Data & Civil Rights: Employment Primer

by Alex Rosenblat, Kate Wikelius, danah boyd, Seeta Peña Gangadharan, and Corrine Yu
Produced for Data & Civil Rights Conference / October 30, 2014

Our nation's employment laws are founded on the principle that every worker should be judged solely on her or his merits. Hardworking Americans should not be kept from supporting their families and making a positive contribution to the economic life of our nation because of characteristics that have no bearing whatsoever on their ability to do a job. Though U.S. workplaces today are more inclusive than ever before, ensuring equal employment opportunity continues to pose significant challenges. These challenges manifest themselves in a number of ways: through barriers to career access and advancement; the persistence of inequitable pay structures; and the persistence of stereotypes and prejudice.

Employees and prospective employees produce more data than ever - in the workplace, on social media, and beyond. Employers and the third party companies that assist them increasingly apply analytical tools to these various data streams to measure factors that influence employee performance, attrition rates, and workplace profitability.[1] While some of the data – such as past performance – are unquestionably relevant to such analysis, other data that produces strong correlations to performance are more surprising. For instance, Evolv, a recruiting software company, analyzed 3 million data points about 30,000 hourly employees and identified that those who installed newer browsers, like Chrome or Firefox, onto their computers stay at their jobs 15% longer than those who use default browsers that come pre-installed on their computers, like Safari for Macs.[2] Job candidates may rightly worry that they will be excluded from or included in job opportunities based on data that seem arbitrary and are outside their field of vision. For example, a job candidate's resume could be excluded from a talent pool because of her online browsing habits,[3] but she is unlikely to find that out directly. The complexity of hiring algorithms which fold all kinds of data into scoring systems make it difficult to detect and therefore challenge hiring decisions, even when outputs appear to disadvantage particular groups within a protected class. When hiring algorithms weigh many factors to reach an unexplained decision, job applicants and outside observers are unable to detect and challenge factors that may have a disparate impact on protected groups.

Background: Discrimination in Employment

Major civil rights concerns and themes in U.S. employment:

- *Underrepresentation*: Women and minorities constitute a significant percentage of the workforce overall, yet they are severely underrepresented in many high-wage occupations and sectors, such as the tech industry.
- *Pay equity*: According to the U.S. Census Bureau, in 2012, women, who make up nearly half of the workforce, earned, on average, only 77 cents for every dollar men earned for comparable jobs. African-American women working full time, year-round typically make

64 cents for every dollar paid to their White male counterparts, and Latina women make only 54 cents.

- *Criminal background checks and credit history checks:* Because people of color are arrested and convicted at rates that far exceed their numbers in the population at large, criminal records-based discrimination has a disproportionate adverse impact on people of color. There is little research that shows any correlation between the existence of a criminal record and the propensity to commit crimes at the workplace.[4] Similarly, the use of poor credit to cut off employment opportunities has had a disparate impact on minorities, who tend to have worse credit, on average, than Whites.[5] Currently, 47% of major employers use credit background checks during the hiring process to screen out employment applicants with poor credit,[6] despite the fact that there is no proven link between personal credit reports or criminal behavior and performance of a specific job.[7]
- *Workplace harassment*: Thousands of harassment charges are filed with the Equal Employment Opportunity Commission ("EEOC") and state and local agencies each year, and the numbers are on the rise.[8]
- *Discrimination against LGBT Communities*: There is no federal law that explicitly protects LGBT people from employment discrimination. Thirty-two states currently lack such explicit protections for gender identity, while 28 states lack explicit protections for sexual orientation.

Employment Algorithms

Job candidates can be assessed by scoring algorithms that use all sorts of information to edge their score higher or lower, which makes it easier for employers or human resource managers to sift through a myriad of applications or a talent pool of candidates. Hiring algorithms produce outputs that are quantitative, not necessarily fair. For instance, Evolv identified that employees who live 0-5 miles from their workplace stay at their jobs 20% longer than those who live further away, although this quality, too, is ripe for cumulative disadvantage since the myriad of reasons why employees leave their jobs can reflect the influence of bias. However, the company decided to exclude the distance a worker lives from their workplace from its hiring algorithms because of concerns over using a factor that could be discriminatory towards members of lower-income communities.[9] As legal scholar Julie Cohen observes, "…opposition to entrenched societal discrimination is hard to reconcile with commitment to the truth-value of information… Effective antidiscrimination policy therefore requires the exercise of moral judgment about the value of information."[10] In other words, a policy of removing relevant criteria from algorithmic calculations may reduce discrimination, but may also confound the accuracy of scoring systems.[11] Moreover, other criteria can act as proxies for the specific, objectionable variables that are removed (i.e. a person who lives further away may have other qualities that define them as part of a lower-income community).

A policy that exercises judgment about the value of information has to take into account the deficiencies of algorithms, such as encoded bias, as well as the advantages they offer, like efficiencies. For example, Entelo, another recruiting software company, developed an anti-discrimination, diversity-hiring algorithm to sort candidates specifically by protected class characteristics in response to employers who expressed a desire to improve the gender balance on their engineering teams.[12] An efficient, data-driven solution to diversity is one method for circumventing bias in the more ingrained hiring and promotion processes of a given employer, but it would also be interesting to see how companies use data to examine their own workplace policies. For example, after recognizing that their workforce is insufficiently diverse, Google examined its own policy for promotion, which is through self-nomination by employees. They found that it favored men over women, who were less likely to self-nominate; subsequently, they used more data to discover that women can be successfully nudged to self-nominate if they are sent emails inviting them to do so.[13]

Some companies apply big data tools to measure workplace performance. Sociometric Solutions, which has clients that include the Bank of America and Cubist Pharmaceuticals, uses sensors placed on employee identification badges to track who is talking to whom, for how long, with what tone of voice, how quickly they speak or when they interrupt, etc. to try to identify what makes for a good team and what does not. In one instance, they discovered that the most important interactions between salespeople outside their teams happened at the coffee machine.[14] Although Sociometric Solutions withholds individualized data from the companies that use its services, workers undergoing increased surveillance have concerns that their privacy and autonomy will be negatively impacted by the data sought about them.[15] How do we clarify the distinctions between client (like Sociometric Solutions) versus supervisor use of the data collected on employees?

Workplace Analytics

Workplace analytics are used to identify what matters to success in the workplace, and the data demonstrates that some of the factors we think of as important are less relevant than we assume, while other factors are more relevant. For example, data analysis is leading some employers to conclude that college degrees are not necessarily strong predictors of employee success; instead, they are looking toward other signals in the data that affect their hiring outlook.[16] Workplace analytics are also creating new insights into the factors that are predictive of employee retention, or the success of a workplace overall. For example, Evolv identified that the walking score of a workplace location is far more predictive of employee tenure than the length of employee commutes.[17] Employees who work within walking distance of parks, restaurants, shops, and other amenities are more likely to stay in their jobs voluntarily and are less likely to get fired; employers with a middling walking score retain employees 58% longer than locations with the lowest walking scores.[18]

In a hypothetical scenario, an employment algorithm that takes into account knowledge of the geographical features of a job candidate, like their commute and childcare options, could help potential employees optimize for workplace satisfaction and retention. For large employers with multiple franchise locations and many lower-level, hourly workers, where employees have high turnover and similar skills, an employer could use an algorithm to personalize schedules or worksite locations according to the geographical points of interest to a job candidate, such as their children's school. To avoid discrimination against job candidates with long or varied commutes and family responsibilities, this optimization (i.e. personalization) could take place after the initial job offer is made, and could be voluntary, with precautions taken to avoid a coercive volunteerism. The tension with this kind of data analysis centers on surveillance; an employee may not want to participate in employment practices that gather personal information unrelated directly to their job performance, like their personal or familial geography.

Algorithmic Discrimination

Hiring algorithms are designed using historical data to create predictions about which qualities correlate to a strong job performance. If workplace discrimination has historically elevated the performance of one group over others, then algorithms derived from such historical data will tend to reinforce that historical bias. This is likely even if the algorithm's designer does not intend to discriminate based on these categories.[19] Even without such historical bias, the accuracy of an algorithm's outputs will still be higher for the dominant statistical group, and lower for the statistical minority, because more information is available about job performance for the larger group.[20] Algorithms that are trained on data sets that include both inaccurate and accurate information output biased or meaningless information; worse, such dirty data can dilute other valuable signals and send inaccuracies rippling through entire systems. Because of limited transparency, unclean data, and the complexity of most algorithms used to do this kind of analysis, it is often difficult to discern the specific reasons for which a job candidate receives a negative score.

The promise of data-driven solutions in the employment sector often downplays the bias and errors in hiring models. Errors in information-processing systems can surface false positives, which then lead to bad hiring decisions and cause employers to reject high-quality candidates. Error may also disproportionately affect members of marginalized groups. For example, the Department of Homeland Security (DHS) offers employers an E-Verify program to confirm that job candidates or employees are authorized to work in the U.S. The information that employers input on employees is cross-referenced with information that DHS and the Social Security Administration have on them. If there is an inconsistency, E-verify will issue a "tentative nonconfirmation" (TNC) notice denying work authorization until the individual worker resolves the inconsistency.[21] Foreign-born employees, and employees with foreign-spelling names, are 20 times more likely to be wrongly flagged with a TNC.[22] If the dominant group in a dataset has Americanized spellings, then the algorithm is trained to accept names with variations that conform in their spelling to those of the dominant group, and to reject or flag spellings that do not conform to the features it is trained to distinguish between as correct or incorrect. Foreign-spelling names can experience a significantly higher error rate in a statistical model than the overall error rate for the group as a whole, which could be much lower. Compounding this, low-wage, hourly workers, whether they are flagged for a spelling error or for other reasons, often lack the time, resources, or legal literacy required to navigate complex bureaucracies to correct misinformation about them in a national database. As a result, they are more likely to have unjustified negative flags, which can cost them a job they are legitimately eligible for.[23]

As Barocas and Selbst observe, informational and categorical accuracy is expensive and difficult to achieve with the level of granularity that can guarantee equal and accurate outcomes for all people. Determining what information is inputted into a sorting mechanism, and how those features are analyzed to produce a particular outcome (like verified or unverified) "…can have serious implications for the treatment of protected classes, if those factors that better account for pertinent statistical variation [like foreign-spelling names] among members of a protected class are not well represented in the set of selected features."[24]

Employment and Health Scoring

Employers increasingly adopt screening tools to improve their workforces and reduce costs, but the opacity of scoring techniques can remove any chance an employee or job applicant has to appeal or negotiate with their assessment. For example, Richfield Management LLC, a waste-disposal company, reduced workers' compensation claims by 68% when they screened out applicants who represent a high disability risk with an initial recruitment test.[25] If a candidate scores poorly, the company will not hire her, without exception. While the law prohibits employers from discriminating against job candidates with disabilities, the screening process creates a regulatory grey area. For instance, personality tests used to screen job candidates may discriminate against those with mental health issues, by requiring them to agree or disagree with statements like, "Over the course of the day, I can experience many mood changes."[26] In the interest of self-protection and efficiency, many employers prefer to simply review outputs of screening tests, which categorize the costliness of prospective employees without knowing the reasons *why*.

In an ecosystem where data mining services can target medical interventions[27] based on credit card data, consumer purchase histories (like Target identifying which of its shoppers are pregnant[28]), or search engine queries,[29] it's not clear that the protections that currently exist to protect patient privacy, like the Health Information Privacy Protection Act (HIPPA), are sufficient to protect medical information from being used in the employment sector. For instance, FICO, the company which assesses consumer credit scores, developed a Medication Adherence Score to predict which patients are likely or unlikely to take prescribed medication based on publicly-available data about them, like home ownership and job status, and without reference to their specific, HIPPA-protected medical records. They came up with that scoring mechanism by analyzing the prescription records of anonymized patient data, and identifying which variables were predictive of high or low adherence.[30] Employers can similarly use a data broker's profile of a job candidate to infer their medical information, conditions or disabilities; subsequently, a complex hiring algorithm can factor in the cost of an inferred medical condition (even if the inference is inaccurate), and label a job candidate with a low overall score if they are identified as likely to be high-cost, without revealing that the specific health factor that resulted in an adverse impact for a job candidate.[31] In essence, employers want a score that predicts costliness; the reasons are irrelevant. As such, if they use only the score, they are necessarily guilty of illegal discrimination.

The "big data" phenomenon ushers in legal challenges that current regulatory regimes are ill-equipped to manage, like algorithmic discrimination against job candidates based on their medical conditions. Harms from networked information stem from the sudden availability of large amounts of data on individuals that is gathered and shared beyond their control. Legal remedies for *individual* harm are not structured in a way that accounts for *networked* harms. An injury must have a close relation to harmful conduct, such that a specific plaintiff can be held accountable by a defendant for a wrongdoing.[32]

Digital Reputations and Algorithms

Employers or the third parties used to assist employers in their staffing needs use algorithms to sift through job candidates' online reputations to identify whether or not they would be a good fit. According to a 2009 survey by Microsoft, 75% of HR professionals and recruiters in the U.S. reported that their companies have formal policies in place to research job candidates' online reputations; 89% of U.S. recruiters and HR professionals seek out professional online data, like

LinkedIn; and 84% think it is appropriate to check online personal data as well. While only 7% of consumers thought their online reputations affected their job prospects, 70% of HR professionals reported that they rejected candidates after mining their data.[33] The ways in which algorithms sort information online, including what might be characterized as negative information about prospective employees, differs significantly from traditional vetting processes, such as word-of-mouth referrals and other reputational tools.

When hiring managers Google job candidates, or when an algorithm does it for them, they scour the web for red flags; for instance, job candidates could be flagged for posting or featuring in sexually explicit material online. However, as Crawford and Gillespie observe, "…a flag is not merely a technical feature: it is a complex interplay between users and platforms, humans and algorithms, and the social norms and regulatory structures of social media."[34] They recall one example where a photo on Facebook of two men kissing was flagged by users as sexual, graphic content and removed by Facebook's community moderators. It was reinstated when other users protested that photos of heterosexual kissing couples were not subject to the same treatment. What an employer might see in that case, prior to the photo's reinstatement, is that a job candidate has posted sexually explicit material on the Internet, rather than the values that went into the flag. Flagging mechanisms essentialize people based on their data by sorting them into binary negative or not-negative categories. When a discriminatory judgment is wrapped up in a flagging mechanism, it can be challenging for job candidates to contest the negative inference attributed to them, even if the flag is based on inaccurate or implicitly biased data.

Online reputations can suffer from inferred connections made by advertisement, even in the absence of a direct red flag, because of how search algorithms sort and rank information in online databases. For instance, the ads that appear next to the names of job candidates reflect implicit racial biases of users who've searched for Black-sounding names before. Former FTC Chief Technologist Latanya Sweeney found that Google's AdSense algorithm turns up ads that suggest possible arrest records when racially-associated names like DeShawn are queried, far more often than for Caucasian-associated names like Jill.[35] In an employment context, these ads might have a disparate impact on job candidates with names associated with certain races. Barocas and Selbst have observed a parallel phenomenon in LinkedIn's featured "Talent Match" algorithm, which scours through 60 million LinkedIn profiles and selects 24 of them as suggestions for each new job posting.[36] If minorities are currently underrepresented in the type of job being advertised, then LinkedIn may reproduce such bias in the candidates it suggests as matches.[37]

Data analytics may also disempower workers in ways that violate our sense of due process. Companies that furnish reports on consumers (or employees) are obliged to comply with the tenets of the Fair Credit Reporting Act (FCRA), which requires consumer reporting agencies to get consumer consent for credit reports that affect employment. First Advantage Corporation, a company that was acquired by LexisNexis, maintained the Esteem database of employees accused of theft, usually in retail settings.[38] It had tens of thousands of subscribers, including major retail companies like CVS and Target.[39] Employers who used the service could populate the database with allegations, not just verified claims. This means that alleged thefts served as sufficient criteria to blacklist employees from future work opportunities.[40] LexisNexis suspended the database after coming to tentative settlement terms with workers who accused them of illegally distributing damaging information about employees.[41] LinkedIn currently faces a class-action suit over its jobs reference tool, which furnishes reference reports on job candidates to employers without the knowledge of job candidates, including the names, locations, and employment history of people who have worked with the job candidate. If FCRA applies, consumers might have the right to know about reports generated about them if they affect their employment prospects.[42]

Questions for Data, Civil Rights, and Employment

1. How do we define fairness?
2. What new kinds of data are employers using in hiring decisions? What newly available data are they choosing *not* to use?
3. Who is liable for discriminatory hiring decisions generated by algorithms?
4. If an algorithm masks a discriminatory bias in hiring, how can an individual challenge it?
5. How do data analytics and informational harm intersect with notions of due process, like informed consent and the right to challenge?
6. What policies or standards should govern the ethical use of data with significant social and legal implications on employment, both in the hiring context and in management situations?
7. Which groups of people are more vulnerable to hiring discrimination or discrimination in the employment context? How does this differ when individual decision makers make the determinations as opposed to determinations that are shaped by algorithm or by publicly available data?
8. Can outcome-oriented scoring systems, like diversity hiring and retentional algorithms or automated performance tools, reduce discrimination or inequity in the workforce?
9. How do the risks and benefits of new, algorithmic methods of assessing employee performance compare to the risks and benefits of more traditional methods?
10. What employment practices can be made fairer or more equitable through the use of data mining techniques? For example, who would benefit from algorithmically generated scheduling?
11. How can data be used to better enable equitable working environments?

[1] Some of this work contains excerpts from "Networked Employment Discrimination" and "Workplace Surveillance." Both articles are by Alex Rosenblat, Tamara Kneese, and danah boyd, from Oct. 2014. Available at http://www.datasociety.net/updates/featured/2014/10/working-papers-from-the-future-of-labor-project/.
[2] "How Might Your Choice of Browser Affect Your Job Prospects?", *The Economist*, April 10, 2013. http://www.economist.com/blogs/economist-explains/2013/04/economist-explains-how-browser-affects-job-prospects.

---

[3] Danielle Keats Citron, "Big Data Should Be Regulated by 'Technical Due Process'," *The New York Times,* August 6, 2014. http://www.nytimes.com/roomfordebate/2014/08/06/is-big-data-spreading-inequality/big-data-should-be-regulated-by-technological-due-process.
[4] National Employment Law Project, "65 Million 'Need Not Apply': The Case for Reforming Criminal Background Checks for Employment," by Michelle N. Rodriguez and Maurice Emsellem, (New York, NY: 2011). http://www.nelp.org/page/-/65_million_need_not_apply.pdf?nocdn=1
[5] *See, e.g.:* Board of Governors of the Federal Reserve System, *Report to the Congress on Credit Scoring and Its Effects on the Availability and Affordability of Credit* (2007); Federal Trade Commission, *Credit-Based Insurance Scores: Impacts on Consumers of Automobile Insurance* (2007); Robert B. Avery, Paul S. Calem, and Glenn B. Canner, "Credit Report Accuracy and Access to Credit," *Federal Reserve Bulletin* (2004); *The Brookings Institution,* "Credit Scores, Reports, and Getting Ahead in America," by Matt Fellowes (Washington, D.C., May 2006).
[6] Society for Human Resources Management, *SHRM Survey Findings: Background Checking The Use of Credit Background Checks in Hiring Decisions* (Alexandria, VA: July 19, 2012).
[7] Demos, "Discredited: How Employment Credit Checks Keep Qualified Workers Out of a Job," by Amy Traub (New York, NY: February, 2013), 8-9. Credit reports were developed to help lenders assess the risks associated with making a loan, not the worthiness of holding a job. A spokesperson for TransUnion, one of the major credit reporting companies, even admitted in 2010: "We don't have any research to show any statistical correlation between what's in somebody's credit report and their job performance or their likelihood to commit fraud."
[8] "Harassment Charges EEOC & FEPAs Combined: FY 1997 - FY 2011," Equal Employment Opportunity Commission, accessed November 11, 2014, http://www.eeoc.gov/eeoc/statistics/enforcement/harassment.cfm.
[9] Dustin Volz, "Silicon Valley Thinks It Has the Answer to Its Diversity Problem," *National Journal*, September 26, 2014. http://www.nationaljournal.com/next-america/economic-empowerment/can-companies-use-big-data-to-fight-racist-and-sexist-hiring-practices-20140926.
[10] Julie E. Cohen, "Privacy, Autonomy, and Information," in *Configuring the Networked Self: Law, Code, and the Play of Everyday Practice* (New Haven, CT: Yale University Press, 2012), 117-118. Cohen's work draws on Frederick Schauer, *Profiles, Probabilities, and Stereotypes* (Cambridge, MA: Harvard University Press, 2003).
[11] Martha Poon, "Discriminating Without Discrimination," in "What Lenders See — a History of the Fair Isaac Scorecard," (PhD dissertation, Department of Sociology, Science Studies Program, UCSD, 2012), 207, cited in Solon Barocas, "Panic Inducing: Data Mining, Fairness, and Privacy" (PhD dissertation, New York University, 2014), 244.
[12] "Announcing Entelo Diversity," *entelo blog*, by Jon Bischke, accessed September 23, 2014, http://blog.entelo.com/announcing-entelo-diversity.
[13] Cecilia Kang, "Google Data-Mines Its Approach to Promoting Women," *The Washington Post*, April 2, 2014. http://www.washingtonpost.com/blogs/the-switch/wp/2014/04/02/google-data-mines-its-women-problem/.
[14] Joshua Brustein, "The Case for Wearing Productivity Sensors on the Job," *Bloomberg Businessweek*, December 19, 2013. http://www.businessweek.com/articles/2013-12-19/sociometric-solutions-ben-waber-on-workers-wearing-sensors.
[15] Brustein, "Wearing Productivity Sensors."
[16] Don Peck, "They're Watching You at Work," *The Atlantic*, November 20, 2013. http://www.theatlantic.com/magazine/archive/2013/12/theyre-watching-you-at-work/354681/.
[17] Evolv, "Q3 2013 Workforce Performance Report," 2013, 13. http://go.evolvondemand.com/rs/evolvondemand/images/Evolv_2013_Q3_WorkforceReport.pdf
[18] Evolv, "Q3 2013 Workforce."
[19] Solon Barocas and Andrew D. Selbst, "Big Data's Disparate Impact," *SSRN*, October 19, 2014, 45-46. http://ssrn.com/abstract=2477899.
[20] Mortiz Hardt, "How Big Data Is Unfair," *Medium*, September 26, 2014. https://medium.com/@mrtz/how-big-data-is-unfair-9aa544d739de.
[21] American Civil Liberties Union. "Prove Yourself to Work: The 10 Big Problems with E-Verify," May, 2013, 2. https://www.aclu.org/files/assets/everify_white_paper.pdf.
[22] American Civil Libterties Union, "Prove Yourself," 7.
[23] Timothy B. Lee, "E-Verify Is Supposed to Stop Undocumented Employment. It Could Also Harm Legal Workers," *The Washington Post*, June 3, 2013. http://www.washingtonpost.com/blogs/wonkblog/wp/2013/06/03/e-verify-is-supposed-to-stop-undocumented-employment-it-could-also-harm-legal-workers/.
[24] Barocas and Selbst, "Big Data," 18.
[25] Barocas and Selbst, "Big Data."

---

[26] Lauren Weber and Elizabeth Dwoskin, “Are Workplace Personality Tests Fair?”, *Wall Street Journal*, September 30, 2014. http://online.wsj.com/articles/are-workplace-personality-tests-fair-1412044257.
[27] Colleen Taylor, “Omada Health Raises $23M Series B Led By Andreessen Horowitz For Digital Health That Actually Works,” *TechCrunch*, April 9, 2014. http://techcrunch.com/2014/04/09/omada-health-raises-23m-series-b-led-by-andreessen-horowitz-for-digital-health-that-actually-works/.
[28] Charles Duhigg, “How Companies Learn Your Secrets,” *New York Times*, February 16, 2012. http://www.nytimes.com/2012/02/19/magazine/shopping-habits.html.
[29] Shannon Pettypiece and Jordan Robertson, “Hospitals Are Mining Patients' Credit Card Data to Predict Who Will Get Sick*,” Businessweek,* July 3, 2014. http://mobile.businessweek.com/articles/2014-07-03/hospitals-are-mining-patients-credit-card-data-to-predict-who-will-get-sick.
[30] Tara Parker-Pope, “Keeping Score on How You Take Your Medicine,” *The New York Times*, June 20, 2011. http://well.blogs.nytimes.com/2011/06/20/keeping-score-on-how-you-take-your-medicine/.
[31] Frank Pasquale, *The Black Box Society* (Cambridge, MA: Harvard University Press, 2015), 36.
[32] danah boyd and Karen Levy, "Networked Harm – PLSC," (unpublished manuscript), 3.
[33] Cross-Tab Marketing Services, “Online Reputation in a Connected World,” January, 2010, 5-7. http://www.microsoft.com/security/resources/research.aspx#reputation.
[34] Kate Crawford and Tarleton L. Gillespie, “What is a Flag for? Social Media Reporting Tools and the Vocabulary of Complaint,” *New Media & Society* (August 5, 2014), 1.
[35] Latanya Sweeney, “Discrimination in Online Ad Delivery,” *Data Privacy Lab*, 2013. http://dataprivacylab.org/projects/onlineads/.
[36] Parker Barille, “Helping you find better jobs - even faster,” *Linkedin Official Blog,* March 22, 2010. http://blog.linkedin.com/2010/03/22/linkedin-profile-match.
[37] Barocas and Selbst, “Big Data,” 41.
[38] Stephanie Clifford and Jessica Silver-Greenberg, “Retailers Use Databases to Track Worker Thefts,” *The New York Times*, April 2, 2013. http://www.nytimes.com/2013/04/03/business/retailers-use-databases-to-track-worker-thefts.html.
[39] Pasquale, *Black Box*, 137.
[40] Pasquale, *Black Box*.
[41] Kelly Knaub, “LexisNexis, Retail Workers Get Nod For $2.38M Settlement,” *Law 360*, August 4, 2014. http://www.law360.com/articles/563583/lexisnexis-retail-workers-get-nod-for-2-38m-settlement.
[42] Alice Truong, “LinkedIn Hit With Class-Action Lawsuit For Its Job Reference Tool,” *Fast Company*, October 13, 2014. http://www.fastcompany.com/3037086/fast-feed/linkedin-hit-with-class-action-lawsuit-for-its-job-reference-tool.