



LGST 242x/642x

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

Fall 2018

Overview

A world of ubiquitous data, subject to ever more sophisticated collection, aggregation, analysis, and use, creates massive opportunities for both financial gain and social good. It also creates dangers such as privacy violations and discrimination, as well as simple hubris about the effectiveness of management by algorithm. This course introduces students to the legal, policy, and ethical dimensions of big data, predictive analytics, and related techniques.

Instructor

Professor Kevin Werbach

Department of Legal Studies and Business Ethics

673 Huntsman Hall. | (215) 898-1222 | werbach@wharton.upenn.edu (best way to reach me)

Office Hours: Wednesday 2-3pm, or by appointment

Course Requirements and Grading

Your grade is comprised of four components: Reading Reactions (20%), a final Algorithmic Accountability Report (40%), Mock Trial materials (20%), and Participation (20%).

Reading Reactions (20%)

You must submit five short response papers (150-200 words) or videos (2-4 minutes in length). These are due on Canvas by 6pm Monday prior to our Tuesday sessions. Identify one argument or theme from the upcoming readings for the week; summarize it; and explain whether you agree or disagree. Your submission should demonstrate that you read and evaluated the materials. If you submit a video, it should reflect comparable thought and organization to a written submission.

The assignments will be graded for completion. Except in rare instances, I will not comment on your submission. Acceptable submissions receive 4 points; half-hearted efforts receive 2 points. You may submit Reading Reactions individually or as a pair with one other student. If you do so as a pair, both students should submit the file to Canvas. You must clearly identify that this was joint work, and the submission should demonstrate that both students materially contributed.

There are seven weeks of class after the first session, so you may choose two weeks to skip the assignment. If you do not receive the full points, you may submit a sixth or seventh Reading

Reaction and receive credit for the highest five scores. Because of this flexibility, no late submissions will be accepted, except in extraordinary cases of illness or personal emergency.

Algorithmic Accountability Report (40%)

The largest component of your grade is due the weekend following the final class session. You must create an “algorithmic accountability report” directed to a specific company or organization. It may be a company discussed in the course, other than Google or Facebook (because they come up so frequently). Or it may be another organization of your choosing.

Your report should do two things:

1. Identify the most significant areas of concern regarding the company’s use of big data, business analytics, machine learning, or similar techniques. Describe specific problems that either have occurred or may occur in the future. Explain why they arise.
2. Make one or more concrete recommendations to the company. Explain the rationale and implementation for each recommendation.

You may choose the format for your submission. It could be a paper or memo (1500-2500 words), a slide deck, a screencast with audio narration, a video, an animation, a blog post, or another way to frame your analysis. (If you have an idea, feel free to check with me before starting.) Regardless of the format, reports will be graded based on quality of presentation, depth of analysis, persuasiveness, organization, research beyond course materials, use of course concepts, and originality. On-time submissions will receive between 25 and 40 points. Late submissions will be marked down 3 points per day until grades are submitted, after which no further submissions will be accepted. Note: I will be running each report through anti-plagiarism software. Plagiarized work will result in severe consequence, per University of Pennsylvania rules.

Mock Trial (20%)

We will do a mock trial during class #8. The in-class exercise itself is not graded. You must submit a pre-trial prep sheet, worth 5 points, which will be graded for completion, using a similar standard to the Reading Responses. After the class, you must submit a post-trial reflection, worth 15 points, that discusses how the exercise illustrated themes or concepts from the course. On-time submissions will receive a score between 8 and 15 based on the quality of analysis and integration of course concepts and materials. Late submissions will be marked down one point per day.

Participation (20%)

You will be assessed on the overall quality of your contributions to the course. Attendance is one factor, but not the primary one. I will record attendance and a participation score each day. If you arrive late, it is your responsibility to let me know after class. There is no way to “make up” an unexcused missed class, but active participation other times will counterbalance the occasional absence. I realize students sometimes miss classes for understandable reasons such as recruiting. All class sessions will be video recorded and available to you on Canvas.

I know that some students are less comfortable speaking in class, or may simply not get called on when they attempt to. Substantive posts to the Canvas discussion boards also count toward the Participation grade. You will receive credit both for comments on class discussions, and for bringing in relevant ideas or materials you encounter outside the class.

Classroom Expectations and Participation

- Please arrive on time. I will start and end the class promptly.
- Sit according to the seating chart, and bring a name tent and display it for each class. (Cards will be provided at the start of the course for undergraduates.)
- Turn off and put away all electronic devices.
- Come prepared to contribute to the class discussion.

This is a discussion-oriented course. I will lecture at times to introduce important concepts, but the heart of most sessions will be conversations. There usually isn't a "right answer," although there are more and less thoughtful contributions. I use a seating chart, and ask you to bring name tents, so that I can learn your names and note your unique perspectives. We have only 14 sessions and I'm teaching two other classes this quarter, so I won't be perfect at that, but I'll do my best.

This class will observe the MBA Program's policy on student use of electronic devices in the classroom. No phones, tablets, or laptops may be used during the class, even for note-taking, unless specifically authorized for an in-class activity. Violations of the rules will be reflected in the class participation aspect of the course grade. As a technology scholar, I was reluctant to ban devices in my classes. It's my job to create an engaging experience that is worth paying attention to. However, it's also my job to establish a [good learning environment](#). Researchers have [found](#) that when students use laptops in class, it decreases their retention *and that of students sitting nearby*. [Studies](#) also show that you learn more when taking notes by hand than on a device.

Learning Objectives

Good data-driven decision-making means not just generating solutions, but understanding how to use them. Sophisticated firms in terms of data science expertise have already gotten into trouble over privacy, security, manipulation, and discrimination. Failure to anticipate such issues can result in ethical lapses, public relations disasters, regulatory sanctions, and even legal liability.

My goal is to help you develop the skills to use analytics in a responsible way, while remaining focused on your business objectives. After completion of the course, you should be able to:

1. Identify where algorithms depend on human judgments or assumptions.
2. Describe legal rules and regulatory obligations, in the U.S. and elsewhere, that may apply to business analytics.
3. Evaluate claims that applications of analytics raise ethical or public policy concerns.
4. Appreciate the perspectives of multiple actors on controversies about privacy, manipulation, and algorithmic bias.
5. Develop thoughtful responses to concerns about the uses of data science.
6. After graduation from Wharton, don't destroy the world, crash the economy, go to jail, or all of the above. (Money-back guarantee not available.)

Syllabus

Unless otherwise noted, hyperlinks are provided below to online versions of all the readings. Where there are questions listed, be prepared to address them in class discussion.

1. THE PROMISE AND THE PERIL

How might data science change the relationships among firms, customers, employees, other firms, and governments? What are some of the legal or ethical concerns that may arise?

Jacob LaReviere et al, [Where Predictive Analytics Is Having the Biggest Impact](#), Harvard Business Review, May 25, 2016

- What are the main business value propositions for analytics and big data?

Will Oremus, [Move Fast and Break Trust](#), Slate, March 7, 2017

- What do the problems with Google's smart speaker and Uber's autonomous driving system have in common?

Optional

Steven Finlay, Predictive Analytics, Data Mining, and Big Data: Myths, Misconceptions, and Methods (2014), chapters 1-2 [available on Study.net]

- A more-detailed introduction to the methods of business analytics, for students without a background in the area.

2. NEUTRALITY

Algorithms rely on human decisions about how data are collected, analyzed, and used. Failure to appreciate this can lead to problems.

Pedro Domingos, [A Few Useful Things to Know about Machine Learning](#), Communications of the ACM, October 2012

- What are the ways that machine learning requires human judgement or intervention?

Zeynep Tufekci, [The Real Bias Built In at Facebook](#), N.Y. Times, May 19, 2016

- Why was Facebook criticized for its Trending Topics?
- Why does Tufekci say that algorithms are not neutral? What does that even mean?
- If algorithms are inherently biased, does that undermine the value of analytics in business?

Zeynep Tufekci, [YouTube, the Great Radicalizer](#), N.Y. Times, March 10, 2018

- Why does YouTube push users to extreme content?
- Should Google do something about it?

3. ACCURACY

The first step to responsible use of analytics is to appreciate limitations of its methods.

Gary Marcus & Ernest Davis, [8 \(No, 9!\) Problems with Big Data](#), N.Y. Times, April 6, 2014

- What are some of the common themes in the authors' list of problems?

David Lazer et al, [The Parable of Google Flu](#), Science, March 14, 2014

- Why was Google Flu Trends so accurate initially, and not subsequently?
- Should the failure of Google Flu make us skeptical about business analytics?

Michael Luca et al, [Algorithms Need Managers, Too](#), Harvard Business Review, January-February 2016

- What is the role for business managers in overseeing the use of analytics?

4. RISK

Algorithmic systems may produce unintended results, which in some cases cause harm. How can organizations using analytics manage those risks?

Andrew Smith, [Franken-Algorithms: The Deadly Consequences of Unpredictable Code](#), The Guardian, August 30, 2014.

- Was Uber irresponsible in putting an autonomous vehicle on a public road?
- How do business incentives and engineering standards each contribute to the dangers of algorithmic systems?

Sapna Maheshwari and Alexandra Stevenson, [Google and Facebook Face Criticism for Ads Targeting Racist Sentiments](#), New York Times, September 15, 2017

- If you were a manager at Google or Facebook responsible for advertising systems, what would you do in light of the controversies described in this article?

Christian Sandvig et al, [An Algorithm Audit](#) (2014)

- How effective do you think the audit method described in the paper can be?

5. TRANSPARENCY

How well can we assess exactly what algorithms are doing, and why?

[Houston Federation of Teachers v. Houston Ind. School District](#) (S.D. Texas, May 4, 2017)

- On what legal basis did the teachers challenge the value added measures system?
- Who won the case, and why?

Cliff Kuang, [Can A.I. Be Taught to Explain Itself?](#), New York Times, November 21, 2017

- Can “explainable AI” techniques address the dangers of analytics?

Kate Crawford and Jason Schultz, [Big Data and Due Process: Toward a Framework to Redress Predictive Privacy Harms](#), 55 Boston College Law Review pp. 93-99, 121-28 (2014)

- How is the “due process” approach different from existing laws we’ve discussed?
- Will this approach be effective?

Julia Powles, [New York City’s Bold, Flawed Attempt to Make Algorithms Accountable](#), The New Yorker, December 20, 2017

- Would the proposed New York algorithmic disclosure mandate be beneficial?
- Should there be similar disclosure when users are private companies or individuals?

6. FAIRNESS

The use of analytics has the potential both to counteract and to reinforce systematic biases.

Kate Crawford, [The Hidden Biases in Big Data](#), Harvard Business Review, April 1, 2013

- What did the services the author describes do wrong?

Solon Barocas & Andrew Selbst, [Big Data's Disparate Impact](#), California Law Review (2015), pp. 677-93

- How can big data can produce results that seem biased against certain populations?

Suresh Venkatesubramanian, [Algorithmic Fairness: From Social Good to a Mathematical Framework](#), LSE Media Policy Project Blog, June 14, 2016

- Can technical solutions such as those described by the author be effective?

Alex Miller, [Want Less-Biased Decisions? Use Algorithms](#), Harvard Business Review, July 26, 2018

- So there's nothing to worry about?

7. DISCRIMINATION

When are differential effects of analytics tantamount to illegitimate or illegal discrimination?

[Ricci v. DeStefano](#), 557 U.S. 557 (2009)

- How did the New Haven fire department respond when it found that white candidates did better on its promotion test?
- How did the court rule on the legal challenge to the fire department's actions, and why?

[Texas Dept. of Housing and Community Affairs v. Inclusive Communities Project](#) (2015)

- How does the court respond to statistical evidence that low-income housing tax credits are offered primarily in non-white areas, arguably worsening segregation?

Julia Angwin et al, [Machine Bias](#), ProPublica, May 23, 2016

- Does the ProPublica report demonstrate unfair or discriminatory outcomes from the use of the COMPAS system for sentencing?
- What might explain the racial variations the researchers found?

8. BUSINESS ANALYTICS ON TRIAL

Based on a "ripped from the headlines" episode of a TV drama, we'll act out a realistic scenario of alleged algorithmic discrimination

Watch the "Good Wife" video segment on the Canvas site.

Prepare to assume your pre-assigned role in a mock trial.

9. DATA COLLECTION AND AGGREGATION

Are there limits on how data should be collected, used, and shared?

Kashmir Hill and Surya Mattu, [How a Company You've Never Heard of Sends You Letters about Your Medical Condition](#), Gizmodo, June 19, 2017

- How does Acurian obtain seemingly private medical information?

Solon Barocas and Helen Nissenbaum, [Big Data's End Run Around Procedural Privacy Protections](#), Communications of the ACM (November 2014)

- What do the authors believe that transparency and consent are insufficient?

The White House, [Big Data: Seizing Opportunities, Preserving Values](#) (2014), pp. 15-21

- What are the key elements of the U.S. approach to privacy law?
- Do you think the U.S. legal framework is effective in general? Will it be effective for the novel challenges of big data and business analytics?

Izaak Crook, [How GDPR Will Affect Data Science](#), Dataconomy.com, April 13, 2018

- How will the European General Data Protection Regulation affect the way companies use analytics?

10. PERILS OF PREDICTION

If sensitive attributes can be inferred from other data, does it even make sense to talk about privacy any more?

Charles Duhigg, [How Companies Learn Your Secrets](#), N.Y. Times Magazine, Feb. 16, 2012

- How does Target analyze customer data to make inferences about customers?
- In your opinion, is the Target system an intrusion on privacy? Why or why not?
- Do Target's actions violate any legal rules?
- Do Target's actions violate any ethical norms?
- Should Target do anything differently?

Michal Kosinski et al, [Private Traits and Attributes are Predictable from Digital Records of Human Behavior](#), Proceedings of the National Academy of Sciences, April 9, 2013

- What kinds of information can be predicted based on Facebook Likes?
- What could possibly go wrong?

11. MANIPULATION

To what extent does analysis itself influence behavior? And what are the limits on using analytics not merely to understand and predict customer actions, but to shape them?

Zeynep Tufekci, [Algorithmic Harms Beyond Facebook and Google: Emergent Challenges of Computational Agency](#), Journal on Telecom. and High-Tech Law (2015), pp. 203-209

- What was Facebook trying to achieve in its emotional contagion study?
- Why were Facebook's actions controversial?
- What is "algorithmic gatekeeping"? Why does Tufekci believe it is a concern?

Rebecca Rosen, [Is This the Grossest Advertising Strategy of All Time?](#), Atlantic, October 3, 2013

- What does the author find new and objectionable about this marketing approach?

Paul Armstrong, [Facebook is Helping Brands Target Teens Who Feel “Worthless.”](#) Forbes, May 1, 2017

- Do you feel any differently about Facebook’s actions than those of the consulting firm in the previous article?

12. MARKET POWER

Should we be concerned about algorithmic monopolies or other anti-competitive practices? And what about those left out by the big data revolution?

Jerry Useem, [How Online Shopping Makes Suckers of Us All](#), Atlantic Monthly, May 2017

- Do you find algorithmic pricing practices troubling?

Ryan Calo, [Digital Market Manipulation](#), 82 George Washington Law Review 995 (2014), pp. 1003-1012, 1020-24

- How do online intermediaries, in Calo’s account, engage in forms of manipulation?

Maurice Stucke, [Here Are All the Reasons It’s a Bad Idea to Let a Few Tech Companies Monopolize Our Data](#), Harvard Business Review, March 27, 2018

- Should there be more aggressive antitrust enforcement against companies exploiting big data and analytics for market power?

13. UTOPIA OR DYSTOPIA?

China’s Social Credit System is the most fully-realized effort to make algorithmic decision-making ubiquitous in society.

Mara Hvistendahl, [Inside China’s Vast Experiment in Social Ranking](#), Wired, Dec. 14, 2017

- What are the major elements of China’s system?

14. ALGORITHMIC ACCOUNTABILITY

How can firms best respond to the challenges we’ve discussed in the course?

[Interview with Mark Van Hollebeke](#), Microsoft (video)

- How does Microsoft address the kinds privacy concerns we’ve discussed in the class

Bethan Cantrell, et al, [Industry Needs to Embrace Data Ethics: How It Could Be Done](#) (2016)

- What do you think of Microsoft’s data ethics principles?

Darrell West, [The Role of Corporations in Addressing AI’s Ethical Dilemmas](#), Brookings, September 13, 2018

- Which of West’s recommendations are relevant to companies other than the big technology platforms like Google and Facebook?

Instructor Bio

Professor Kevin Werbach is an expert on legal, business, and policy implications of emerging technologies such as broadband, big data, gamification, and blockchain. He served on the Obama Administration's Presidential Transition Team, founded the Supernova Group (a technology conference and consulting firm), helped develop the U.S. government's approach to internet policy during the Clinton Administration, and created one of the most successful massive open online courses, with over 450,000 enrollments. His books include *For the Win: How Game Thinking Can Revolutionize Your Business* and *The Blockchain and the New Architecture of Trust*. In 2010, he was named Wharton's first-ever "Iron Prof".