



# Asymmetric Information, Big Data, and Algorithmic Economic Decision-Making: A Simple Lesson on the Consequences for Statistical Discrimination and Civil Liberties

The information age of big data has increased the use of predictive data analytics for decision-making, leading to complex societal and legal debates with respect to statistical discrimination and the protection of civil liberties. This paper presents a simple pedagogic narrative in the context of asymmetric information. Students walk through a predictive-algorithmic decision-making process by a principal and agent. A simple predictive algorithm is constructed that leads to statistical discrimination and incentivizes principals to acquire more information confronting civil liberties of the right to privacy and protection from unnecessary surveillance. The lesson can be applied to labor, financial, and insurance markets.

**Mary Mervat**

<sup>†</sup>John Cabot University

## 1. Introduction

Integrating the theoretical material covered in an introductory economics course with the digital economy is becoming more imperative (Lazovic et Al., 2021). Developing teaching strategies that strengthen student engagement by connecting their learning with contemporary issues confronting society (Shanahan & Meyer, 2001; Taylor & Parsons, 2011) would contribute to fostering this integration. In this paper, we provide a simple lesson that links the increasing use of predictive algorithms for probabilistic decision-making by firms with the legal and moral issues that arise from their use (Kleinberg et al., 2018; Favaretto, De Clercq, & Elger, 2019; Tilcsik 2021). We focus on two important societal and legal outcomes. The first is statistical discrimination which occurs when decision-makers use the average characteristics of a group to inform their economic decisions about the expected behavior of individuals (Phelps, 1972; Aigner & Cain, 1977). This can lead to the perpetuation of unfair biases. The second is erosion of civil liberties with respect to the right to privacy and protection from unnecessary search and seizure that arises from conditions generated by an algorithmic decision-making process (Barocas et Al., 2017).<sup>1</sup> The use of predictive algorithms incentivizes firms to acquire more information on individuals to achieve economic goals. How that information is obtained and who has access to it, confronts issues of privacy and surveillance.

The format of the lesson is to use a simple pedagogic narrative to walk students through a probabilistic decision-making process under conditions of asymmetric information. The narrative follows the spirit of the great physicist, Richard Feynman: use simple examples to illustrate more complex ideas (Phillips, 2013) that can then form the basis for exploring their many implications. The narrative has two protagonists, a principal and an agent. The principal can be an employer making hiring decisions, a commercial bank deciding on credit allocation, or an insurance company determining rates. The agent is the individual who is the recipient of the decision. The lesson has the following specific student learning goals:

- (i) Understand why asymmetric information may require decision-makers, such as firms, to use predictive algorithms to place probabilistic bets to achieve an economic goal;
- (ii) Develop a simple predictive algorithm using a single variable population linear regression to illustrate how information contained in big data is used to make decisions based upon conditional means;
- (iii) Learn how predictive algorithms can result in statistical discrimination and understand the differences between discrimination generated by tastes or systemic structures;
- (iv) Realize and explore the incentives that arise in the decision-making process to acquire or provide more information by both the principal and agent and show how this is linked to the protection of civil liberties; and
- (v) Raise awareness on how the gathering and use of big data in predictive algorithms can create moral and legal issues in society.

The core lesson provides definitions and analyses using a neutral example as a foundation so that the instructor may easily adapt it to economic decisions related to employment,

---

<sup>1</sup> Civil rights are protections for equal treatment of groups (e.g., age, gender, religion, race) in the application of civil liberties.

insurance pricing, and credit allocation. The outline of the lesson follows.

## 2. Lesson Outline

The structure of the lesson follows a pedagogic narrative of the decision-making process to introduce key definitions and analytical concepts step by step.

### 2.1 Prerequisite Knowledge

The prerequisite knowledge for students is listed as follows:

- (i) know how to compute an average;
- (ii) know the basics of algebra to use two points to find the slope and intercept of a linear equation;
- (iii) understand the basics of probability: selecting a specific event out of a set of all possible events.

For more advanced economic classes, where students have had a basic statistics course, the lecture can elaborate on extending a simple regression to a multiple regression, incorporating a discussion of Bayesian learning processes of updating prior probabilities or beliefs when presented with new information, and demonstrating how probability distributions for two groups may or may not overlap.

The lecture begins with the instructor defining asymmetric information and providing examples of who is a principal and who is an agent. To maintain the focus on the topic of economic decision-making and its consequences for statistical discrimination and civil liberties, it is helpful for the instructor to provide a very brief overview of discrimination that can arise for other reasons. The first is the “taste for discrimination” as developed by Becker (1957). Taste for discrimination generates differential decisions based upon the tastes or dislikes of different groups by the principal and agent. The second is the complex economic-legal-societal factors in a country, region, or historic period that generate systemic differences in group economic outcomes. Acknowledging other factors that underlie economic discrimination clarifies for students that the lesson’s singular focus on statistical discrimination does not exclude these other reasons.<sup>2</sup>

### 2.2 The Objective of the Decision-maker and Initial Conditions of Asymmetric Information

The narrative first establishes the economic objective of the principal, in this case, a firm, under conditions of asymmetric information. The firm has determined they need a tall individual as the necessary hiring qualification to achieve their economic goal of stocking shelves when they do not have a ladder. Selecting the tallest individual will allow the firm to maximize its profits and subsequent survivability. Including survivability allows the instructor to reinforce that some employers may also feel a great deal of responsibility for the well-being of their existing workforce. The firm has a pool of applications to choose from. Given the context of asymmetric information, individuals know their height, but the firm only knows their age. An individual is either a 10-year-old or a 12-year-old.

---

<sup>2</sup> It is easy enough to show here how distributions across groups can vary and overlap to highlight that we are not addressing the reasons that generate differences in distributions.

### 2.3 Development of a Simple Predictive Algorithm

How does the employer make the hiring decision? This part of the lesson introduces students to what is meant by a predictive algorithm. The firm needs a predictive algorithm for height. A simple population regression model can be used as a rudimentary predictive algorithm. The use of a population regression is purposeful as it avoids confounding the lesson with the greater complexities of inferential statistics.<sup>3</sup>

Development of a predictive algorithm requires data that highlights to students the connection between the decision-making process and, usually historic, big data. It's best to use some simple data on age and height that facilitates computations. In this example, we use  $n = 8$  observations for two age groups: age = 10 and age = 12. The data on age and height are provided in Table 1.

**Table 1: Population Data on Height and Age**

Height in meters by age		
	Age 10	Age 12
	0.85	1.00
	0.90	1.05
	0.95	1.10
	1.00	1.15
	1.05	1.20
	1.10	1.25
	1.15	1.30
	1.20	1.35
<b>Average</b>	<b>1.025</b>	<b>1.175</b>

Before introducing a simple linear population regression model, ask the students to compute the average height for the two age groups. After that, introduce a simple linear population regression line as the predictive algorithm. The dependent variable  $Y_i$  is height in meters and the independent variable  $X_i$  is age in years. We leave off the error term as we want students to focus on learning that the predictive outcome of a regression is a conditional mean. The regression equation is given by Equation (1):

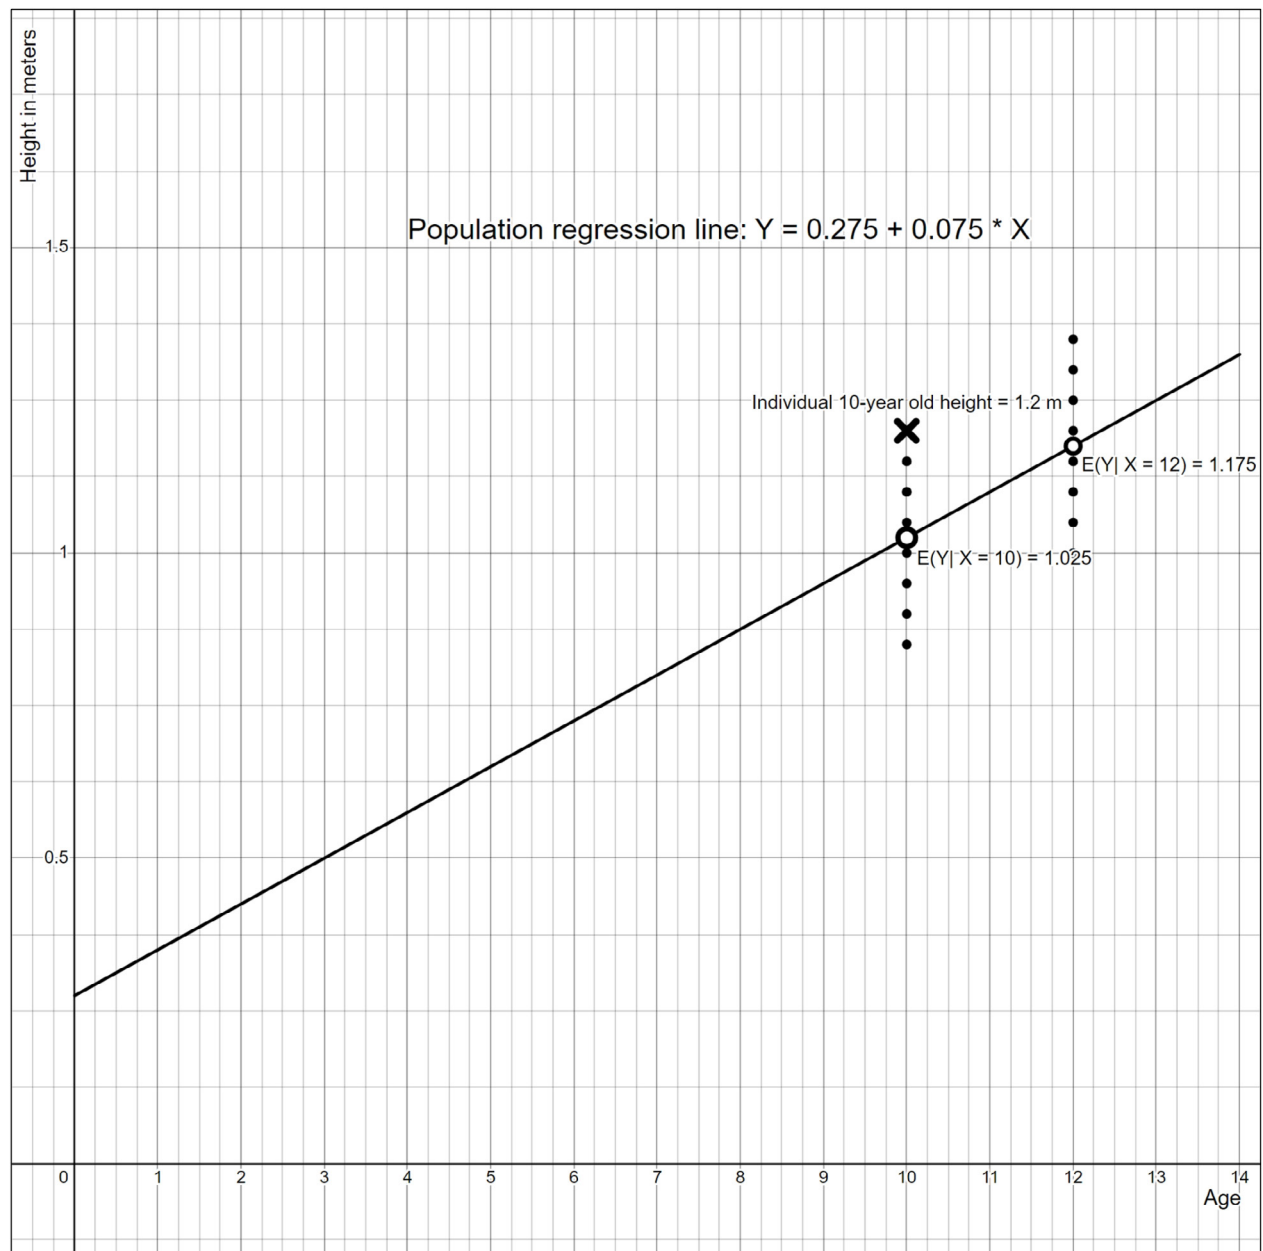
$$(1) Y_i = \beta_0 + \beta_1 X_i \text{ or } E(Y_i | X_i) = \beta_0 + \beta_1 X_i$$

Have students find the slope and intercept for the two points of age and average height. While simplistic, students have just built a predictive algorithm. The population regression line is given in Equation 2. The data and regression line are both shown in Graph 1.

<sup>3</sup> If students in an economics course have already had statistics, it still makes sense to keep this part of the lecture simple as later, when discussing the need for better predictions, the increase in data can be framed within inferential statistics.

$$(2) E(Y | X_i) = \beta_0 + \beta_1 X_i = 0.275 + 0.075 \bullet X_i$$

**Graph 1: Data and Population Regression Result for Height and Age**



Using the population regression line given by Equation (2), ask students to predict the heights for both 10-year-olds and 12-year-olds and have them verify that the predictions are the conditional means for each group. This reinforces for students that the outcome of a regression is a conditional mean which is a key technical point of the lesson. It is helpful to have students look at the data to be clear that each observation is a unique individual that may not conform to their group's average. This sets the next stage for the pedagogic narrative: using the predictive algorithm for decision-making.

## 2.4 *Using the Predictive Algorithm for Decision-Making: Introducing Statistical Discrimination and Civil Liberties*

The firm, lacking complete information and having the objective of hiring the tallest person, uses the predictive algorithm to make a hiring decision. The regression results show the conditional mean for the age 12 group is higher at 1.125 meters compared to 1.025 meters for the age 10 group. Based on this, students should come up with a decision rule for the firm: select an individual from the group that has the highest average height. Note to the students that the firm has no preference or taste for one group or another. They simply want to hire a tall person and using the predictive algorithm indicates the firm should select an individual from the group of 12-year-olds. This decision-making process using the predictive algorithm excludes a priori anyone who is 10 years old. This is what is meant by statistical discrimination against 10-year-olds. Some 10-year-olds are taller than 12-year-olds but, unfortunately, they will never be selected.

At this stage in the lesson, students can be prompted to ask if statistical discrimination could be avoided by eliminating the condition of asymmetric information where the firm only knows if an individual is 10 or 12 years old. The students can be guided to reason that if the firm knew the height of every individual, they would simply pick the tallest person. Getting more information appears to be a solution that reduces statistical discrimination and improves the prediction of the algorithm. It raises another problem, however, as it requires every individual participating in the selection pool to have their heights measured, verified, and recorded for the firm's use. Students can now begin to connect the firm's incentive to gathering more information that confronts the erosion of individual civil liberties of the right to privacy and protection from unreasonable surveillance.

## 2.5 *Exploring the Incentives by the Principal and Agent to Acquire or Provide More Information*

The incentives to acquire or provide more information by both the principal and some agents and its connection to statistical discrimination and civil liberties can now be further explored. The firm's use of a predictive algorithm in this decision-making process incentivizes them to acquire more information. Individuals also have incentives to reveal (or not) their private information. This is more subtle as the incentive to "show one's cards" is non-neutral with respect to their group identifier. If heights cannot be verified, short individuals who are age 12 have an incentive to hide their height. Tall individuals who are age 10, have an incentive to reveal their height. Additionally, there are international differences in the protection of civil liberties concerning data privacy and ownership. For example, in some countries, individuals may be coerced to provide their private information. We divide the next part of the lesson to analyze the information responses for the firm and then for the individual.

## 2.6 *Incentives and Implications for the Firm to Acquire More Information*

To understand the incentives for the firm to acquire more information, we introduce students to the concept of a "probabilistic bet" which also provides them with a better understanding of the risks faced by firms. The algorithm used by the firm led them to randomly select an individual from the group of 12-year-olds. While selecting from the group of 12-year-olds is the best rational approach to decision-making when information is incomplete, the predictive model does expose the firm to a 1/8 probability of ending up with the shortest 12-year-old who is 1.0 meters. If this happened, the firm could go out of business! To reduce this risk the firm could attempt to add other variables to predict height instead of using one variable, age. This turns the students' focus to the role of the independent variables and big data.

To ensure students understand the role of the independent variables in a regression model, it can be helpful to term the collection of independent variables as an “identity box.” This provides students with a clear understanding that the set of independent variables is defining a group identity to compute a conditional mean. It is the conditional group mean that is used for the subsequent prediction of economic decisions. Reminding students that classifying an individual by their “identity box”<sup>4</sup> may be helpful for prediction, but it ignores their unique characteristics. No one may be average.<sup>5</sup> Students’ curiosity can be provoked by exploring with them how much data is available on individuals and who has the right to collect and use it. This is a simple engaging question that motivates students to consider the subtleties of how information is collected. Is private data a commodity that can be exchanged for access to platforms or is private data protected as a general human right?

To make the lesson more international, students could be made aware that countries may differ in their protection of data privacy with respect to civil liberties. In the EU, personal data protection is considered a human right under the General Data Protection Regulation (GDPR).<sup>6</sup> In the US personal data is a commodity that can be traded for services such as platform access (Custers & Malgieri, 2022). Personal data may also be confiscated by a government as exemplified by China’s Social Credit System (Jakob, 2021) which could give them an AI advantage. Li, Tong, and Xiao (2021) write, “China’s lack of clear policies and regulations in areas such as privacy can explain how it caught up so rapidly in certain AI application fields. For example, the ubiquity of surveillance cameras in China creates a big market for AI firms specializing in visual and facial recognition.”

Once students understand the incentives by firms to acquire more information, the instructor can re-connect this to the regression model and loosely illustrate that, as more independent variables are added to the algorithm, the predictive conditional mean is for smaller group identities or profiles. An exceptionally large list of variables can be used to illustrate that in the limit, you can avoid both statistical discrimination and perhaps the errors of probabilistic bets if you can exactly measure, collect, and monitor the characteristics of individual observation. Using our original example on age and height, simply pick the tallest person.<sup>7</sup> In sum, while more information gives the principal greater predictive power for their algorithmic decision-making and reduces the extent of statistical discrimination, individuals may or may not have the agency to provide their personal information. The extent of the erosion of their civil liberties would depend upon whether the individual consents or not to release their private information.<sup>8</sup> This can generate a lively discussion as students have direct experience with this on social media platforms.

## 2.7 *Incentives and Implications for the Individual to Provide More Information*

Do individuals want to release their private information? Does it depend on the identity of the group the algorithm puts them in? Ask students to shift their perspective from the firm to an individual. Returning to the original data of age and height, consider an individual who is in

<sup>4</sup> Using “identity” as a term here stimulates interest as students may have heard of identity politics Heyes (2020).

<sup>5</sup> Identity and self are a big topic in psychology and interested students may be referred to (Stryker and Burke 2000).

<sup>6</sup> Human Rights Watch (2018) has a very nice summary of the GDPR that is accessible to students.

<sup>7</sup> It can be helpful to remind students of the classical statistic caveat of being unable to predict the future behavior of an individual.

<sup>8</sup> It is nice to remind students that protection of privacy varies by country and should not be taken for granted. They can be directed to explore the international differences in privacy protection that are summarized in international human rights laws by the UN Commission on Human Rights ([International Standards | OHCHR](#)).



the 10-year-old group and is 1.2 meters tall. Not only are they taller than the 1.125-meter average height of 12-year-olds, but they are taller than three of the 12-year-olds whose heights are 1.0, 1.05, and 1.10 meters, respectively. As we noted, using the algorithm, this tall 10-year-old will never be selected because they belong to the group with the lower conditional mean for height. This individual will always be judged by their group's average and never by their individual characteristics. This creates the persistence of discriminatory biases. Can this be addressed if individuals provide more private information?

It is interesting to see if the students begin to discover the non-neutrality of providing private information. Tall individuals in the group of 10-year-olds would have the most incentive to reveal personal information to get a job.<sup>9</sup> However, short individuals in the 12-year-old group would not only have the most incentive to hide their personal information, but they also have an incentive to prevent the tall 10-year-olds from revealing their private information as this reduces the "rent" received from being in the right group. This part of the lesson can be more challenging as it can easily be related to social privileges or class entitlements. This is what makes economic reasoning interesting and relevant for students. The instructor should be prepared to manage the discussion by reminding students the objective is to provide them with a greater analytical framework to address all the important observations and insights that this topic generates.

### 3. Economic Examples

Once students have walked through the simple decision-making process using the age and height example, the concepts can be applied analytically to specific economic examples. Below are a few ideas.

#### 3.1 Credit Allocation

The principal is a commercial bank facing the decision to approve a loan to limit defaults. Ask the students what factors would be used to predict whether an economic agent would default on a loan. A fun way to do this is to use one of the many applications for the determination of a FICO score such as the free FICO Score Estimator: [Free FICO Score Estimator | myFICO](#). Have students run through it answering the questions for someone who is highly likely to default to focus on individuals who may be excluded by the algorithm but are good credit risks.<sup>10</sup> The resulting statistical discrimination can be used to illustrate behavioral responses by individuals who are facing negative statistical discrimination and opportunities for financial innovation to identify these individuals and open new markets.

Excluded individuals have an incentive to find ways to manipulate the independent variables in the algorithm to improve their credit scores. Showing students with a simple Google search on "How to improve my credit score" raises an issue about the validity of a measure as a good predictor when the agent can manipulate it. This is an example of Goodhart's law that is often formulated as: "When a measure becomes a target, it ceases to be a good measure." This can motivate firms to find better measures. A Google search can illustrate to students how many start-ups are formed to use AI to find better credit scores and predictive algorithms to sell to commercial banks and other lenders.<sup>11</sup> The excluded individuals who may simply not have a

---

<sup>9</sup> Students may note that tall 10-year-olds would also have more incentives to acquire market signals (Spence 1973).

<sup>10</sup> As suggested by a reviewer, it could be interesting to have students produce credit scores for different group profiles and compare them.

<sup>11</sup> Interested students could be referred to a summary of AI and credit scoring. See Addy et al., (2024).



credit card or credit history are also a new market. Students may be interested in learning more about financial inclusion and how a digital economy has spurred financial innovations when credit information is lacking. Students may enjoy reading about Kenya's telecoms operator Safaricom and the digital bank M-Pesa (Murray, 2023).

### *3.2 Interview and Hiring Decisions*

Employment decisions that use predictive algorithms can result in statistical discrimination. For example, the classic study by Bertrand and Mullainathan (2004) shows differential outcomes in interview callbacks based on whether a name on a resume is perceived as a White or African American name. Students may find discussions on the biases and discrimination introduced by AI algorithms used in job search platforms. The article from MIT Technology Review could be used here: "LinkedIn's job-matching AI was biased. The company's solution? More AI" (Wall & Schellmann, 2021). The instructor can also demonstrate the behavioral reaction of agents to statistical discrimination by doing a Google search on "How to improve LinkedIn profiles to match AI algorithms."

If the instructor wants to tilt the lesson towards statistics, then introducing sequential decision making and Bayesian processes can be included. Do firms update their views of the distribution of characteristics across groups as they hire more individuals from both groups? The article on Bayesian Bigots (Pager & Karafin, 2009) can be a useful reference and is accessible for students to read. For example, returning to the age-height employment decision of the firm, if a firm did hire a 10-year-old who was tall would they change their views about the heights of 10-year-olds or consider the individual a subtype? How are prior beliefs updated if a firm is small compared to a large firm? Large firms could potentially use sequential hires to learn following a Bayesian process of updating prior beliefs as new information becomes available. Small firms, however, may not make enough sequential hires to learn and update prior beliefs. This may result in smaller firms relying on their biased perceptions during an interview process regarding who would be the best fit for the needs of the firm. Students may be asked if this may lead to the persistence of discrimination. Regardless of size, all firms are attempting to hire people they believe will get the job done but none of them know a priori whether an individual candidate is that person.

### *3.3 Insurance Pricing*

Students can find meaningful connections to the lesson if you ask them about car insurance. Who are the riskier drivers and, therefore, should pay more for insurance? How does the firm identify these drivers? Connecting this back to the age variable in the simple model, if younger drivers are riskier compared to older drivers, then we would conclude that younger people should pay higher insurance rates. Connecting this to predictive algorithms and statistical discrimination, if someone is young but is a safe driver, they may feel it is unfair to pay the higher rates just because of their group's characteristics. At the same time, older people who are not safe drivers would not want that information revealed; they are happy being in the "safe" group. How does this relate to the incentives concerning car tracking devices? Who would select to have one installed? Students may respond that drivers who know they are less risky would have an incentive to trade their privacy for lower rates. Students may realize that if this is the case, then someone who is low risk but refuses to be tracked because they value their privacy, would effectively pay an implicit price for this choice. This is often a lively topic as it is accessible to students and allows them to make wider connections on the trade-offs between statistical discrimination and the right to privacy.

#### **4. Concluding Remarks**

This basic introductory lesson was designed to link predictive data analytics to decision-making leading to the complex moral and legal debates surrounding statistical discrimination and the protection of civil liberties. This was achieved with a simple lesson using a pedagogic narrative that walked students through a firm's decision-making process in the presence of asymmetric information that requires the use of predictive algorithms. Connecting the use of predictive algorithms in the decision-making process creates behavioral incentives for both the principle and agent with respect to information gathering or giving. The lesson was designed to focus on fostering student engagement by providing relatable connections between economic behavior in a digitalized big data world and its consequences for statistical discrimination and civil liberties.

## References

- Addy, W. A., Ajayi-Nifise, A. O., Bello, B. G., Tula, S. T., Odeyemi, O., & Falaiye, T. 2024. AI in credit scoring: A comprehensive review of models and predictive analytics. *Global Journal of Engineering and Technology Advances*, 18(2), 118–129. DOI: [10.30574/gjeta.2024.18.2.0029](https://doi.org/10.30574/gjeta.2024.18.2.0029)
- Aigner, D. J., & Cain, G. G. 1977. Statistical theories of discrimination in labor markets. *Industrial and Labor Relations Review*, 30(2), 175–187. DOI: [10.2307/2522871](https://doi.org/10.2307/2522871)
- Barocas, S., Bradley, E., Honvar, V., & Fostor, P. 2017. Big data, data science, and civil rights. *arXiv:1706.03102*, 1-2. DOI: [10.48550/arXiv.1706.03102](https://doi.org/10.48550/arXiv.1706.03102)
- Becker, G. 1957. *The economics of discrimination*, University of Chicago Press.
- Bertrand, M., & Mullainathan, S. 2004. Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. *The American Economic Review*, 94(4), 991-1013. DOI: [10.1257/0002828042002561](https://doi.org/10.1257/0002828042002561)
- Custers, B., & Malgieri, G. 2022. Priceless data: why the EU fundamental right to data protection is at odds with trade in personal data. *Computer Law & Security Review*, 45, 105683, ISSN 0267-3649, DOI: [10.1016/j.clsr.2022.105683](https://doi.org/10.1016/j.clsr.2022.105683)
- Favaretto, M., De Clercq, E. & Elger, B. S. 2019. Big data and discrimination: perils, promises and solutions. A systematic review. *Journal of Big Data*, 6, 12. DOI: [10.1186/s40537-019-0177-4](https://doi.org/10.1186/s40537-019-0177-4)
- Heyes, C. 2020. Identity politics, *The Stanford Encyclopedia of Philosophy*. Fall Edition, Edward N. Zalta (ed.), <https://plato.stanford.edu/archives/fall2020/entries/identity-politics/>
- Human Rights Watch. 2018. The EU general data protection regulation questions and answers. June 6. <https://www.hrw.org/news/2018/06/06/eu-general-data-protection-regulation>
- Jakob, S. 2021. The corporate social credit system in China and its transnational impact. *Transnational Legal Theory*, 12(2), 294-314. DOI: [10.1080/20414005.2021.1977019](https://doi.org/10.1080/20414005.2021.1977019)
- Kleinberg, J., Ludwig, J., Mullainathan, S., & Sunstein, C. R. 2018. Discrimination in the age of algorithms. *Journal of Legal Analysis*, 10, 113–174. DOI: [10.1093/jla/laz001](https://doi.org/10.1093/jla/laz001)
- Lazovic, V., Rondovic, B., Lazovic, D., & Djurickovic, T. 2021. Is economic theory, presented in basic academic textbooks, applicable to the digital economy? *Sustainability* 13(22), 1-23, November. DOI: [10.3390/su132212705](https://doi.org/10.3390/su132212705)
- Li, D., Tong, T. W., & Xiao, Y. 2021. Is China emerging as the global leader in AI? *Harvard Business Review*, February 18. <https://hbr.org/2021/02/is-china-emerging-as-the-global-leader-in-ai>
- Murray, S. 2023. How mobile money fosters financial inclusion. *Knowledge at Wharton*, December 15. <https://knowledge.wharton.upenn.edu/article/how-mobile-money-fosters-financial-inclusion/>
- Pager, D., & Karafin, D. 2009. Bayesian bigot? statistical discrimination, stereotypes, and employer decision making. *The Annals of the American Academy of Political and Social Science*, 621(1), 70–93. DOI: [10.1177/0002716208324628](https://doi.org/10.1177/0002716208324628)

Phelps, E. S. 1972. The statistical theory of racism and sexism. *The American Economic Review*, 62(4), 659–661. <http://www.jstor.org/stable/1806107>

Phillips, R. 2013. In retrospect: The Feynman lectures on physics. *Nature*, 504, 30–31. DOI: [10.1038/504030a](https://doi.org/10.1038/504030a)

Shanahan, M. P., & Meyer, J. H. F.. 2001. A student learning inventory for economics based on the students' experience of learning: A preliminary study. *The Journal of Economic Education*, 32(3), 259–267. DOI: [10.2307/1183383](https://doi.org/10.2307/1183383)

Spence, M. 1973. Job market signaling. *The Quarterly Journal of Economics*, 87(3), 355–374. DOI: [10.2307/1882010](https://doi.org/10.2307/1882010)

Taylor, L., & Parsons, J. 2011. Improving student engagement. *Current Issues in Education*, 14(1). Retrieved from <http://cie.asu.edu/>

Stryker, S., and Burke, P. J. 2000. The past, present, and future of an identity theory. *Social Psychology Quarterly*, 63(4), 284–297. DOI: [10.2307/2695840](https://doi.org/10.2307/2695840)

Tilcsik, A. 2021. Statistical discrimination and the rationalization of stereotypes. *American Sociological Review*, 86(1), 93–122. DOI: [10.1177/0003122420969399](https://doi.org/10.1177/0003122420969399)

Wall, S., & Schellmann, H. 2021. LinkedIn's job-matching AI was biased. The company's solution? More AI." *MIT Technology Review*, June 23. <https://www.technologyreview.com/2021/06/23/1026825/linkedin-ai-bias-ziprecruiter-monster-artificial-intelligence/>