Foretelling the Future

A Critical Perspective on the Use of Predictive Analytics in Child Welfare

By Kelly Capatosto, Research Associate

KIRWAN INSTITUTE RESEARCH REPORT | FEBRUARY 2017

Introduction

We are living in the era of Big Data. Our current ability to amass huge data sets combined with innovative methods of analysis has led to an unprecedented push toward the use of analytic tools in both the public and private sector. One of the most recent developments in the world of Big Data is the use of predictive analytics as a decision making tool, which has been described as a way to "predict the future using data from the past" (Davenport, 2014, p. 1). These predictions require analyses that sift through enormous sets of data in order to identify patterns. Although there is no standard method for the analysis, these predictions often rely on statistical algorithms and machine learning. Both the public and private sectors already employ predictive analytics to make key decisions in a variety of industries, including advertising, insurance, education, and, of particular interest, child welfare.

The field of child welfare has a long history of using risk analysis to guide institutional decision-making (Russell, 2015). Many in the field look toward predictive analytics as the next big innovation for understanding the risks associated with child maltreatment. Proponents of predictive analytics point to a variety of potential benefits, such as the ability to access hidden patterns, streamline service delivery, and decrease operation budgets. Beyond these benefits, the biggest push for predictive analytics use comes from the potential to prevent youth maltreatment before it occurs by identifying who is most likely to need care (Russell, 2015).

The Problem

While there has been a lot of enthusiasm surrounding predictive analytics and their possible benefit in the area of child welfare, others have begun to voice concerns regarding their use. As discussed in this white paper, there are reasons to be wary of the widespread use of predictive analytics. The risk of perpetuating cognitive and structural biases is among them. While this white paper does not to condemn the use of predictive analytics, it does hope to promote a critical assessment of these tools and the emergence of other Big Data applications.

A Perspective on Predictive Analytics that is Uniquely Kirwan

At the Kirwan Institute for the Study of Race and Ethnicity, our mission is to ensure that all people and communities have the opportunity to succeed. Through this work, the Institute developed a framework for analyzing inequity that considers both 1) cognitive and 2) structural barriers, defined below. In tandem, the operation of these barriers explains how inequity can persist in various institutions and systems, even in the absence of intentional prejudice or discrimination.

Cognitive Barriers: The role of individual-level thoughts and actions in maintaining structures of inequity.

Rather than focusing on explicit, intentional discrimination, the Kirwan Institute highlights the importance of implicit bias and other unconscious psychological processes. Generally, implicit bias is understood as the automatically activated evaluations or stereotypes that affect individual's understanding, actions, and decisions in an unconscious manner (Staats, 2013). All humans exhibit implicit bias, and having these biases does not reflect the intent to cause harm.

Structural Barriers: The influence of history on policies, practices, and values that perpetuate inequity (Davies, Reece, Rogers, & Rudd).

Although our society has made efforts to address racism, sexism, and other forms of discrimination, our nation's institutions remain rooted in a legacy of legally endorsed discrimination. For example, redlining—which purposefully devalued homes in minority neighborhoods by limiting access to financing—was a common practice until the Fair Housing Act of 1968 was passed (Olinger, Capatosto, & McKay, 2016). The enduring harmful impact of these practices is evident in the racial disparities found in the current housing landscape.

Thus, by considering both social forces—structural and cognitive, this white paper aims to do the following:

- Uplift concerns related to the use of predictive analytics in child welfare and other systems
- Examine concerns according to the inputs, outputs and application of predictive analytics
- Propose suggestions for the future of predictive analytics in child welfare

Concerns Related to the Use of Predictive Analytics in the Child Welfare System

Models of predictive analytics proceed in three stages. First, data goes into the model. Second, the model, with algorithms and/or statistical analyses, creates an output. Finally, individuals apply the model's outputs to decision-making at the field level. The following analysis critically examines concerns at each stage of this process—the inputs, outputs, and application of these models—regarding cognitive and structural factors that could be at play.

Inputs

Cognitive: Humans Encode Cognitive Biases into Machines

"Our own values and desires influence our choices, from the data we choose to collect to the questions we ask. Models are opinions embedded in mathematics."

– Cathy O'Neil, Weapons of Math Destruction, 2016

All humans rely on a variety of automatic mental processes to make sense of the world around us. One example of these processes is the operation of implicit biases. People are typically unaware of the implicit biases they possess, and these biases often do not align with one's explicit intentions to be egalitarian (Greenwald & Krieger, 2006; Nosek & Hansen, 2008). As such, all people make decisions that unintentionally rely on faulty or biased information. These decisions can have huge ramifications for our ability to safeguard opportunities for individuals of various genders, races, and ability statuses. To illustrate, one study demonstrated that resumes with White sounding names were nearly 50% more likely to get a call back than resumes with Black sounding names, despite controlling for all other factors, including work experience (Bertrand & Mullainathan, 2004).

As research continues to demonstrate how human bias can disrupt attempts to achieve equity, many have looked to the use of technology to ensure that decisions are made more objectively. Predictive analytics, like other data-based decision-making tools, have received considerable support for their potential to combat biases and provide opportunities for marginalized groups (Federal Trade Commission, 2016). However, human beings encode our values, beliefs, and biases into these analytic tools by determining what data is used and for what purpose. The data that institutions choose to use reveal what variables and reporting mechanism are valued most.

The advancement in technology has certainly improved the ability of child welfare systems to harness data to prevent child abuse and neglect. Large-scale data collection and reporting have made it easier than ever to manage cases and communicate among child welfare agencies. Despite this overall progress, the quality and consistency of the data used within the child welfare system remains an area of concern (Russell, 2015). Much of the data that predictive analytics tools use are derived from field-level reports. These may include self-reports on a variety of socioemotional factors or clinician reports on a child's background and intake experience (Commission to Eliminate Child Abuse and Neglect Fatalities, 2014). It is impossible to remove all subjectivity from personal reporting tools. Moreover, those who work in child welfare are often working in environments of high ambiguity, time constraints, and stress—all of which increase the likelihood of relying on implicit factors during decision-making (in general, see Mitchell, Banaji, & Nosek, 2003; Van Knippenberg, Dijksterhuis, & Vermeulen, 1999). As such, it is critical to acknowledge that human biases can limit the integrity of the data that informs predictive analytic models.

Structural: Previous Marginalization as a Predictor for Future Risk

"The math-powered applications powering the data economy were based on choices made by fallible human beings. Some of these choices were no doubt made with the best intentions. Nevertheless, many of these models encoded human prejudice, misunderstanding, and bias into the software systems that increasingly managed our lives."

– Cathy O'Neil, Weapons of Math Destruction, 2016

The adage "garbage in, garbage out" never holds truer than in the field of predictive analytics. When determining the inputs for these analyses, it is almost impossible to avoid incorporating longstanding patterns of inequity that exist in our society. Yet it may be difficult to identify where these patterns originate. For example, some predictive models currently used in the child welfare system assign a numeric indicator of risk associated with youth outcomes. In some cases, this risk informs how likely the child is to be reunited with the family, or how resilient the child is (Russell, 2015; Sledjeski, Dierker, Brigham, & Breslin, 2008; Toche-Manley, Dietzen, Nankin, & Beigel, 2013). Builders of such models would almost certainly avoid overt, illegal discrimination such as reliance on data that linked a family's race to the probability of child maltreatment or the assignment of a higher risk rating for racial minorities. However, because of past discrimination and historical inequities, subtle biases can emerge when seemingly "race neutral" data acts as a proxy for social categories. For example, data related to neighborhood characteristics are profoundly connected to historic practices of racial exclusion and discrimination. Thus personal information, such as one's zip code or diet, are deeply connected to racial identity (Sen & Wasow, 2016). Data that is ostensibly used to rate risk to child well-being can serve as a proxy for race or other past oppression, thereby over-representing those who have suffered from past marginalization as more risky. In this way, poor and marginalized communities are often disproportionately penalized by rating systems, even if the data feeding into these models would take considerable—perhaps, unreasonable—time and effort to alter, such as changing one's credit score or zip code (O'Neil, 2016).

Even more troubling is the *omission* of information for youth who do not enter the child welfare system as a counterbalance for these predictions of risk. It is impossible to know how many children who are never maltreated and whom would not properly be assessed as "high-risk" for maltreatment under these factors. This type of control group data simply does not (and should not) exist. Thus, it is important to acknowledge that the assumptions often built into data analytic models predict maltreatment with incomplete data.

Outputs

Cognitive: Overconfidence in the Objectivity of Outputs

“Data and data sets are not objective; they are creations of human design. We give numbers their voice, draw inferences from them, and define their meaning through our interpretations. Hidden biases in both the collection and analysis stages present considerable risks, and are as important to the big-data equation as the numbers themselves”

– Kate Crawford, The Hidden Biases in Big Data, 2016

The allure of predictive analytics is their potential for identifying and correcting for human biases that may arise during important child welfare decisions by lessening reliance on individual judgments. However, algorithms alone are no panacea to subjectivity. As discussed earlier, these models can unintentionally encode the same biases reflected in our society. Thus, one of the most serious dangers of predictive analytics is our overconfidence in the objectivity of their outputs.

When tools rely on vast quantities of data and complex analyses, it can be difficult or even impossible to be aware of the cognitive mechanisms influencing a model’s predictions. For example, if field workers are required to enter a particular score related to a child’s recidivism into the foster care system, it is highly unlikely that they will have the opportunity or authority to question the objectivity of that score at a later date, or exercise the discretion to make an exception to the score. For this reason, it can be very difficult to retroactively identify or correct instances where these seemingly objective outputs act as a gatekeeping mechanism—steering children and families toward service options that are not a suitable match for their individual needs.

Structural: Predictive Analytics Can Perpetuate Existing Structural Disparities

“The creators of these new models confuse correlation with causation. They punish the poor, and especially racial and ethnic minorities. And they back up their analysis with reams of statistics, which give them the studied air of evenhanded science.”

– Cathy O’Neil, Weapons of Math Destruction, 2016

Beyond relying on data inputs that reflect existing biases, predictive analytics may also exacerbate these structures of inequity through their outputs. For example, the tendency for algorithms to “digitally redline,” sometimes referred to as “weblining,” garnered significant attention from the federal government in the 2014 Big Data report (Executive Office of the President, 2014, p. 53). In the same way that ubiquitous redlining practices restricted loans and devalued homes in minority neighborhoods, weblining occurs when public and private institutions use opaque scoring algorithms to restrict communication and services to certain groups of people.

In the private sector, this most often occurs in the form of targeted advertising—matching consumers with products that the data reveals are most relevant to them. This matching process can involve various practices such as tracking consumers' purchase history or grouping consumers with similar attributes in order to offer goods and services that most likely match that population. Although largely viewed as a beneficial use of predictive analytics, targeted marketing has produced instances of racial discrimination. In one example, a test preparation program's marketing algorithm offered different pricing among geographic regions, which resulted in Asian families being subjected to higher prices (Angwin, Mattu, & Larson, 2015). In another, search engine searches for Black-sounding names were more likely to reveal advertisements offering arrest records than searches with White-sounding names (Sweeney, 2013). Importantly, no one intended for these algorithms to produce discriminatory outputs. Instead, as discussed above, these search processes learned the biases from user patterns, and then played a role in perpetuating these biases through consumer behavior.

A highly controversial example of predictive analytics playing out in the public sector is targeted policing (Executive Office of the President, 2014). While some applaud the use of predictive models for focusing limited resources on high crime areas, the practice has been criticized for justifying an outsized police presence in poor neighborhoods with large minority populations (For a general overview, see joint statement on predictive policing: American Civil Liberties Union, 2016). Some predictive policing efforts have gone so far as to form lists and engage in active surveillance of those deemed most likely to commit a crime (Executive Office of the President, 2014). Individual-level predictions such as these are more likely to target people for who they are (race, proximity to crime, class, education level, etc.…) rather than on the basis of observable behavior (O'Neil, 2016).

Predictive analytics tools in child welfare can operate the same way as the predictive policing scenario—by classifying individuals and families based on individual risk profiles for maltreatment. To illustrate, one predictive analytic tool utilized data from youth self-reports to determine the variables most related to youth resiliency. Youth received a resilience score based on 11 indicators in order to assist service workers in developing their treatment goals (Toche-Manley et al., 2013). Even though the identification of these risk factors is empirically valid, research has yet to show the link between these resiliency scores and treatment outcomes. Thus, this type of scoring may have the potential to impose a punitive system of gatekeeping on less-resilient youth who are denied opportunities more resilient youth are routinely offered. This is just one example of predictive analytics efforts, though research-based, that may not generalize into effective field use. Moreover, if tools such as these *do* get utilized in the field, their application may actually perpetuate existing structural disparities by restricting necessary services to certain families or neighborhoods.

How Predictive Analytics Impact Decisions in Child Welfare

The prior sections of this paper addressed *what* these models encode and produce. This final section will discuss the importance of *how* these tools can influence decision-making in child welfare systems that reproduce inequity. Predictive analytics are already governing real-world decision-making across many social services fields; many of these decisions literally involve life and death. Although the Kirwan Institute strongly supports the use of data and empirical research to inform organizational behavior, it is essential to remain cautious to the potential drawbacks of how predictive analytics are applied.

First, when relying on predictive analytic models, users can fall into the common trap of confusing correlation with causation. For example, one analysis conducted in New Zealand analyzed all maltreatment cases for five-year-olds and discovered that 83% had been enrolled in the public benefit system before they were two; they concluded that the receipt of public benefits predicted future child maltreatment (Vaithianathan, Maloney, et al., 2013, p. 354). Similarly, other studies have concluded that prior experience with child protective services (CPS) was the best predictor of recurrent maltreatment (DePanfilis & Zuravin, 1999; Fluke, Shusterman, Hollinshead, & Yuan, 2008; Sledjeski et al., 2008). While this information is empirically valid and instructive for understanding the recurrent nature of child welfare involvement, it does not provide information on the circumstances that determined the families' need for these services. Moreover, these analyses cannot conclude why youth who a) received public benefits or had CPS contact and b) did not experience maltreatment, fared better than some of their peers. These findings illustrate a classic example of the adage that correlation does not equal causation; public benefits and CPS involvement acted as confounding variables because both were highly correlated to later child maltreatment, but were not the underlying cause. In this case, the lack of a causal relationship between these risk factors and the outcome of maltreatment should seem apparent. **However, other variables identified by these models are not held to the same standard of scrutiny; many of these so-called risk factors are then targeted as a potential point of intervention without ever knowing whether they contribute to outcomes in any meaningful way**. Equally troublesome is the possibility that these models will comb through vast quantities of data only to reveal what the child welfare system has known for decades—poverty and lack of opportunity are detrimental to families. For example, if a predicative analytics model reveals a risk factor: R, it is necessary to evaluate whether R is truly the source of maltreatment: M (i.e. $R \Rightarrow M$), or if R is just another product of an underlying variable, such as poverty: P (i.e. $P \Rightarrow R$ & M).

Child welfare systems have a clear obligation to invest in the most effective methods for mitigating the risk of youth maltreatment. Thus, it is necessary to take a deeper look into the process by which predicative analytic models inform decision-making, especially if they take the place of other decision-making tools (e.g. experimental research literature, staff surveys, etc.) to determine families' access to benefits that exist to promote child well-being. In short, while predictive analytics can help identify important intervention points and patterns of targeted need, it may prove difficult (or even potentially impossible) to reap the full benefits of predictive analytics without first addressing systemic factors such as poverty and discrimination. For these reasons, this paper closes by recommending ways to utilize predictive analytics within the context of a historically informed framework for understanding social inequities.

Suggestions for the Future of Predictive Analytics in Child Welfare

Although this paper casts a wary eye on predictive analytics, these Big Data tools still have much to offer the field of child welfare. They can reveal patterns of social disparities, and help determine the most effective use of limited public and private resources. Nevertheless, it is important to be aware that longstanding and deeply embedded systemic and cognitive inequities can limit the effectiveness of data analytic tools and the conclusions they reach. Thus, when institutions utilize Big Data tools they should be conscious of how the models interact with pre-existing structures of and barriers to opportunity. Thus, the following recommendations are not directed toward reforming predictive analytics in general. Instead, these suggestions focus on ways to help safeguard child welfare agencies and the populations they serve from the misuse of predictive analytic tools.

Develop a Code of Ethics

Among other impacts, predictive analytics use have immense social, legal, and financial ramifications. As such, it may prove beneficial to develop a comprehensive code of ethics to help guide how the field uses predictive analytics and other Big Data applications. Child welfare systems involve a wide range of actors with varying expertise. Thus, to ensure that multiple perspectives weigh in on the ethical considerations of predictive analytics, an interdisciplinary committee or task force should be formed. Examples of potential representatives include human services employees, computer scientists, and social science researchers (For table on potential interdisciplinary connections for Big Data projects, see Staab, Stalla-Bourdillon, & Carmichael, 2016, pp. 24-25). Moreover, the implementation of child welfare services often involves multiple levels of governance. Thus, to ensure the consistency of ethical standards, the committee should focus on general guidelines at the state and national level, while also representing local interests and concerns whenever possible.

Increase Accountability

"Black box" algorithms are characteristically difficult to understand; the inputs and outputs are observable but the internal processes are ambiguous and complex (Staab et al., 2016, p. 7). Predictive analytics tools that utilize black box algorithms have the potential to restrict transparency and accountability in decision-making.

If families are denied services or interventions based on a predictive analytics algorithm, they should be able to understand what factors contributed to that outcome and have a recourse for disputing that decision—especially as these decisions are so vital for determining outcomes for families and children in need. Additionally, if child welfare agencies are unsure of how these algorithms operate, they may find themselves in vicious loop where those working on the ground are unable to provide the appropriate feedback to help these models improve and adapt.

Assess Equity Impact

As part of a comprehensive effort to address disparities in child welfare, predictive analytic models should undergo an evaluation to gauge their equity impact. Simply identifying accurate predictive factors and using these factors to make decisions about service delivery does not guarantee the interventions will be implemented equitably. As such, those who determine the benefits of predictive analytics should be trained to look for existing structures of inequity that may limit the effectiveness of the resulting interventions. For example, racial groups have different experiences when encountering health and social services professionals. Research demonstrates that practitioners' implicit racial biases may lead to different quality interactions and treatment decisions for Black and White patients (Green et al., 2007; Johnson, Roter, Powe, & Cooper, 2004; Penner et al., 2010). Awareness that practitioner bias can inhibit child welfare interventions should be factored into the operation of application of predictive analytics model.

Broaden the Scope

As the prior sections noted, individual-level risk predictions depends on a variety of factors, many of which are difficult or impossible to control (e.g. race, education, geography, and socioeconomic status). Thus, those who promote the use of predictive analytics in child welfare should consider broadening the scope to include neighborhood and citywide predictions. For example, tools like opportunity mapping can identify geographic regions that would benefit most from additional resources. By targeting neighborhoods for child welfare interventions rather than families alone, it is easier to combat the systemic sources of these risk factors such as poverty and lack of opportunity. Moreover, it is important to acknowledge that neighborhood-level and individual interventions are not in opposition to one another; addressing poverty for a whole community can bolster family-level efforts to promote child welfare. For example, a neighborhood-targeted approach may focus on preventative efforts (such as workforce development programs or a public health campaign to decrease teenage pregnancy), in addition to individual-level supports like therapeutic services for families. ■

References

American Civil Liberties Union. (2016). ***Statement of concern about predictive policing by ACLU and 16 civil rights privacy, racial justice, and technology organizations***. Retrieved from https://www.aclu.org/other/statement-concern-about-predictive-policing-aclu-and-16-civil-rights-privacy-racial-justice

Angwin, J., Mattu, S., & Larson, J. (2015). ***The tiger mom tax: Asians are nearly twice as likely to get a higher price from Princeton Review***. Retrieved from https://www.propublica.org/article/asians-nearly-twice-as-likely-to-get-higher-price-from-princeton-review

Bertrand, M., & Mullainathan, S. (2004). Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. *The American Economic Review*, *94*(4), 991–1013.

Commission to Eliminate Child Abuse and Neglect Fatalities. (2014). T*he dissenting report of the honorable Judge Patricia M. Martin CECANF commissioner*. Retrieved from http://nccpr.org/reports/JudgeMartinDissent.pdf.

Crawford, K. (2013). The hidden biases in Big Data. *Harvard Business Review*. Retrieved from https://hbr.org/2013/04/the-hidden-biases-in-big-data.

Davenport, T. H. (2014). A predictive analytics primer. *Harvard Business Review*. Retrieved from https://hbr.org/2014/09/a-predictive-analytics-primer .

Davies, S., Reece, J., Rogers, C., & Rudd, T. (n.d.) Structural racialization: A systems approach to understanding the causes and consequences of racial inequity. *The Kirwan Institute*. Retrieved from http://kirwaninstitute.osu.edu/docs/NewSR-brochure-FINAL.pdf

DePanfilis, D., & Zuravin, S. J. (1999). Predicting child maltreatment recurrences during treatment. *Child Abuse & Neglect*, *23*(8), 729–743.

Executive Office of the President. (2014). Big data: Seizing opportunities, preserving values. *The White House*. Retrieved from https://www.whitehouse.gov/sites/default/files/docs/big_data_privacy_report_5.1.14_final_print.pdf.

Ramirez, E., Brill, J., Ohlhausen, M.K., & McSweeny, T. (2016). Big Data: A tool for inclusion or exclusion. *Federal Trade Commission*. Retrieved from https://www.ftc.gov/system/files/documents/reports/big-data-tool-inclusion-or-exclusion-understanding-issues/160106big-data-rpt.pdf.

Fluke, J. D., Shusterman, G. R., Hollinshead, D. M., & Yuan, Y.-Y. T. (2008). Longitudinal analysis of repeated child abuse reporting and victimization: Multistate analysis of associated factors. *Child Maltreatment*, *13*(1), 76–88.

Green, A. R., Carney, D. R., Pallin, D. J., Ngo, L. H., Raymond, K. L., Iezzoni, L. I., & Banaji, M. R. (2007). Implicit bias among physicians and its prediction of thrombolysis decisions for Black and White patients. *Journal of General Internal Medicine*, *22*(9), 1231–1238.

Greenwald, A. G., & Krieger, L. H. (2006). Implicit bias: Scientific foundations. *California Law Review*, *94*(4), 945–967.

Johnson, R. L., Roter, D., Powe, N. R., & Cooper, L. A. (2004). Patient race/ethnicity and quality of patient-physician communication during medical visits. *American Journal of Public Health*, *94*(12), 2084–2090.

Mitchell, J. P., Banaji, M. R., & Nosek, B. A. (2003). Contextual variations in implicit evaluation. *Journal of Experimental Psychology: General*, *132*(3), 455–469.

Nosek, B. A., & Hansen, J. J. (2008). The associations in our heads belong to us: Searching for attitudes and knowledge in implicit evaluation. *Cognition and Emotion*, *22*(4), 553–594.

O'Neil, C. (2016). *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. New York: Crown Publishers.

Olinger, J., Capatosto, K., & McKay, M. A. (2016). Challenging race as risk: Implicit bias in housing. *The Kirwan Insitute*. Retrieved from http://kirwaninstitute.osu.edu/my-product/challenging-race-as-risk-implicit-bias-in-housing/.

Penner, L. A., Dovidio, J. F., West, T. V., Gaertner, S. L., Albrecht, T. L., Dailey, R. K., & Markova, T. (2010). Aversive racism and medical interactions with Black patients: A field study. *Journal of Experimental Social Psychology*, *46*(2), 436–440.

Russell, J. (2015). Predictive analytics and child protection: Constraints and opportunities. *Child Abuse and Neglect*, *46*, 182–189.

Sen, M., & Wasow, O. (2016). Race as a bundle of sticks: Designs that estimate effects of seemingly immutable characteristics. *Annual Review of Political Science*, *19*, 499–522.

Sledjeski, E. M., Dierker, L. C., Brigham, R., & Breslin, E. (2008). The use of risk assessment to predict recurrent maltreatment: A classification and regression tree analysis (CART). *Prevention Science, 9*(1), 28–37.

Staab, S., Stalla-Bourdillon, S., & Carmichael, L. (2016). Part I: Observing and recommending from a social web with biases. *Web Science Institute*. Retrieved from https://arxiv.org/ftp/arxiv/papers/1604/1604.07180.pdf.

Staats, C. (2013). State of the science: Implicit bias review. *The Kirwan Institute*. Retrieved from http://www.kirwaninstitute.osu.edu/reports/2013/03_2013_SOTS-Implicit_Bias.pdf.

Sweeney, L. (2013). Discrimination in online ad delivery. *Queue, 11*(3), 10.

Toche-Manley, L. L., Dietzen, L., Nankin, J., & Beigel, A. (2013). Revolutionizing child welfare with outcomes management. *The Journal of Behavioral Health Services & Research*, *40*(3), 317–329.

Van Knippenberg, Dijksterhuis, A., & Vermeulen, D. (1999). Judgement and memory of a criminal act: The effects of stereotypes and cognitive load. *European Journal of Social Psychology*, *29*(2), 191–201.

This publication was produced by the Kirwan Institute for the Study of Race and Ethnicity at The Ohio State University. As a university-wide, interdisciplinary research institute, the Kirwan Institute works to deepen understanding of the causes of—and solutions to—racial and ethnic disparities worldwide and to bring about a society that is fair and just for all people.

Kirwan Institute research is designed to be actively used to solve problems in society. Its research and staff expertise are shared through an extensive network of colleagues and partners—ranging from other researchers, grassroots social justice advocates, policymakers, and community leaders nationally and globally, who can quickly put ideas into action.

For More Information

The Kirwan Institute for the Study of Race and Ethnicity at The Ohio State University is known and respected nationally and deeply engaged in social issues. We are focused on projects that are integrated with sound research, strategic communication, and advocacy. To learn more, visit www.kirwaninstitute.osu.edu.

THE OHIO STATE UNIVERSITY
KIRWAN INSTITUTE FOR THE
STUDY OF RACE AND ETHNICITY