



LGST 242/642

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

Fall 2020

---

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.

*Note: This is a fast-evolving subject area. I revise the course frequently. Make sure to check Canvas for the most up-to-date readings and materials.*

## Instructor

Professor Kevin Werbach  
Department of Legal Studies & Business Ethics  
[werbach@wharton.upenn.edu](mailto:werbach@wharton.upenn.edu)

## 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 the limitations and flaws of algorithmic decision-making systems.
2. Anticipate legal or ethical controversies arising from applications of business analytics.
3. Evaluate mechanisms to promote algorithmic accountability from a variety of perspectives.
4. After graduation, don't destroy the world, crash the economy, or go to jail. (*Money-back guarantee not available.*)

## Course Structure

This is a fast-paced course. It is only 14 sessions long, and there are a number of assignments you need to complete. These are part of the learning experience, not just tests. I have designed in significant flexibility to make the workload is manageable, but it's your responsibility to stay on top of the requirements.

I am using a flipped hybrid course model. For most class session, I've recorded several short videos. This is the "lecture" component, to complete on your own prior to class, with embedded quiz questions to check your comprehension..

Class time is for interaction—live discussions, groupwork, and activities. The synchronous class sessions (via Zoom) will be shortened to 50 minutes, because you will have already watched the lecture portions. Every session (with associated videos and activities) is a Module on the Canvas site. The synchronous sessions will be recorded, and there will be alternate opportunities for those unable to attend live.

For every session, there is a Canvas discussion board. I encourage you to use these as a means of class participation. I will contribute to the discussion boards as well from time to time. In addition to comments based on course materials and sessions, highlighting recent events and articles relevant to the course topics is strongly encouraged. Substantive responses to posts by other students are encouraged, and may even earn you more participation credit than brief original posts. Participation credit will be assessed holistically: quality matters, not just quantity. Adding 10 comments at the end of the quarter will not be viewed as favorably as engaging when topics are highlighted in class or through current events.

## Course Requirements and Grading

### Lecture Video Quizzes [10%]

There will be multiple-choice questions associated with the pre-recorded lecture videos, to gauge your comprehension.

### Black Mirror Reflections [20%]

For most class sessions, I have selected an episode of the dystopian TV series Black Mirror, which is available on Netflix. (If you do not have a subscription, you can get a free trial for one month.) I recommend you watch the episode after doing the readings and watching the videos for the relevant class session. Note that episode lengths vary; some are almost two hours.

In 150-400 words, identify connections between the episode and one or more of the assigned readings or videos for the associated class. Please note: the assignment is to focus on the materials for that day's upcoming class. If you wish, you are welcome to draw connections with earlier sessions, or to general themes of the course, once you have done so. Be specific regarding what parts of the readings or videos you are referencing.

You only need to submit five Black Mirror Reflections. Which ones you select are entirely up to you. The Reflections will be peer graded by your fellow students. Part of the credit you receive for the assignment will be for completing peer assessments. If you submit more than five Black Mirror Reflections, you will receive one additional point toward your final grade for each one that earns a score of 4/4.

Each Black Mirror episode involves a slightly futuristic scenario that crystallizes the kinds of issues we discuss in the course. Although the series is fictional, students in the past found this exercise very helpful for understanding what's at stake with the technologies we examine. Be aware that some episodes are disturbing, include sexual situations, or involve violence.

There are multiple possible connection points with each episode. Your submission should demonstrate that you reviewed and evaluated the materials -- both the show and the readings. Formal citations are not required, but you should make clear which readings/videos you're talking about.

#### Strategic Algorithmic Audit [40%]

As a final project, you must create a report directed to a specific company or organization. It may be a company discussed in the course, other than Google, Amazon, or Facebook (because they come up so frequently). Preferably, though, it should be another organization of your choosing.

Your report should do two things:

1. Identify 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 that may occur in the future. Explain why they did, or are likely to, arise.
2. Make concrete recommendations to the company. Explain the rationale and implementation for each recommendation. Your recommendations should include at least one technical mechanism, one reference to a legal or regulatory requirement, and one operational response describing practices the company should adopt.

You may choose the format for your submission. It could be a paper or memo (1500-3000 words), a slide deck, a screencast with audio narration, a video, an animation, a blog post, or another way to frame your analysis that demonstrates a similar level of effort and depth of 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, use of course concepts, research beyond course materials, and originality.

Your paper should be specific to the company selected, and should illustrate how the course helped you identify problems and solutions. Note: I will be running reports through anti-plagiarism software. Plagiarized work will result in severe consequence, per University of Pennsylvania rules.

#### Mock Trial [15%]

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 Reactions. After the class, you must submit a post-trial reflection, worth 10 points, that discusses how the exercise illustrated themes or concepts from the course.

#### Participation [15%]

Participation is an important aspect of this course. It is not just a matter of memorizing information; I want you to engage with these novel issues and form your own conclusions. I realize that the current environment of social distancing, and the various challenges associated with the Covid-19 pandemic, will make participation difficult at times. Students may not be able to attend synchronous sessions, and even for those who do, it may be more challenging to engage in class discussion.

Therefore, I will provide a range of both synchronous and asynchronous participation opportunities. Those unable to participate synchronously will not be penalized, but you must devote effort to engaging with the course through other means. You will be assessed based on your overall contributions to the course.

## Deadlines

I will always accept late submissions, up to the final class session of the course, with a graduated markdown. Late submission of the final paper will similarly incur a graduated markdown until grades are submitted. If you encounter difficulty completing an assignment in a timely manner, please contact me before the deadline. I understand you may be facing unusual challenges in the current environment. I am always willing to listen and consider accommodations. Maximizing your learning is my primary goal.

## 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.*

| <b>Class</b><br><i>Dates listed for both LGST 242 (Q1) and LGST 642 (Q2)</i>  | <b>Assignments</b>   |
|---|--|
| September 2   October 26<br><br><b>I, FOR ONE, WELCOME OUR NEW ROBOT OVERLORDS</b><br>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? | <u>Read</u><br><br>Jacob LaReviere et al, <a href="#">Where Predictive Analytics Is Having the Biggest Impact</a> , Harvard Business Review, May 25, 2016 <ul style="list-style-type: none"> <li>• What are the main business value propositions for analytics and big data?</li> </ul><br>Matthew Hutson, <a href="#">AI Could ID Gang Crimes. Opening an Ethical Firestorm</a> , Science, February 28, 2018 <ul style="list-style-type: none"> <li>• Should this system be implemented?</li> </ul>   |
| September 9   October 28<br><br><b>ALGORITHMIC DECISIONMAKING</b><br>Algorithms rely on human decisions about how data are collected, analyzed, and used. Failure to appreciate this can lead to problems.  | <u>Black Mirror</u> : Entire History of You (series 1 episode 3)<br><i>*Reaction paper may cover first session readings</i><br><br><u>Read</u><br><br>Pedro Domingos, <a href="#">A Few Useful Things to Know about Machine Learning</a> , Communications of the ACM, October 2012 <ul style="list-style-type: none"> <li>• What are the ways that machine learning requires human judgement or intervention?</li> </ul><br>Graham Ruddick, <a href="#">Admiral to Price Car Insurance Based on Facebook Posts</a> , The Guardian, Nov. 1, 2016 <ul style="list-style-type: none"> <li>• Would you feel comfortable using the Admiral system for your car insurance?</li> </ul><br>Zeynep Tufekci, <a href="#">YouTube, the Great Radicalizer</a> , N.Y. Times, March 10, 2018 <ul style="list-style-type: none"> <li>• Why does YouTube push users to extreme content?</li> <li>• Should Google do something about it? Can it?</li> </ul> |
| September 14   November 2<br><br><b>ACCURACY</b><br>The first step to responsible use of analytics is to appreciate limitations of its  | <u>Black Mirror</u> : Be Right Back (Series 2 episode 1)<br><br><u>Read</u><br><br>Gary Marcus & Ernest Davis, <a href="#">8 (No, 9!) Problems with Big Data</a> , N.Y. Times, April 6, 2014   |

|   |  |
|---|--|
| <p>methods. Algorithmic decision-making is powerful, but not always effective or robust to changes.</p>   | <ul style="list-style-type: none"> <li>• What are some of the common themes in the authors' list of problems?</li> </ul> <p>David Lazer et al, <a href="#">The Parable of Google Flu</a>, Science, March 14, 2014</p> <ul style="list-style-type: none"> <li>• Why was Google Flu Trends so accurate initially, and not subsequently?</li> <li>• Should the failure of Google Flu make us skeptical about business analytics?</li> </ul> <p>Kashmir Hill, <a href="#">Wrongfully Accused by an Algorithm</a>, NY Times, June 24, 2020</p> <ul style="list-style-type: none"> <li>• Is the face recognition system used in Detroit a "racist algorithm"?</li> <li>• What exactly went wrong that led to Robert Williams' mistaken arrest?</li> </ul> <p>Michael Luca et al, <a href="#">Algorithms Need Managers, Too</a>, Harvard Business Review, January-February 2016</p> <ul style="list-style-type: none"> <li>• What is the role for business managers in overseeing the use of analytics?</li> </ul>  |
| <p>September 16   November 4</p> <p><b>TRANSPARENCY</b><br/>How well can we assess exactly what algorithms are doing, and why?</p>  | <p><u>Black Mirror</u>: Hang the DJ (Series 4 episode 4)</p> <p><u>Read</u></p> <p><a href="#">Houston Federation of Teachers v. Houston Ind. School District</a> (S.D. Texas, May 4, 2017) – <i>edited version</i></p> <ul style="list-style-type: none"> <li>• On what legal basis did the teachers challenge the value added measures system?</li> <li>• Who won the case, and why?</li> </ul> <p>Jenna Burrell, <a href="#">How the Machine 'Thinks': Understanding Opacity in Machine Learning Algorithms</a>, Big Data &amp; Society, January-June, 2016</p> <ul style="list-style-type: none"> <li>• Why is it difficult to identify the bases for recommendations generated by machine learning?</li> </ul> <p>Julia Powles, <a href="#">New York City's Bold, Flawed Attempt to Make Algorithms Accountable</a>, New Yorker, December 20, 2017</p> <ul style="list-style-type: none"> <li>• Would the proposed New York algorithmic disclosure mandate be beneficial?</li> <li>• Should there be similar disclosure when users are private companies or individuals?</li> </ul> |
| <p>September 21   November 9</p> <p><b>RISK AND RESPONSIBILITY</b><br/>Algorithmic systems may produce unintended results, which in some cases cause harm. Who should be held</p> | <p><u>Black Mirror</u>: Hated in the Nation (Series 3 episode 6)</p> <p><u>Read</u></p> <p>National Transportation Safety Board, <a href="#">Preliminary Report</a>, May 24, 2018</p>  |

|  |   |
|--|---|
| <p>liable? How can organizations using analytics manage those risks?</p>   | <ul style="list-style-type: none"> <li>Based on these findings, who if anyone should be held responsible for the death of Elaine Herzberg, the pedestrian struck by the Uber self-driving car?</li> <li>How could the risk of future accidents be mitigated?</li> </ul> <p>National Transportation Safety Board, <a href="#">Preliminary Report</a>, May 24, 2018</p> <ul style="list-style-type: none"> <li>Based on these findings, who if anyone should be held responsible for the death of Elaine Herzberg, the pedestrian struck by the Uber self-driving car?</li> <li>How could the risk of future accidents be mitigated?</li> </ul> <p>Karen Hao, <a href="#">When Algorithms Mess Up, the Nearest Human Gets the Blame</a>, MIT Technology Review, May 28, 2019</p> <ul style="list-style-type: none"> <li>What is the concept of “moral crumple zones”? How does it relate to responsibility for harms involving algorithmic systems?</li> </ul> <p>Future of Privacy Forum, <a href="#">Beyond Explainability: A Practical Guide to Managing Risk in Machine Learning Models</a> (June 2018)</p> <ul style="list-style-type: none"> <li>Are these steps sufficient to address the risks of algorithmic decision-making?</li> <li>Is it realistic to expect companies to adopt such processes?</li> </ul> |
| <p>September 23   November 11</p> <p><b>FAIRNESS</b><br/>The use of analytics has the potential both to counteract and to reinforce systematic biases. But what exactly does it mean for an algorithmic system to be “fair”?</p> | <p><u>Black Mirror</u>: Men Against Fire (Season 3, Episode 5)</p> <p><u>Read</u></p> <p>Kate Crawford, <a href="#">The Hidden Biases in Big Data</a>, Harvard Business Review, April 1, 2013</p> <ul style="list-style-type: none"> <li>What did the services the author describes do wrong?</li> </ul> <p>Harini Suresh and John Guttag, <a href="#">A Framework for Understanding Unintended Consequences of Machine Learning</a> (January 2019)</p> <ul style="list-style-type: none"> <li>What are the different forms of fairness the authors describe?</li> <li>Is it possible to build a fair system using machine learning?</li> </ul> <p>Julia Angwin et al, <a href="#">Machine Bias</a>, ProPublica, May 23, 2016</p> <ul style="list-style-type: none"> <li>Does the ProPublica report demonstrate unfair outcomes from the use of the COMPAS system for sentencing?</li> <li>What might explain the racial variations the researchers found?</li> </ul>   |
| <p>September 28   November 16</p> <p><b>DISCRIMINATION</b><br/>When are differential effects of analytics tantamount to illegitimate or illegal discrimination?</p>  | <p><u>Black Mirror</u>: The Black Museum (Season 4, Episode 6)</p> <p><u>Read</u></p> <p>Aaron Klein, <a href="#">Credit Denial in the Age of AI</a>, Brookings.com, April 11, 2019</p> <ul style="list-style-type: none"> <li>How can algorithms potentially create legally actionable discrimination in credit decisions?</li> </ul>  |

|   |   |
|---|---|
|   | <p><a href="#">Ricci v. DeStefano</a>, 557 U.S. 557 (2009) – <i>edited version</i></p> <ul style="list-style-type: none"> <li>• How did the New Haven fire department respond when it found that white candidates did better on its promotion test?</li> <li>• How did the court rule on the legal challenge to the fire department’s actions, and why?</li> </ul> <p><a href="#">Texas Dept. of Housing and Community Affairs v. Inclusive Communities Project</a> (2015) – <i>edited version</i></p> <ul style="list-style-type: none"> <li>• How does the court respond to statistical evidence that low-income housing tax credits are offered primarily in non-white areas, arguably worsening segregation?</li> </ul> <p>Ruha Benjamin, <a href="#">Assessing Risk, Automating Racism</a>, Science, October 25, 2019</p> <ul style="list-style-type: none"> <li>• Should the differential treatment caused by the algorithm described in the article be legally prohibited? how?</li> </ul> |
| <p>September 30   November 18</p> <p><b>BUSINESS ANALYTICS ON TRIAL</b></p> <p>Based on a “ripped from the headlines” episode of a TV drama, we’ll act out a realistic scenario of alleged algorithmic discrimination</p> | <p>Watch the “Good Wife” <a href="#">video segment</a> on the Canvas site.</p> <ul style="list-style-type: none"> <li>• Prepare to assume your pre-assigned role in a mock trial.</li> </ul> <p><u>Read:</u></p> <p><a href="#">Legal Standard for Mock Trial</a></p> <p>Brian Clifton et al, <a href="#">Predictive Policing for White-Collar Crime</a></p> <ul style="list-style-type: none"> <li>• Should the system proposed in this paper be deployed widely?</li> </ul>   |
| <p>October 5</p> <p><b>DATA COLLECTION AND AGGREGATION</b></p> <p>Are there limits on how data should be collected, used, and shared?</p>   | <p><a href="#">Black Mirror</a>: Arkangel (Series 4 episode 2)</p> <p><u>Read</u></p> <p>Solon Barocas and Helen Nissenbaum, <a href="#">Big Data’s End Run Around Procedural Privacy Protections</a>, Communications of the ACM (November 2014)</p> <ul style="list-style-type: none"> <li>• Why do the authors believe that transparency and consent are insufficient?</li> </ul> <p>Kashmir Hill, <a href="#">The Secretive Company that Might End Privacy as We know It</a>, N.Y. Time, January 18, 2020</p> <ul style="list-style-type: none"> <li>• Do you find Clearview AI troubling? Is it a more serious threat to privacy than other existing systems?</li> </ul> <p>Kashmir Hill and Surya Mattu, <a href="#">How a Company You’ve Never Heard of Sends You Letters about Your Medical Condition</a>, Gizmodo, June 19, 2017</p> <ul style="list-style-type: none"> <li>• How can Acurian obtain seemingly private medical info?</li> </ul>   |
| <p>October 7   November 23</p>  | <p><a href="#">Black Mirror</a>: Shut Up and Dance (Series 3, Episode 3)</p>  |

|   |   |
|---|---|
| <p><b>INFERENCEAL PRIVACY</b><br/>If sensitive attributes can be inferred from other data, does it even make sense to talk about privacy any more?</p>  | <p><u>Read</u></p> <p>Charles Duhigg, <a href="#">How Companies Learn Your Secrets</a>, N.Y. Times Magazine, Feb. 16, 2012</p> <ul style="list-style-type: none"> <li>• How does Target analyze customer data to make inferences about customers?</li> <li>• In your opinion, is the Target system an intrusion on privacy? Why or why not?</li> <li>• Do Target’s actions violate any legal rules?</li> <li>• Do Target’s actions violate any ethical norms?</li> <li>• Should Target do anything differently?</li> </ul> <p>Nuala O’Connor, <a href="#">Reforming the U.S. Approach to Data Protection and Privacy</a>, (2018)</p> <ul style="list-style-type: none"> <li>• What are the limitations of the current U.S. legal approach to privacy? What are the benefits?</li> <li>• How does O’Connor suggest U.S. law should change?</li> </ul> <p>GDPR.eu, <a href="#">What Is GDPR, the EU’s New Data Protection Law?</a></p> <ul style="list-style-type: none"> <li>• Do you find the European GDPR approach superior to the American one?</li> </ul> |
| <p>October 12   November 30</p> <p><b>MANIPULATION</b><br/>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?</p> | <p><u>Black Mirror</u>: White Christmas (2014 special)</p> <p><u>Read</u></p> <p>Zeynep Tufekci, <a href="#">Algorithmic Harms Beyond Facebook and Google: Emergent Challenges of Computational Agency</a>, Journal on Telecom. and High-Tech Law (2015), pp. 203-209</p> <ul style="list-style-type: none"> <li>• Why was Facebook’s emotional contagion study controversial?</li> <li>• What is “algorithmic gatekeeping”? Why does Tufekci believe it is a concern?</li> </ul> <p>Parmy Olson, <a href="#">For \$29, This Man Will Help Manipulate Your Loved Ones With Targeted Facebook And Browser Links</a>, Forbes (2019)</p> <ul style="list-style-type: none"> <li>• Is The Spinner just creepy, or is it the kind of thing we should have ethical or legal qualms about?</li> </ul>  |
| <p>October 14   December 2</p> <p><b>MARKET POWER</b><br/>Should we be concerned about algorithmic monopolies or other anti-competitive practices?</p>  | <p><u>Black Mirror</u>: 15 Million Merits</p> <p><u>Read</u></p> <p>Jerry Useem, <a href="#">How Online Shopping Makes Suckers of Us All</a>, Atlantic Monthly, May 2017</p> <ul style="list-style-type: none"> <li>• Do you find algorithmic pricing practices troubling?</li> </ul> <p>Data Mattioli, <a href="#">Amazon Changed Search Algorithm in Ways That Boost Its Own Products</a>, Wall St. Journal, September 16, 2019</p>   |



|   |  |
|---|--|
|   | <ul style="list-style-type: none"> <li>• What are the best arguments for and against the changes Amazon made to its search algorithm?</li> </ul> <p>Helen Coffey, <a href="#">Airlines Face Crack Down on Use of "Exploitative" Algorithm That Splits Up Families on Flights</a>, The Independent, November 19, 2018</p> <ul style="list-style-type: none"> <li>• Is it right to describe the airline practice in question an "exploitative algorithm"?</li> </ul>   |
| <p>October 19   December 7</p> <p><b>THE SCORED SOCIETY</b><br/>China's Social Credit System is the most ambitious effort to make algorithmic decision-making ubiquitous in society. Is it concerning, and is it all that different from algorithmic scoring systems elsewhere?</p> | <p><u>Black Mirror</u>: Nosedive (series 3 episode 1)</p> <p><u>Read</u></p> <p>Mara Hvistendahl, <a href="#">Inside China's Vast Experiment in Social Ranking</a>, Wired, Dec. 14, 2017</p> <ul style="list-style-type: none"> <li>• What are the major elements of China's system?</li> </ul> <p>Khadeeja Safdar, <a href="#">On Hold for 45 Minutes? It Might Be Your Secret Customer Score</a>, Wall St. Journal (2018)</p> <ul style="list-style-type: none"> <li>• How are these trust scores similar to, and different than, the social credit scores in China?</li> </ul>  |
| <p>October 21   December 9</p> <p><b>ALGORITHMIC ACCOUNTABILITY</b><br/>How can firms best respond to the challenges we've discussed in the course?</p>   | <p><u>Read:</u></p> <p><a href="#">Principled Artificial Intelligence: A Map of Ethical and Rights-Based Approaches</a>, July 4, 2019 draft</p> <ul style="list-style-type: none"> <li>• Does seeing all these different AI ethics frameworks compared make you feel better or works about the prospects for solutions?</li> </ul> <p>Bethan Cantrell, et al, <a href="#">Industry Needs to Embrace Data Ethics: How It Could Be Done</a> (2016)</p> <ul style="list-style-type: none"> <li>• What do you think of Microsoft's data ethics principles?</li> </ul> <p><a href="#">Algorithmic Accountability Act of 2019</a> (Senate bill)</p> <ul style="list-style-type: none"> <li>• What would this law require?</li> <li>• How effective would it be in addressing the concerns we've discussed in the class?</li> </ul> |

## 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 roughly 500,000 enrollments. His books include *After the Digital Tornado: Networks, Algorithms, Humanity, 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" for his research.