

References

- American Psychological Association. (2010). *Ethical principles of psychologists and code of conduct*. Washington, DC: American Psychological Association.
- Guzzo, R. A., Fink, A. A., King, E., Tonidandel, S., & Landis, R. S. (2015). Big data recommendations for industrial–organizational psychology. *Industrial and Organizational Psychology: Perspectives on Science and Practice*, 8(4), 491–508.
- McCune, E. A., Dekas, K., Anderson, A., Kasmatis Singleton, M., MacNiven, S., & Sinnett, S. (2015, April). *Guidelines for ethical research in the age of big data*. Panel presented at the 30th Annual Conference of the Society for Industrial and Organizational Psychology, Philadelphia, PA.
- Rousseau, D. M. (1995). *Psychological contracts in organizations: Understanding written and unwritten agreements*. Thousand Oaks, CA: Sage.

Big Data in I-O Psychology: Privacy Considerations and Discriminatory Algorithms

A. James Illingworth

Geode People, Inc., Decatur, Georgia

Introduction

The “big data” movement is forcing many fields to establish best practices for the collection, analysis, and application of big data, and the field of industrial–organizational (I-O) psychology is not exempt from this disruptive influence. Over the last several years, I-O scientists and practitioners have grappled with questions related to the definition, application, and interpretation of big data (e.g., Doverspike, 2013; Maurath, 2014; Morrison & Abraham, 2015; Poepelman, Blacksmith, & Yang, 2013). The focal article by Guzzo, Fink, King, Tonidandel, and Landis (2015) continues this discussion and represents one of the first attempts to establish a formal set of recommendations for working with big data in ways that are consistent with I-O psychology’s professional guidelines and ethics requirements.

The big data issues discussed by Guzzo et al. are not unique to I-O psychology. In fact, they overlap significantly with similar discussions occurring among computer scientists, technologists, privacy advocates, and policymakers about the challenges of maintaining privacy, informed consent, and analytical rigor in the big data era. That so many other fields are engaged in a similar conversation provides a tremendous opportunity for I-O psychology to draw on insights from this larger dialogue to shorten the big data learning curve, ensure alignment with current thought in other fields,

A. James Illingworth, Geode People, Inc., Decatur, Georgia.

Correspondence concerning this article should be addressed to A. James Illingworth, 2955 Fantasy Lane, Decatur, GA 30033. E-mail: ajillingworth@geodepeople.com

and enhance the development of its own professional practices and recommendations. Therefore, the purpose of this commentary is to provide the perspective of this larger community on two big data issues discussed by Guzzo et al.: privacy considerations associated with the use of big data and the potential for discriminatory algorithms derived from big data analyses.

Privacy Considerations: Data Collection Versus Usage

There is an active debate among technologists, privacy advocates, and policymakers about the best approach for protecting individual privacy in the big data era. At one end of the continuum are those who argue for preserving privacy during the collection and storage of big data; this perspective underlies almost all contemporary strategies for preserving data privacy and is the cornerstone for current domestic and international privacy laws (Executive Office of the President, 2012). Guzzo et al. describe this approach in their recommendations to establish data privacy plans and use informed consent. Informed consent provides people with control of their personal information at the point of data *collection*, whereas data privacy plans describe the aggregation of *stored* data, procedures for controlling data access, and methods to anonymize or remove personal and sensitive information from data.

However, managing personal information and data privacy with collection and storage strategies is difficult to implement in practice and may be obsolete in its current form given recent developments in big data analytics (Barocas & Nissenbaum, 2014; Kagal & Abelson, 2010). The use of informed consent or “notice and consent” cannot possibly account for the myriad ways in which a person’s original data will be used or the wide variety of “new data” (e.g., inferences, attributes) that arises from the integration of multiple data sets and their subsequent analysis (Mundie, 2014). Moreover, it would be impossible for people to authorize every request to use their data, whether collected voluntarily (e.g., agreeing to terms of use for web-based services), passively (e.g., GPS location tracking, video footage in public spaces), or through the recombinant analysis of integrated data sets, with one estimate suggesting that it would take approximately one month each year for individuals to review website privacy policies associated with their online activities (Masnick, 2012). In addition, it is well known that anonymity, the defining characteristic of most data privacy plans, can be and has been undermined through the integration and fusion of multiple data sets, which inadvertently reveals sensitive and personally identifying information that was otherwise not available in any individual file.

For these reasons, it has been suggested that the focus of big data privacy should shift from a sole emphasis on maintaining privacy during the collection and storage of big data to also mandating the responsible *use* of

big data (e.g., Kagal & Abelson, 2010; Mayer-Schonberger & Cukier, 2013; Mundie, 2014). The argument behind this approach is that data themselves are not harmful but could result in harm, depending on the context of their usage, when they are aggregated, analyzed, and interpreted for a specific purpose. This means that notions of privacy and the mechanisms for protecting privacy would transition from a model based on “privacy by consent” from the public to one of “privacy through accountability” for data users (Mayer-Schonberger & Cukier, 2013). Although proponents of this approach are still discussing the legal and technological infrastructure that would support usage-based privacy for data users, it generally includes transparency, accountability, enforcement, and privacy-preserving technology to facilitate the efficient sharing of data while at the same time maintaining security and increasing the amount of control people have over how their data are used (Kagal & Abelson, 2010; Lohr, 2015; Mayer-Schonberger & Cukier, 2013). It is important to note that use-based methods for preserving big data privacy are not intended to eliminate collection and storage controls; on the contrary, the hope is that they will enhance and complement them in a way that is less onerous for the public and data users.

There are several reasons why I-O psychology should consider use-based privacy protections for big data when developing its big data practice recommendations. First, this form of privacy protection has been endorsed by the Obama Administration (Executive Office of the President, 2014; President’s Council of Advisors on Science and Technology, 2014) and the World Economic Forum (2013) as the recommended method for handling big data privacy in the future. An endorsement is a long way from codification in public policies, regulations, and laws, but it does suggest that changes are looming in the requirements for managing data privacy that could impact I-O psychology’s responsibilities when handling big data. By maintaining an awareness of the larger dialogue around data privacy, I-O psychology can have a more informed discussion about data management practices and proactively respond to potential changes in the privacy landscape.

Second, I-O psychologists are in a unique position as both the collectors and users of big data, which is in direct contrast to professions that can be clearly distinguished as either collectors (e.g., data aggregators) or users (e.g., data analysts). As a result, I-O psychologists are responsible for maintaining privacy throughout the “data chain of custody.” Furthermore, responsibility for the usage part of the data chain is only going to grow in the coming years as I-O psychologists increasingly rely on data mining and predictive analytics to investigate and solve a wide range of talent management issues (e.g., recruiting, selection, engagement, retention). These analyses will be conducted on an ever-increasing number of disparate data sets containing all manner of data (e.g., structured, unstructured), each with the potential

to compromise individual or organizational privacy. Therefore, it is important that any discussion of recommendations for handling big data privacy include a consideration of data usage.

Finally, the rapid rise of big data and its resulting impact on the practice of I-O psychology may have exceeded the ability of existing professional guidelines to account for its influence. For example, the responsibility for privacy during data collection and use is implicitly covered by the American Psychological Association's (2010) ethical guidelines regarding confidentiality (Sections 4.01, 4.02) and the avoidance of harm to others (Section 3.01). These guidelines are sufficient for traditional data collection and storage privacy strategies, but it is unclear whether they account for privacy needs related to big data usage. This is primarily due to an expansion of the traditional "schema-based" research model that involves the collection of data to investigate specific a priori questions to the inclusion of "schema-less" approaches where research questions are developed post hoc after a wide range of data is collected or made available (e.g., Croll, 2012). In the future, as the big data privacy debate continues and use-based privacy controls potentially take hold, it may be necessary to more explicitly describe ethical requirements and expectations associated with using big data in research and applied practice.

Big Data Analyses: Potential for Discriminatory Algorithms

The utility of big data is not so much the data themselves but the identification of previously unknown relationships, inferences, and predictions obtained from the data through the use of data mining and algorithms. A review of any field, including I-O psychology, will find numerous examples of these big data analysis techniques solving previously intractable problems or developing new, unanticipated insights into more contemporary issues. Despite the many benefits resulting from data mining and algorithms based on big data, there is growing concern that these analysis tools may discriminate against protected classes (e.g., Barocas & Selbst, *in press*; Croll, 2012; Rieke, Robinson, & Yu, 2014; Schrage, 2014). Guzzo et al. briefly discuss several risks associated with using big data algorithms in employment contexts, such as more homogenous workplaces and reduced employment opportunities for subsets of the applicant pool, but do not mention the very real potential for these types of algorithms to violate employment antidiscrimination laws (i.e., Title VII of the Civil Rights Act of 1964) or the need for I-O psychology to establish practice recommendations specific to the development and evaluation of big data algorithms.

The idea of discriminatory algorithms is somewhat counterintuitive because I-O psychology has a long history of relying on mechanical techniques to increase the objectivity and utility of selection tools while at the same time

considering their adverse impact (e.g., De Corte, Lievens, & Sackett, 2007; Kuncel, Klieger, Connelly, & Ones, 2013). But an article by Barocas and Selbst (in press) highlights the need for I-O psychologists to seriously consider the discriminatory impact of big data algorithms on protected groups. The article is organized into three parts that discuss (a) how bias may inadvertently be incorporated into the data mining process for developing big data algorithms; (b) whether the data mining process and resulting algorithms violate antidiscrimination employment law, specifically as it relates to disparate treatment and disparate impact under Title VII of the Civil Rights Act of 1964; and (c) the difficulties in modifying employment discrimination law to account for the discriminatory challenges posed by data mining and its algorithms.

From their analysis, Barocas and Selbst reach two conclusions that are critically important to I-O psychologists. First, they indicate that at every point in the data mining process for supervised machine learning there are multiple opportunities for the resulting algorithm to become biased and produce outcomes that discriminate against protected groups. This may occur when defining the target variable of interest (e.g., potential for job success), collecting the algorithm training data and classifying them into examples of different groups (e.g., qualified vs. not qualified), and selecting the input variables to be included in the analysis (e.g., background information, assessment scores, job performance, training outcomes, social media data). For example, an algorithm's classification decisions could become biased if the training data from different subgroups are not representative (e.g., undersampling) or the data includes past instances of biased decision making (e.g., interview scores). As a result, the algorithm may simply reproduce the inherent bias in the data when evaluating and weighting relationships among variables or making classification decisions.

Second, Barocas and Selbst determine that it may be challenging to find liability under Title VII for the data mining process and resulting algorithms. In terms of disparate treatment, unless a variable representing a protected class is included as an input variable or there is evidence that decisions were made during the data mining process to intentionally discriminate against members of a protected class explicitly or through masking, it may be difficult to support a claim of disparate treatment given the nature of the data mining process. However, there is a somewhat better chance of establishing a claim of disparate impact because the data mining process and resulting algorithm could be evaluated in ways similar to traditional assessments (e.g., cognitive measures, physical ability tests). This would involve questioning the job-relatedness of the target variables included in the algorithm, evaluating evidence of their validity, and assessing the consideration of alternative variables that are potentially less discriminatory. Yet even in this situation it

may be difficult to support a disparate impact claim given the large number and type of variables that could be included in the analysis, many of which are likely to be job-related traits, characteristics, or attributes (although this must still be demonstrated), and the underlying goal of data mining to find significant relationships and maximize their predictions. Under both types of discrimination, it is difficult to establish liability because current antidiscrimination laws were not written to account for the data mining process and resulting algorithms.

The conclusions from Barocas and Selbst are compelling enough to suggest that I-O psychology should establish practice recommendations regarding the development and evaluation of big data algorithms. But there are additional reasons for supporting this initiative as well. In the future, the need for I-O psychologists to understand how the development of big data algorithms can lead to discrimination against protected groups is only going to increase as these types of algorithms are applied to every aspect of talent selection and assessment. In the area of recruitment, for example, some organizations are beginning to move away from traditional, passive methods of sourcing applicants to more proactive techniques that rely on big data algorithms to scan social media websites such as Facebook, LinkedIn, or Twitter and identify qualified applicants (e.g., Dormehl, 2014; Lam, 2015; Lohr, 2013; Miller, 2015; Preston, 2011; Richtel, 2013). As more organizations deploy these algorithms to recruit, assess, and select potential employees, I-O psychologists will find themselves responsible for developing and maintaining big data algorithms and providing guidance about the implications and risks of making employment decisions based on algorithmic results.

I-O psychologists will also have greater involvement in the evaluation of big data algorithms used to make employment decisions. At a very basic level, this will involve auditing algorithms developed by organizations or third-party vendors to ensure the data sources, analysis methods, results, and interpretations are accurate and valid (e.g., Mayer-Schonberger & Cukier, 2013), with a particular emphasis on the search for bias during the development process and consideration of discriminatory outcomes. More broadly, however, I-O psychologists must also be able to explain how the algorithms work. As the algorithms become more complex in terms of their data sources and analyses, they could develop into “black boxes” whose contents are unknown and operations are incomprehensible. In other applications of big data algorithms it may be sufficient to accept that they work without understanding how (e.g., marketing, retail), but in I-O psychology any tool used to make employment decisions must conform to specific guidelines and standards related to its development, validation, and implementation. If an employment outcome derived from a big data algorithm is challenged, I-O psychologists will be responsible for explaining how that outcome was

determined and what steps were taken to ensure the appropriateness of the data sources and variables on which the outcome was based.

Summary

Guzzo et al. have taken the first steps to develop formal recommendations regarding big data in I-O psychology. Although their efforts are important for informing and standardizing the field's approach to big data, there is an opportunity to elaborate further on these recommendations with insights from other fields engaged in a similar discussion. The experiences of these other disciplines can serve as benchmarks for I-O psychology to critically examine its own theoretical, analytical, and ethical approaches to big data and, where appropriate, revise them to incorporate this new information or address previously unknown shortcomings. This commentary attempted to demonstrate the merits of this approach by describing how I-O psychology's current understanding and handling of two big data issues, privacy and discriminatory algorithms, could be enhanced with input from other fields. The hope is that this information has broadened Guzzo et al.'s recommendations related to big data privacy and stimulated a new discussion about the recommendations that may be required as I-O psychologists become more involved in the development and evaluation of big data algorithms.

References

- American Psychological Association. (2010). *Ethical principles of psychologists and code of conduct*. Washington, DC: American Psychological Association.
- Barocas, S., & Nissenbaum, H. (2014). Big data's end run around anonymity and consent. In J. Lane, V. Stodden, S. Bender, & H. Nissenbaum (Eds.), *Privacy, big data, and the public good: Frameworks for engagement* (pp. 44–75). New York, NY: Cambridge University Press.
- Barocas, S., & Selbst, A. D. (in press). Big data's disparate impact. *California Law Review*. Retrieved from <http://ssrn.com/abstract=2477899>
- Croll, A. (2012, August). Big data is our generation's civil rights issue, and we don't know it: What the data is must be linked to how it can be used. *Radar*. Retrieved from <http://radar.oreilly.com/2012/08/big-data-is-our-generations-civil-rights-issue-and-we-dont-know-it.html>
- De Corte, W., Lievens, F., & Sackett, P. R. (2007). Combining predictors to achieve optimal trade-offs between selection quality and adverse impact. *Journal of Applied Psychology*, 92(5), 1380–1393.
- Dormehl, L. (2014, May). Your web presence just picked your next job. *Wired*. Retrieved from <http://www.wired.co.uk/magazine/archive/2014/05/features/web-presence-employment>
- Doverspike, D. (2013). A small, quick commentary on big data and predictive analytics. *Assessment Council News*, 1, 11–12.
- Executive Office of the President. (2012, February). *Consumer data privacy in a networked world: A framework for protecting privacy and promoting innovation in the global digital economy*. Washington, DC: The White House.

- Executive Office of the President. (2014, May). *Big data: Seizing opportunities, preserving values*. Washington, DC: The White House.
- Guzzo, R. A., Fink, A. A., King, E., Tonidandel, S., & Landis, R. S. (2015). Big data recommendations for industrial–organizational psychology. *Industrial and Organizational Psychology: Perspectives on Science and Practice*, 8(4), 491–508.
- Kagal, L., & Abelson, H. (2010, July). *Access control is an inadequate framework for privacy protection*. Paper presented at W3C Workshop on Privacy for Advanced Web APIs, London, England.
- Kuncel, N., Klieger, D. M., Connelly, B. S., & Ones, D. S. (2013). Mechanical versus clinical data combination in selection and admissions decisions: A meta-analysis. *Journal of Applied Psychology*, 98(6), 1060–1072.
- Lam, B. (2015, June). For more workplace diversity, should algorithms make hiring decisions? Some recruiters are hoping that software can somehow compensate for human failings. *The Atlantic*. Retrieved from <http://www.theatlantic.com/business/archive/2015/06/algorithm-hiring-diversity-HR/396374/>
- Lohr, S. (2013, April). Big data, trying to build better workers. *The New York Times*. Retrieved from <http://www.nytimes.com/2013/04/21/technology/big-data-trying-to-build-better-workers.html>
- Lohr, S. (2015). *Data-ism: The revolution transforming decision making, consumer behavior, and almost everything else*. New York, NY: HarperCollins.
- Masnick, M. (2012, April). To read all of the privacy policies you encounter, you'd need to take a month off from work each year. *Techdirt*. Retrieved from <https://www.techdirt.com/articles/20120420/10560418585/to-read-all-privacy-policies-you-encounter-you-d-need-to-take-month-off-work-each-year.shtml>
- Maurath, D. (2014). A critical incident for big data. *The Industrial–Organizational Psychologist*, 51(3), 15–25.
- Mayer-Schonberger, V., & Cukier, K. (2013). *Big data: A revolution that will transform how we live, work, and think*. New York, NY: Houghton Mifflin Harcourt.
- Miller, C. C. (2015, June). Can an algorithm hire better than a human? *The New York Times*. Retrieved from <http://www.nytimes.com/2015/06/26/upshot/can-an-algorithm-hire-better-than-a-human.html>
- Morrison, J. D., & Abraham, J. D. (2015). Reasons for enthusiasm and caution regarding big data in applied selection research. *The Industrial–Organizational Psychologist*, 52(3), 134–139.
- Mundie, C. (2014, March/April). Privacy pragmatism: Focus on data use, not data collection. *Foreign Affairs*. Retrieved from <https://www.foreignaffairs.com/articles/2014-02-12/privacy-pragmatism>
- Poepelman, T., Blacksmith, N., & Yang, Y. (2013). “Big data” technologies: Problem or solution? *The Industrial–Organizational Psychologist*, 51(2), 119–126.
- President's Council of Advisors on Science and Technology. (2014, May). *Big data and privacy: A technological perspective*. Washington, DC: Executive Office of the President of the United States.
- Preston, J. (2011, July). Social media history becomes a new job hurdle. *The New York Times*. Retrieved from <http://www.nytimes.com/2011/07/21/technology/social-media-history-becomes-a-new-job-hurdle.html>
- Richtel, M. (2013, April). How big data is playing recruiter for specialized workers. *The New York Times*. Retrieved from <http://www.nytimes.com/2013/04/28/technology/how-big-data-is-playing-recruiter-for-specialized-workers.html>

- Rieke, A., Robinson, D., & Yu, H. (2014). *Civil rights, big data, and our algorithmic future: A September 2014 report on social justice and technology* (Version 1.2). Washington, DC: Upturn.
- Schrage, M. (2014, January). Big data's dangerous new era of discrimination. *Harvard Business Review Magazine*. Retrieved from <https://hbr.org/2014/01/big-datas-dangerous-new-era-of-discrimination/>
- World Economic Forum. (2013). *Unlocking the value of personal data: From collection to usage*. Geneva, Switzerland: World Economic Forum.