3. Evaluate the risks and opportunities of relying on algorithmic decision making.

Course Structure

We will cover a lot of ground in just 6 weeks. I have attempted to keep the readings manageable, with the goal that everyone comes to the seminar fully prepared to participate actively. Each session will be divided into two parts. First, is a discussion of the readings. Second, is a close analysis of a particular case study. I will share a short memo to prepare you for the case study discussion. In class, I will ask students to adopt the perspectives of different stakeholders and engage in a structured debate of the issue.

Course Requirements and Grading

Strategic Memo [50%]

The final project is a strategic memo directed to a specific company or organization that recommends whether to adopt an algorithmic decision-making tool to make business decisions. The memo should be 2000-3000 words (provide a word count at the end) and should describe the context for the decision, identify the relevant legal and ethical challenges, and make concrete recommendations for what action the business should take. The memos will be graded based on the quality of the analysis, engagement with the ethical considerations, persuasiveness, and organization. I will ask for a draft introductory paragraph for the memo two weeks before the deadline.

Reading Responses [25%]

By 10a on the day of our seminar (Thursday) please submit three or four bullet points sharing your reactions to or questions about one or more of the assigned readings.

Participation [25%]

Participation credit will be assessed holistically: quality matters, not just quantity. The seminar structure is particularly dependent on active participation, respectful listening, and open dialogue.

Deadlines

Deadlines are an inevitable part of life. Please respect them out of regard for yourself and your peers. If you encounter difficulties, I am open to finding any accommodations, but I need sufficient notice before the deadline or there will be a grade penalty.

Weekly Schedule

Week 1: Overview

When Does Predictive Technology Become Unethical? Eric Siegel, Harvard Business Review, Oct. 2020: <https://hbr.org/2020/10/when-does-predictive-technology-become-unethical>

Putting Responsible AI Into Practice, Rumman Chowdhury et al., MIT Sloan Management Review, Oct. 2020: <https://sloanreview.mit.edu/article/putting-responsible-ai-into-practice/>

Case study: Student Data

Pre-read: Virtual Classrooms and Real Harms: Remote Learning at U.S. Universities, Cohney et al., USENIX SOUPS 2021: <https://arxiv.org/abs/2012.05867>.

Week 2: Fundamentals

Ethical Lenses in Technology Practice, Shannon Vallor, Irina Raicu, Brian Green, Markkula Center for Applied Ethics, July 2020: <https://www.scu.edu/ethics-in-technology-practice/ethical-lenses/>

Introduction: Fairness & Machine Learning, Solon Barocas, Moritz Hardt, Arvind Narayanan: <https://fairmlbook.org/introduction.html>

Case Study: Content Moderation

Pre-read: Content moderation, AI, and the question of scale, Tarleton Gillespie, *Big Data & Society*, Aug. 2020: <https://journals.sagepub.com/doi/full/10.1177/2053951720943234>

Week 3: Fairness & Discrimination

Assessing risk, automating racism, Ruha Benjamin, *Science* (Oct. 25, 2019):
<https://science.sciencemag.org/content/366/6464/421>.

Big Data: A Report on Algorithmic Systems, Opportunity, and Civil Rights, White House Report (May 2016):
https://obamawhitehouse.archives.gov/sites/default/files/microsites/ostp/2016_0504_data_discrimination.pdf

Optional: Watch *Coded Bias* on Netflix

Case Study: Apple Card

Pre-read: New York Department of Financial Services Report on the Apple Card Investigation, Mar. 2021:
https://www.dfs.ny.gov/system/files/documents/2021/03/rpt_202103_apple_card_investigation.pdf

Week 4: Dark Patterns

Big Data: A Tool for Inclusion or Exclusion? Understanding the Issues, Federal Trade Commission (Jan. 2016): <https://www.ftc.gov/reports/big-data-tool-inclusion-or-exclusion-understanding-issues-ftc-report>

Dark Patterns: Past, Present, and Future, Narayanan et al., Communications of the ACM, <https://cacm.acm.org/magazines/2020/9/246937-dark-patterns/fulltext>

Case Study: Online Advertising

Pre-read: Department of Housing & Urban Development lawsuit against Facebook:
https://www.hud.gov/sites/dfiles/Main/documents/HUD_v_Facebook.pdf

Week 5: Privacy & Data Protection

Big data, artificial intelligence, machine learning and data protection, ICO United Kingdom (2017), <https://ico.org.uk/media/for-organisations/documents/2013559/big-data-ai-ml-and-data-protection.pdf>

Case study: Algorithmic Hiring

Pre-read: Princeton Dialogue on AI and Ethics, Hiring by Machine Case Study:
<https://aiethics.princeton.edu/wp-content/uploads/sites/587/2018/12/Princeton-AI-Ethics-Case-Study-5.pdf>

Week 6: Accountability & Data Governance

2021 EU AI Regulation Proposal:
<https://ec.europa.eu/newsroom/dae/redirection/document/75792>.

Optional Additional Materials:

AI Nation Podcast: <https://whyy.org/programs/ai-nation/>

Black Mirror on Netflix

Movies: Minority Report; Moneyball, Margin Call, Coded Bias, The Social Dilemma,
The Circle