

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

LGST 6420; Fall 2022

Instructor: [Mihir Kshirsagar](#)

Contact: mkshirsa@wharton.upenn.edu

Office Hours: 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 require careful thinking about how to make sound, ethical decisions. 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 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 cover a lot of ground in just 6 classes. I have attempted to keep the readings manageable, with the goal that everyone comes to the class fully prepared to participate actively. Each session will be divided into two parts. First, is a lecture and discussion of the readings. Second, is a case study discussion. I will share a short memo to prepare you for the case study. 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

Final Paper [50%]

The final project is a strategic memo directed to a specific company or organization that recommends whether and how to adopt a data-driven decision-making tool. The memo should be ~2000 words (provide a word count at the end) and should describe the context for the decision, identify the relevant legal or ethical challenges, and make concrete recommendations for what action the business should take to address those challenges. The memos will be graded based on the quality of the analysis, engagement with the ethical considerations, persuasiveness, and organization. I will offer an opportunity to get feedback on a draft introductory paragraph and outline for the memo ahead of the deadline.

Weekly Responses [25%]

By 10a on the day of our seminar (Thursday) please submit three or four bullet points sharing your reactions to the weekly prompt. Completing the assignment on time will score you full points.

Participation [25%]

Participation credit will be assessed holistically: quality matters, not just quantity. The course structure is particularly dependent on active participation, respectful listening, and open dialogue. Please use the feature on Canvas to record your attendance and request any absences via email.

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 accommodations, but I need sufficient notice before the deadline or there will be a grade penalty.

Weekly Schedule

Week 1: Introduction

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/>

Elizabeth M. Renieris, David Kiron, and Steven Mills, *To Be a Responsible AI Leader, Focus on Being Responsible*," MIT Sloan Management Review and Boston Consulting Group, September 2022. <https://sloanreview.mit.edu/projects/to-be-a-responsible-ai-leader-focus-on-being-responsible/>

Week 2: Frameworks

OSTP, Blueprint for an AI Bill of Rights, <https://www.whitehouse.gov/wp-content/uploads/2022/10/Blueprint-for-an-AI-Bill-of-Rights.pdf>

Concrete Problems in AI Safety Interview Transcript, <https://futureoflife.org/uncategorized/transcript-concrete-problems-ai-safety-dario-amodei-seth-baum/#:~:text=There%20are%20five%20concrete%20problems,and%20robustness%20to%20distributional%20shift> (original academic paper is available here: <https://arxiv.org/abs/1606.06565>.)

Sayash Kapoor and Arvind Narayanan, *Eighteen pitfalls to beware of in AI journalism, A checklist for avoiding hype*, <https://aisnakeoil.substack.com/p/eighteen-pitfalls-to-beware-of-in>

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>.

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

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: Consumer Protection

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

FTC, *Aiming for truth, fairness, and equity in your company's use of AI*,
<https://www.ftc.gov/business-guidance/blog/2021/04/aiming-truth-fairness-equity-your-companys-use-ai>

FTC *Dark Patterns Report*,
https://www.ftc.gov/system/files/ftc_gov/pdf/P214800%20Dark%20Patterns%20Report%209.14.2022%20-%20FINAL.pdf

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

Privacy and human behavior in the age of information, Alessandro Acquisti, Laura Brandimarte, and George Loewenstein, Science 2015,
<https://www.science.org/doi/10.1126/science.aaa1465>

A Contextual Approach to Privacy Online, Helen Nissenbaum, Daedalus (2011),
<https://www.amacad.org/publication/contextual-approach-privacy-online>.

Congressional Research Service, *Overview of the American Data Privacy and Protection Act*, H.R. 8152, <https://crsreports.congress.gov/product/pdf/LSB/LSB10776>

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

Brent Mittelstadt, *Principles alone cannot guarantee ethical AI*,
<https://www.nature.com/articles/s42256-019-0114-4> (Nature Machine Intelligence, 2019)

Rumman Chowdhury, Bogdana Rakova, Henriette Cramer, and Jingying Yang, *Putting Responsible AI Into Practice*, (October 22, 2020),
<https://sloanreview.mit.edu/article/putting-responsible-ai-into-practice/>

David Robinson, *The Kidney Transplant Algorithm's Surprising Lessons for Ethical A.I.*,
<https://slate.com/technology/2022/08/kidney-allocation-algorithm-ai-ethics.html>

Optional Additional Materials:

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

Black Mirror on Netflix

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