Networked Employment Discrimination

Alex Rosenblat, Tamara Kneese, and danah boyd

Data & Society Working Paper, October 8, 2014

Prepared for: Future of Work Project supported by Open Society Foundations

Brief Description:

Employers often struggle to assess qualified applicants, particularly in contexts where they receive hundreds of applications for job openings. In an effort to increase efficiency and improve the process, many have begun employing new tools to sift through these applications, looking for signals that a candidate is “the best fit.” Some companies use tools that offer algorithmic assessments of workforce data to identify the variables that lead to stronger employee performance, or to high employee attrition rates, while others turn to third party ranking services to identify the top applicants in a labor pool. Still others eschew automated systems, but rely heavily on publicly available data to assess candidates beyond their applications. For example, some HR managers turn to LinkedIn to determine if a candidate knows other employees or to identify additional information about them or their networks. Although most companies do not intentionally engage in discriminatory hiring practices (particularly on the basis of protected classes), their reliance on automated systems, algorithms, and existing networks systematically benefits some at the expense of others, often without employers even recognizing the biases of such mechanisms. The intersection of hiring practices and the Big Data phenomenon has not produced inherently new challenges. While this paper addresses issues of privacy, fairness, transparency, accuracy, and inequality under the rubric of discrimination, it does not pivot solely around the legal definitions of discrimination under current federal anti-discrimination law. Rather, it describes a number of areas where issues of inherent bias intersect with, or come into conflict with, socio-cultural notions of fairness.

Detailed Topic Description:

Many large employers use (or rely on) staffing agencies that use Internet-based Vendor Management Systems (VMS) to procure both permanent employees and (more often) to contract with and manage contingent or temporary workforce labor (Johnson, 2014). Both large companies and staffing agencies use Applicant Tracking Systems (ATS) software to electronically recruit applicants and manage their information. ATS can score and sort resumes and other job application materials from a central database they compile or have access to; then they rank applicants in order to achieve the best fit between a job opening and available job candidates (Rangel, 2014). Recruiters don’t

generally see beyond the information that applicants input, although they can make additional queries. Many do; according to a recent Microsoft survey, 80% of employers search for job candidates on Google during their review (Citron, 2014a).

Large companies, particularly in the service and manufacturing industries, also use different kinds of Computer Business Systems (CBS) to micromanage singular employees and groups of employees by tracking and monitoring their actions (Skidelsky, 2014). For example, a CBS can act as digital control system for managing a large workforce by linking the workstations of employees together, storing information on their actions in central databases, and having a small group of managers activate and supervise automated management processes (Skidelsky, 2014). Because so many processes are digitized, employees produce much data about their actions and this data can be used for further analyses of the workplace. For example, Evolv, a recruiting software company, analyzed 3 million data points about 30,000 employees and identified that those who installed newer browsers, like Chrome or Firefox, onto their computers stayed at their jobs 15% longer than those who used default browsers that come pre-installed on their computers, like Safari (*Economist*, 2013).

Data-driven employment opportunities are changing the ways that candidates are assessed for employment. These can be direct, like psychometric tests, or indirect, through databases that host the names of people suspected of employee misconduct, which can cause candidates to be screened out from the application process. Candidates may also be negatively or positively assessed by the stray, personal data available on them through a Google search; some negative association might even be attributed to them by the advertisements that surface next to the search results.

Employers are interested in harnessing many kinds of data to improve efficiencies in their businesses, but it's not always clear what kinds of data are relevant to employment, or what they should be optimizing for. Increasingly, job candidates are scored in ways that go far beyond traditional notions of credentials, like educational certificates. Non-work related data is often used in these scoring algorithms, but it is not always clear what data is included or how much control users have over how they are being interpreted. As law professor Danielle Keats Citron (2014b) has pointed out, consumers are largely unaware that their online browsing activity could result in their resumes being sifted out of the eligible talent pool.

The design of data-driven hiring systems may be fraught with issues that pose a greater entry barrier for some applicants more so than others. For instance, there is some concern that particularly sensitive data, like health information, will factor into complex hiring algorithms that weed out employees whose real or perceived medical conditions represent a high-cost to their employers. While data mining tools can promote both inclusion and exclusion from job opportunities, current legal frameworks can make it challenging to remedy hiring tools that have a disparate impact on marginalized groups.

Applicant Tracking Systems (ATS)

While the Internet has made the circulation of job advertisements easier and more accessible to a wider number of people, the applicant tracking systems (ATS) that manage electronic recruitment have narrowed whose applications employers see. These recruiting mechanisms can be valuable in terms of efficiency for employers, but these systems also have a number of drawbacks.

Most broadly, large companies often rely on digital recruiting systems that only accept Internet-based job applications, which can act as a barrier to entry for people who do not have Internet access at home, or who must rely on libraries or other places with limited hours and access to computers to fill out applications. Although companies often rationalize Internet-based or Internet-only job application processes as a way to test for basic computer skills, numerous jobs do not require these skills, such as custodial, food service, and home healthcare personnel positions (Kelly, 2012). In this sense, Internet-based applications act as a misleading proxy for an intelligence test, which can end up becoming counter-productive. A teenager may be better able to fill out the application, but a candidate who did not grow up around computerized devices might be more competent at the job.

Outside of hiring tools that are developing more novel approaches to sorting and ranking applicants to promote specific goals, like workplace diversity, basic ATS' rely primarily on keyword searches to sift through resumes to match the keywords found in the job posting. A good resume or a job candidate with solid credentials by traditional standards may be poorly equipped to satisfy ATS software (Hansen, 2013). To complicate matters further, some ATS cannot scan italicized words, and others can only read text but not graphics. Not all ATS can read documents with narrow margins, formatting that cramps the text (Fertig, 2013), or documents saved in particular formats, like PDF or some versions of Microsoft Word, etc. (UIC Office of Career Services, n.d.). A failed applicant may never find out that the reason she wasn't selected for an interview was because of some small omission that didn't gel with the ATS (Pasquale, 2015, p. 42).

An ATS will sometimes incorporate pre-employment psychological, personality, cognitive, or other kinds of tests into their hiring process, particularly in screening potential retail employees. A 2009 survey found that these tests were used in 16% of retail hiring decisions (Pasquale, 2015, p. 43). The types of questions they ask and the selections that candidates make from the available responses may poorly reflect how capable they are of doing their job; still, this measurement will factor into how their application is scored, and can ultimately determine if they get an interview. For example, a question could ask if a potential employee ever thinks it's okay to steal, and with response options of: *Strongly Agree*, *Agree*, *Disagree* or *Strongly Disagree* (Ehrenreich, 2001, p. 29). A candidate who, for example, allows for contextual factors and answers 'Disagree' instead of 'Strongly Disagree' may be faulted for lack of integrity in their honest response—and be penalized for it—simply because the test is not designed to be cognitive of context.

Such pitfalls in the application-sorting process can in turn render the impression that there is a skills shortage in the eligible pool of candidates rather than a sorting system with flawed granularity or design (Barocas & Selbst, 2014, p. 46), or a disconnect in the management theory that relies on automated systems. However, it can be challenging to interrogate the variables that go into a calculation, or the design of the sorting mechanism, both of which are often treated with the safeguard of ‘trade secret’.

Wharton professor Peter Cappelli argues that data-driven employment is sometimes a deficient stratagem for recruiting labor because it feeds into the notion that with the right calculation, the perfect employee can be found. This is particularly problematic in a labor market with high unemployment. Cappelli contends that companies are leaving positions vacant for long periods of time as they search for candidates who fit a long list of micro-requirements, instead of opting to train employees for their jobs instead (Buchanan, 2012), although he supports workplace analytics that improve upon established HR practices (Lohr, 2014).

The founder of Entelo, a recruiting software company that conducts workforce analytics, lauds temp-to-perm hiring practices, such that employees and employers can test-drive each other before making permanent employment relationships by having a trial working arrangement first (Max, 2014). Hiring algorithms that can quickly produce new laborers with the right data profiles makes these practices easier to implement on a wide scale. Temp-to-perm hiring can mitigate vacancies and risk, but the practice generally undermines job security for workers who might be competing with several other ‘trial’ workers for a permanent position. An employee has to be an immediate rockstar in order to get a permanent position; in essence, they have to arrive pre-trained because an employer would rather be set-up with a great match than an untrained entry-level worker. Innovative hiring algorithms and practices are generally evolving in ways that support contingency and contract labor, and which empower employers to select from an increasingly analyzed labor pool.

Networked Information: Social Media and Hiring Practices

Applicant tracking systems (ATS) can come with a social recruiting component and a digital on-boarding process.¹ “Social recruiting” refers both to using social media platforms to locate passive potential job candidates through direct sourcing (thus circumventing the possibility that HR will be inundated with job applications) and the distribution of job postings on social media networks and platforms or through Vendor Management Systems (VMS) that market job opportunities.¹ While ATS systems may work primarily with the data that is inputted into its forms, the networked nature of information is raising the likelihood that employment algorithms will rely on new hordes of information on job candidates that will be sorted and categorized in even more obscure ways. According to a 2009 survey by Microsoft of HR professionals, recruiters, and consumers, 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 (Microsoft, 2009, p. 5). 89% of U.S. recruiters and HR professionals seek out professional online data, like LinkedIn, while 84% think it is appropriate to check online personal data as well (Microsoft, 2009, p. 7). Employers scan social media for signs or flags of negative behavior, but they can also use the information available from public profiles to discriminate against job candidates on the basis of protected class information, like religion (Valentino-DeVries, 2013).

To prevent employers from making hiring decisions on the basis of information they are prohibited from including in their job candidate assessments, a Fair Credit Reporting Act (FCRA)-compliant company called Social Intelligence offers to perform social media background checks on prospective employees on behalf of employers. Their Social Insight Executive Report blacks out information related to protected class categories, like race or age, from the written material it collects on job candidates' online reputations, although the reports include photos that reveal information like gender, sex, race, etc., so the efficacy of their method is debatable. The following is a list of negative filters they use to raise red flags about job candidates:

Negative Filters

- ☒ Animal Activism
- ☒ At Risk Populations
 - ☒ At Risk Profile
- ☒ Potentially Unlawful Activity
 - ☒ Photos, videos or references to drugs (All others)
 - ☒ Photos, videos or references to drugs (Marijuana)
 - ☒ Possible other unlawful activity
 - ☒ References to theft, fraud, or other unlawful activities
 - ☒ Solicitation
 - ☒ Underage alcohol consumption
- ☒ Potentially Violent Behavior
 - ☒ Aggressive online behavior/cyber bullying
 - ☒ Display or use of force or violence
 - ☒ Flagrant display of weapons
 - ☒ Other potential violent behaviors or actions
 - ☒ Potentially aggressive verbiage
- ☒ Racism and/or Demonstrations of Intolerance
 - ☒ Participation in groups or discussion boards that are potentially derogatory or intolerant against or towards a
 - ☒ Potentially derogatory behavior against or towards a protected group of people
 - ☒ Potentially derogatory photos or videos against or towards a protected group of people
 - ☒ Potentially derogatory verbiage against or towards a protected group of people
- ☒ Sexually Explicit Material
 - ☒ Audio
 - ☒ Pictures
 - ☒ Text
 - ☒ Video
- ☒ Social Media Policy Violation

Continue

Cancel Changes

Select All Filters

Screenshot taken by Data & Society 2014-08-22. Content from Social Intelligence Corp.

While the efficacy of their services is unclear, it *is* clear that job applicants are potentially subject to greater scrutiny of their private lives on the Internet than ever before. Crawford and Gillespie (In Press) recall one example in which a photo of two men kissing on Facebook was flagged as graphic sexual content, until protests from Facebook users, who pointed out that that straight kisses were not subject to the same treatment, prompted Facebook to remove the flag. As they 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" (p. 1). What an employer would see is that a job candidate has posted sexually explicit material on the Internet, rather than the values that went into the flag. Employers have to trust that the mechanisms of flagging are competent in order to sort through job

candidates efficiently, but this may come at a significant cost to the job prospects of candidates whose private lives have been communicated in negative ways to employers.

The persistence of negative information on the Internet and how it is sorted into databases is a departure from traditional methods of word-of-mouth referrals and other reputational tools. As law professor Frank Pasquale (2015) writes, “Unlike a recommendation letter written for one or a few readers, or a phone call that is almost never heard by anyone other than the callers, Internet-based rumors and lies are frequently persistent, searchable, replicable, and accessible to any decision maker with access to the right software or database” (p. 140). For example, Google search queries might turn up decontextualized or misleading contexts for a job candidate’s reputation online. Harvard researcher Latanya Sweeney’s (2013) study found that Google’s AdSense algorithm turns up ads that suggest possible arrest records when racially-associated names like DeShawn are queried, as opposed to Caucasian-associated names like Jill. While Google may not have intended to produce defamatory results in its ads, the ads may appear in response to the frequency that they are clicked on, such that the hidden biases of users who query names are made explicit through the mediation of an algorithm.

In an employment context, these ads might have a disparate impact on job candidates with names associated with certain races. Barocas and Selbst (2014) 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 (Barille, 2010). It may show candidates to potential employers that effectively reproduce a hidden bias, and which may result from how often employers click on the proffered candidates (Barocas & Selbst, 2014, p. 41). The recommendation engine can thus replicate the biases of employers overall without accounting for the harm that this does to the candidates who are rendered less visible.

A hiring algorithm that operates in a social media space could predict good matches based on the increasingly visible social networks that employers and job candidates have in common. Who you know is knowledge that is valued in Google’s recruiting process: Google asks job candidates to make a list of all the people they are already acquainted with who work there, presumably on the premise that similarly-networked individuals will be a good fit. Even though our social networks can be strong predictors of our behavior or situations, individuals who are not already ensconced in desirable social networks may find that they are excluded from job opportunities, which could disparately impact marginalized communities.

The use of hiring algorithms that rely on seemingly innocuous data to weed out higher-cost employees through effective proxies (or a simple lack of information on said employees) means that employers can use familiar scoring mechanisms to penalize or benefit job candidates, and do so by using rationales that constitute a black box of mysterious calculations to the job candidate (Pasquale, 2015, p. 137). The costs of

networked employment models can thus reproduce illegal or unpermitted (not illegal, but still uncomfortable or unethical) forms of discrimination against job candidates.¹

While these cases introduce possible mechanisms for abuse, identifying precise harm or holding companies accountable for their procedures turns out to be challenging. In short, too little is currently known about exactly how these processes are employed in the workforce. It's important to remember that although algorithms are potentially useful replacements for 'hunch'-based hiring, they are fallible in other ways that often get lost in the excitement about the Big Data phenomenon and its potential to automate and perfect scored hiring processes (Walker, 2012).

Data Analytics and Inequalities

ATS are characteristic or prototypical of what has been described as the 'scoring society,' a term that encompasses the way that much of daily life, along with specific activities at work, act as signals to feed into algorithmic systems that determine allocations of rewards or punishments to individuals (or their data-doubles) (Pasquale, 2015, p. 32). However, the logics of efficiency and optimization, as well as the promise of algorithmic systems to be neutral and fair mechanisms for assessing employability, inform the way that data analytics practices can amplify the problematic scoring systems inherent to ATS.

If employers only look at applicants above a certain score threshold, there is reason to be concerned about the ways in which algorithms might obfuscate opportunities through their sorting mechanisms, or through the databases and data they rely on to score candidates. The Big Data phenomenon is raising concerns that job candidates can be included or excluded in employment opportunities in ways that are unfamiliar, and potentially unfair.

When the often-impenetrable logics of scoring systems are implemented automatically, there's often no room for flexibility, as employers don't see additional applicant information to consider trade-offs. For instance, Richfield Management LLC, a waste-disposal company, has reduced workers' compensation claims by 68% by using a recruitment test to screen out applicants who represent a high disability risk (Walker, 2012). The test asks candidates to agree or disagree with statements like "When I'm working for a company I take pride in making it as profitable as possible." If a candidate scores poorly, the company does not hire them, without exception.

Using data analytics, Evolv, the recruiting software company discussed above, discovered that workers whose homes are far from their workplaces are more likely to quit their jobs. However, it deliberately excludes the distance a worker lives from their workplace from its hiring algorithms because of concerns that including such information could disparately impact the members of disadvantaged socio-economic

¹ Theoretically, employers could use the urinalysis results from routine drug tests to similarly determine medical conditions or pharmaceutical drug uses linked to certain disease treatment protocols.

communities (Volz, 2014). Evidently, data analytics can be a tool for inclusion as well as exclusion in employment opportunities, but creating accountability or incentives for measures that benefit inclusion rather than exclusion is challenging.

Not all data, or databases, are made equal, and the way they sort job candidates into positive or negative categories can be just as harmful as errant flags in social media background checks. Job candidates may be paying a high opportunity cost for inaccurate information, small errors, or flawed recruitment software, that affects how they score both in singular job applications, and over the course of their careers, even if these costs are not immediately visible. For example, First Advantage Corporation maintains the Esteem database that 'blacklists' employees who are accused of theft, usually in retail settings. As Clifford and Silver-Greenberg (2013) explain, First Advantage Corporation has tens of thousands of subscribers, including major retail companies like CVS and Target. In this system, it is the existence of allegations that matters, rather than the veracity of the claim. Employees are largely unaware that the 'records' of their alleged misconduct are inputted to a larger database that could prevent them from getting a job later on because companies may view them as potential problem candidates. The Fair Credit Reporting Act does not regulate this database, but the system nonetheless puts a permanent black mark on the records of employees. The accuracy and implications of these records can be harder to contest than the veracity of information in one's credit file because there are no legally-mandated procedures for doing so. From a business perspective, employers have a concerted interest in combatting employee theft, which accounts for about 44% of theft altogether in the retail industry, and back channeling information about problem employees is one method for doing so. However, the methods by which employers optimize for efficiencies can come into conflict with societal notions of fairness and due process.

Fuzzy Nudges and Diversity

Generally, relying on scoring systems to hire individuals to perform a discrete series of tasks can undermine efforts to form a diverse workforce. Best practices guides for diversity hiring emphasize cluster hiring, meaning that a group of people with complementary skills who are able to float between different areas of expertise are employed together. Cluster hiring can also minimize instances where a hiring practice aimed at integrating more diversity into the workforce effectively isolates individuals by their singular and representative presence (Guenter-Schlesinger & Kunle, 2009).

Typically, an ATS ranks job candidates by criteria such that, out of thousands of resumes submitted for a job posting, just the top 10 or 20 or perhaps only a few are read by a hiring manager, even though the ATS ranking may be a poor indication of the availability of qualified applicants in the pool. The top resumes that are put in front of the hiring manager effectively direct them towards those applicants at the exclusion of others. Google uses a recruitment algorithm that matches an applicant's resume with about 10 jobs based on the probability that the traits or experiences in their resume will match well with the job descriptions (Soltani, 2014, p. 57). Hiring managers are thus

nudged into selecting applicants from the pool of pre-scored candidates who appear on their desk(tops). However, these algorithms can obscure discrimination. While hiring managers may not be making decisions based on protected class information, an algorithm that sorts resumes according to fuzzy, non-transparent criteria may well be making correlations between protected-class characteristics like race or sex, and the traits that predict employee success (Soltani, 2014).

Entelo developed a 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. Employers can thus be nudged toward diversity hiring, although Entelo insists that employers should always recruit the best person for the job (Bischke, 2014). A diversity-algorithm may offer a convenient technical solution to a non-technical, social problem, which might make diversity easier to achieve quickly, but this outcome-orientation may not address the workplace cultures that are inimical to diversity, or the best ways to achieve it. In other words, erasing an embarrassing absence of diversity through a ‘hack’ will not automatically generate the conditions that breed equality in the workforce. A quick fix is potentially a disincentive to examine more closely the issues that create an environment that is hostile to diversity in the first place. As such, an algorithmic solution can support a framing fallacy: just because Entelo provides a diversity solution through recruiting software does not mean that the fundamental problem spurring workplace homogeny is the absence of a good algorithm. However, it’s possible, though not a guarantee, that having more diverse people at work will itself re-orient conversations around workplace diversity in meaningful ways. As well, the potentialities of data analytics will create room for debate on the best ways to do so.

Although individuals can be overwhelmed by the prospect of challenging a scoring algorithm or a database, data analytics can also be applied to identify patterns of negative bias in the workplace (Volz, 2014), and to identify the most effective areas for intervention. More data can create a broader and more granular view of labor conditions and relations, such that a seemingly innocuous item can be identified as having a greater predictive value for its impact on equality or inequality.

Employment and Health Data: Reflections on Potential Harms

The Big Data phenomenon generally implies that both public and more private, sensitive, localized, or otherwise restricted data are becoming more available for purposes that far exceed the original intent of their collection or use. There is some speculation that health data would be valuable to employers who could use it to assess whether job candidates will be high or low cost employees. Law professor Sharona Hoffman (2010) argues that employers, or the third-parties that employers rely on to fulfill staffing needs, could develop complex scoring algorithms that factor health needs into the cost of hiring or retaining a worker (p. 422), which could result in a disparate impact on workers with, or who develop, disabilities. According to Hoffman, over ten million workers sign authorizations releasing certain medical records to their employers

during the job application process (p. 415). The circumstances under which employers can request medical information, which kinds, and what they can do with that data is governed and restricted to some extent by US law. For example, it is illegal to make hiring decisions on the basis of genetic or disability information, with few exceptions.² But employers *are* sifting through these types of sensitive health data, and it is readily available to them—often with the inadvertent consent of employees and prospects. For instance, the urinalysis results from drug tests can also reveal certain medical conditions, like uncontrolled diabetes and pregnancy, as well as the presence of pharmaceutical drug uses linked to certain disease treatment protocols (Privacy Commissioner of Canada, 2012).

The health data that is collected can legally be used for a variety of purposes, including to assess a person's fitness for duty, reasonable accommodations in the workplace for disabilities, workers' compensation and insurance claims, and to process Family and Medical Leave Act requests (Hoffman, 2010, p. 415). In such situations, having this type of information on employment candidates can be risky for employers, especially if their hiring processes could be called into question legally. In the interest of self-protection, many prefer to simply be told which candidates would be costly to the company if employed without knowing the reasons *why*. Commenting on Hoffman's research, Pasquale (2015, p. 137) notes that an employer won't explicitly say that a candidate was not hired because of their diabetes. Employers may not be looking to deliberately violate existing laws, but they *are* looking to determine which of their current or prospective employees will be high-cost workers. Although they do not need to find high-tech solutions to do so, external scoring leaves employers less vulnerable to questions about their hiring practices, particularly if they do not appear to be actively looking for medical correlations.

In an ecosystem where data mining services are able to anticipate medical interventions (Taylor, 2014) based on credit card data or search engine results (Pettypiece & Robertson, 2014a), 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. Many types of sensitive medical conditions can be inferred from data that is not protected in its collection or dissemination by HIPPA in the first place, which can make it more accessible to employers, even if it is unverified. Data analytics practices of collecting, compiling, sorting, and creating easy access to huge amounts of information undermines restrictions that are put in place to ensure that information can be collected only for the purposes to which an individual initially consented. As boyd and Levy (2014) have described, networked harms stem from the sudden availability of large amounts of data on individuals that is gathered and shared

² There is a grey-zone between permitted and illegal discrimination when it comes to medical issues. For example, an employee could be fired for poor hygiene for emitting foul smells, but if they have halitosis and there is sufficient proof to render it representative of a disability, then they could potentially be protected from employment discrimination under law.

beyond their control. The problem is two-fold; in addition to the initial breach, there are no real legal remedies available for those whose privacy rights or expectations have been violated by a lapse in system-design or in the collapse of context around sensitive information. The sheer availability of information leaves many potential applicants vulnerable, even if some protections have been put in place.

Job candidates do not need to undergo direct screening to be implicated with an expensive disease that can potentially harm their employment prospects. For example, Dan Abate, a technology worker, had his name and address posted online under the header Sampler Diabetes Mailing List, even though he does not have diabetes. However, Acxiom Corp, a leading data broker that resells the information it compiles on millions of people to other data brokers, often to marketers, catalogued his information as ‘of diabetes interest’, which is how he wound up on a mailing list for diabetes (Pettypiece & Robertson, 2014b). While in a marketing context, the harm this misinformation and mis-categorization can have seems small, the line between marketing data and employment data, is blurry. If a Google search can turn up reputation data that gets flagged and harms job prospects, is there something about health data that is distinctly different in its production, dissemination, or prospective effect on employability?

Although penalizing potential employees for their probable or possible medical conditions may be viewed as unethical if not illegal, giving preferential treatment to candidates willing to track and voluntarily share medical information is not viewed as negatively. Quantified-self tools, like the Fitbit and Nike’s Fuelband, are being readily adopted by countless health-conscious individuals already. Increasingly, companies are already encouraging employees to adopt health and wellness programs that involve personalized health tracking (Hendrickson, 2014). The start-up company Appirio gave Fitbits to around 400 employees in 2013-2014, and based on the data it collected, it convinced its insurance provider to reduce their rates by 5%, which amounted to a savings of \$280,000 for the company; in 2013, BP America purchased Fitbits for 14,000 employees, 6,000 spouses, and 4,000 retirees (Bort, 2014). There is clearly a strong cost-savings incentive for companies to encourage employees to track and score their fitness activities, and potential health benefits to employees. Many companies have started to encourage adoption of these tools, offering discounts and opportunities to those who participate. It is not hard to imagine that candidates who are willing to offer such information might fare better in scoring systems than those whose data are missing.

Legal Challenges

It’s very difficult to challenge hiring algorithms and strategies that tout greater efficiency, productivity, and quantifiable qualities, especially when they give the impression of equalization or fairness because the barrier to entry is universally applied. In other words, it’s hard to assert that two people with a high score should not get the job compared to the person with a low or middling score, even if there are substantive inequalities inherent to the assessment tests that produce the lower-scoring candidate.

This fraught notion of fairness is reflected in legal precedent, particularly in the case of *Ricci v. DeStefano* (Barocas & Selbst, 2014, p. 44). In 2003, the city of New Haven held a civil service exam to determine promotions for its firefighters. One Hispanic and seventeen white firefighters scored disproportionately higher than their black counterparts, none of whom would merit promotion based on the results. As a result of this disparate impact, the city refused to certify the examination results, arguing that this would constitute a violation of Title VII of the Civil Rights Act of 1964. With the exam results called into question and ultimately not counted, none of the qualifying firefighters were promoted. They sued the city of New Haven as a group, and the case went up to the Supreme Court in 2009. The Court ruled in a 5-4 vote that the city could not ignore the test results because there was no strong basis in evidence that certifying the results would have made the city liable for disparate impact lawsuits. The city could have changed the test design before it was given in order to remedy potential disparate impact results, but effectively, using race as a reason to disqualify the results is discriminatory against the successful candidates (United States Department of Labor, n.d.).

If a legal precedent like *Ricci v. DeStefano* serves to inform how algorithms are held accountable for discriminatory effects resulting from a test, Barocas and Selbst (2014) argue that hiring algorithms (and others) will similarly confound attempts to *remedy* the disparate impact effects they might produce. It is difficult to find ways of correcting for prejudicial results in advance of any resulting disparate impact because the notion of what the correct classificatory label is for the information that algorithms sort is fuzzy (Calders & Zliobaite, 2013). Hiring algorithms are designed using historical data to create predictions about which qualities correlate to a strong job performance. If the historical data has an implicit bias that favors one group over another, then the measure of a strong worker may have a strong correlation to race, sex, or another protected class, even if the algorithm-designer has no intention of replicating past hiring decisions that marginalize groups of people based on these categories (Barocas & Selbst, 2014, p. 45-46). An algorithm could also rely on proxies for these, which could make it difficult for job candidates to hold employers accountable for intentional or unintentional discrimination.

Cases of employment discrimination are notoriously difficult to prove, even without algorithms to contend with. For instance, a 2004 survey of court cases in which plaintiffs brought suits against their employers under the Americans with Disabilities Act prevailed successfully in only 3% of cases (Allbright, 2004). Thus, it may be cost-effective for employers to continue using algorithmic calculations that could theoretically make them accountable for discrimination because the likelihood that they will be is so low. It can also be exceedingly difficult for employees or job candidates to prove that the discrimination was related to their consensual release of medical information in particular, especially if an employer does not act based upon that information immediately (Hoffman, 2010, p. 423). 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 (boyd & Levy, 2014). In the case of algorithmic hiring models, it would be challenging to detect where the precise harm is if the source of damage is located in proxies rather than in obvious points of contention, and it is more challenging still to determine who or which platform or network is liable for its functions or outcomes (given that a hiring algorithm would operate in co-ordination with other organizational mandates). At what point does the cost to individuals in negotiating their own employment prospects merit a re-evaluation of the system that causes their particular penalties (Barocas & Selbst, 2014)?

Conclusion

While a lot of the questions that networked employment discrimination provokes are speculative at this stage, they invite us to consider the ways in which such discrimination can be obscured or re-invigorated using algorithmic hiring practices. The networked nature of information signifies that employers potentially have access to more comprehensive electronic profiles on job candidates than has been traditionally available to them, which can expose job candidates to a greater scrutiny of their personal lives. Online reputations can be flagged in ways that are misleading, and even advertisements can create negative associations between job candidates in ways that are subtle, but present. Moreover, a small item in an employee's history may be the tipping point between a low and a high cost to an employer, but those factors are ill understood. It is unclear what employers are optimizing for in their attempts to maximize efficiencies and minimize risk and costs, nor is it easy to determine whether job candidates who successfully achieve good or clean scores are participating in a fair process. There are natural barriers to transparency in said process, and the argument against transparency is that if the rationales behind algorithmic calculations are released, applicants will try to "game the system." While workplace analytics and scoring represent opportunities for inclusion and exclusion in the workplace, the scoring and ranking systems they use to discriminate between job candidates requires careful consideration. Data analytics can be used to improve equality of outcome, but they can also mask hiring processes that contain prejudicial elements.

Acknowledgments

This document was produced as a part of the Future of Work Project at Data & Society Research Institute. This effort is supported by the Open Society Foundations' U.S. Programs Future of Work inquiry, which is bringing together a cross-disciplinary and diverse group of thinkers to address some of the biggest questions about how work is transforming and what working will look like 20-30 years from now. The inquiry is exploring how the transformation of work, jobs and income will affect the most vulnerable communities, and what can be done to alter the course of events for the better.

We are grateful to Patrick Davison, Tim Hwang, Rosa Jurjevics, Karen Levy, and Seth Young for their help with different aspects of this project.

Bibliography

- Allbright, Amy L. (2004). *Employment Decisions Under the ADA Title I – Survey Update*, 29 Mental & Physical Disability L. Rep. 513, 513 (July / August 2005).
- Barille, Parker. (2010). "Helping you find better jobs - even faster." *LinkedIn Official Blog*, March 22, 2010. <http://blog.linkedin.com/2010/03/22/linkedin-profile-match>
- Barocas, Solon, & Andrew Selbst. (2014). "Big Data's Disparate Impact." Available at SSRN: <http://ssrn.com/abstract=2477899>
- Bischke, Jon. (2014). "Announcing Entelo Diversity." Entelo. Accessed September 23, 2014. <http://blog.entelo.com/announcing-entelo-diversity>
- Bort, Julie. (2014). "This Company Saved A Lot Of Money By Tracking Their Employees With Fitbits." *Yahoo! Finance*, August 23, 2014. <http://finance.yahoo.com/news/company-saved-lot-money-tracking-005949318.html>
- boyd, danah, and Karen Levy. (2014). "Networked Rights and Networked Harms." Presented at Privacy Law School Conference (June 6, 2014) and Data & Discrimination (May 14, 2014). <http://www.datasociety.net/initiatives/privacy-and-harm-in-a-networked-society/>
- Buchanan, Leigh. (2012). "Unemployment is Up. Why is it so Hard to Find the Right Hires?" *Inc. Magazine*, June 1, 2012. <http://www.inc.com/leigh-buchanan/hiring-recruiting-unemployment-wharton-peter-cappelli.html>
- Calders, Toon & Indrė Žliobaitė. (2013). "Why Unbiased Computational Processes Can Lead to Discriminative Decision Procedures." In Bart Custers, Toon Calders, Bart Schermer & Tal Zarsky (Eds.). *Discrimination And Privacy In The Information Society: Data Mining And Profiling In Large Databases 3*, 43-57.
- Citron, Danielle Keats. (2014a). "Advocacy." Accessed August 22, 2014. <http://www.daniellecitron.com/advocacy.html>
- Citron, Danielle Keats. (2014b). "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>
- Clifford, Stephanie, and Jessica Silver-Greenberg. (2013). "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>
- Crawford, Kate and Tarleton L. Gillespie. (In Press). "What is a Flag for? Social Media Reporting Tools and the Vocabulary of Complaint." *New Media & Society*.
- Economist, The*. (2013). "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>
- Ehrenreich, Barbara. (2001). *Nickel and Dimed: On (Not) Getting by in America*. New York, NY: Metropolitan Books.
- Fertig, Arnie. (2013). "Ensuring Your Resume Avoids Applicant Tracking System Pitfalls." *U.S. News*, October 22, 2013. <http://money.usnews.com/money/blogs/outside-voices-careers/2013/10/22/ensuring-your-resume-avoids-applicant-tracking-system-pitfalls>
- Guenter-Schlesinger, Sue & Ojikutu Kunle. "Best Practices: Recruiting & Retaining Faculty and Staff of Color." *Western Washington University*. http://www.wvu.edu/eoo/docs/BestPractices_Recruiting_and_Retaining_Staff_of_Color.pdf

- Hansen, Katherine. (2013). "Have Applicant Tracking Systems (ATS) Ruined Recruiting, Hiring, and Job Search? A Quintessential Careers Annual Report 2013." *Quintessential Careers*.
http://www.quintcareers.com/applicant_tracking_systems_report.html
- Hendrickson, Barbara. (2014). "Increasing Employee Participation in Corporate Wellness Programs." *Occupational Health and Safety*, August 21, 2014.
<http://ohsonline.com/Articles/2013/09/01/Increasing-Employee-Participation-in-Corporate-Wellness-Programs.aspx>
- Hoffman, Sharona. (2010). "Employing E-Health: The Impact of Electronic Health Records on the Workplace." *Case Legal Studies Research Paper No. 2010-1*. Available at SSRN: <http://ssrn.com/abstract=1531265>
- Johnson, Craig. (2014). "David and Goliath: With the Right Strategy, Staffing Firms May Succeed Without Going Through a VMS." *Staffing Industry Review*, August 1, 2014.
<http://www.staffingindustry.com/row/Research-Publications/Publications/Staffing-Industry-Review/July-August-2014/David-Goliath-With-the-right-strategy-staffing-firms-may-succeed-without-going-through-a-VMS>
- Kelly, Lila. (2012). "The Hidden Face of Discrimination for Internet Applicants." <https://www.lilakelly.com/applicants>
- Lohr, Steve. (2013). "Big Data, Trying to Build Better Workers." *The New York Times*, April 20, 2013. <http://www.nytimes.com/2013/04/21/technology/big-data-trying-to-build-better-workers.html>
- Max, Sarah. (2014). "Uncertain About Hiring, Some Companies Try 'Test Drives.'" *The New York Times*, September 10, 2014.
<http://www.nytimes.com/2014/09/11/business/smallbusiness/uncertain-about-hiring-some-companies-try-test-drives.html>
- Microsoft. (2010). "Online Reputation in a Connected World." Online Safety Research.
<http://www.microsoft.com/security/resources/research.aspx#onlinerep>
- Pasquale, Frank. (2015). *The Black Box Society*. Cambridge, MA: Harvard University Press.
- Pettypiece, Shannon & Jordan Robertson. (2014a). "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>
- Pettypiece, Shannon & Jordan Robertson. (2014b) "Did You Know You Had Diabetes? It's All Over the Internet." *Bloomberg*, September 11, 2014. <http://www.bloomberg.com/news/2014-09-11/how-big-data-peers-inside-your-medicine-chest.html>
- Privacy Commissioner of Canada, The. (1990). "Drug Testing and Privacy."
https://www.priv.gc.ca/information/02_05_12_e.pdf
- Poon, Martha. (2012). *What Lenders See — a History of the Fair Isaac Scorecard* (Doctoral dissertation). University of California, San Diego, CA.
- Rangel, Lisa. (2014). "The Easy How-To Guide to Formatting Resumes for Applicant Tracking Systems." LinkedIn. Accessed August 12, 2014.
<http://premium.linkedin.com/jobsearch/articles/the-easy-how-to-guide-for-formatting-resumes-for-applicant-tracking-systems.html>
- Skidelsky, Robert. (2014). "The Programmed Prospect Before Us." *The New York Review of Books*, April 3, 2014. <http://www.nybooks.com/articles/archives/2014/apr/03/programmed-prospect-us/>
- Social Intelligence. (n.d.) *Hiring Product: Select Filters*. Accessed August 2014.
https://app.socialintel.com/flow/external/controller/configuration/select_filters

- Soltani, Ashkan. (2014). "Alternative Scoring Products." Federal Trade Commission Spring Privacy Series, March 19, 2014.
http://www.ftc.gov/system/files/documents/public_events/182261/alternative-scoring-products_final-transcript.pdf
- Sweeney, Latanya. (2013). "Discrimination in Online Ad Delivery." Data Privacy Lab.
<http://dataprivacylab.org/projects/onlineads/2013>
- Taylor, Colleen. (2014). "Omada Health Raises \$23M Series B Led By Andreessen Horowitz For Digital Health That Actually Works." *TechCrunch*, April 9, 2014. Accessed August 22, 2014.
<http://techcrunch.com/2014/04/09/omada-health-raises-23m-series-b-led-by-andreessen-horowitz-for-digital-health-that-actually-works/>
- UIC Office of Career Services. (n.d.) "Optimizing Resumes for Applicant Tracking Systems (ATS)." August 22 2014.
http://www.uic.edu/depts/ocs/docs/Optimizing_Resumes_for_Applicant_Tracking_Systems.pdf
- United States Department of Labor. (n.d.). "Ricci FAQs."
http://www.dol.gov/ofccp/regs/compliance/faqs/Ricci_FAQ.htm
- Volz, Dustin. (2014). "Silicon Valley Thinks It Has the Answer to Its Diversity Problem." *National Journal*. Accessed October 1, 2014. <http://www.nationaljournal.com/next-america/economic-empowerment/can-companies-use-big-data-to-fight-racist-and-sexist-hiring-practices-20140926>
- Walker, Joseph. (2012). "Meet the New Boss: Big Data." *Wall Street Journal*, September 20, 2012.
<http://online.wsj.com/news/articles/SB10000872396390443890304578006252019616768>