ALGORITHMIC BIAS EXPLAINED

How Automated Decision-Making Becomes Automated Discrimination

THE GREENLINING INSTITUTE

Table of Contents

Introduction

Over the last decade, algorithms have replaced decision-makers at all levels of society. Judges, doctors and hiring managers are shifting their responsibilities onto powerful algorithms that promise more data-driven, efficient, accurate and fairer decision-making. However, poorly designed algorithms threaten to amplify systemic racism by reproducing patterns of discrimination and bias that are found in the data algorithms use to learn and make decisions.

> "We find it important to state that the benefits of any technology should be felt by all of us. Too often, the challenges presented by new technology spell out yet another tale of racism, sexism, gender inequality, ableism and lack of consent within digital culture."[1]
>
> —Mimi Onuoha and Mother Cyborg, authors, "A People's Guide to A.I."

The goal of this report is to help advocates and policymakers develop a baseline understanding of algorithmic bias and its impact as it relates to socioeconomic opportunity across multiple sectors. To this end, the report examines biased algorithms in healthcare, at the workplace, within government, in the housing market, in finance, education and in the pricing of goods and services. The report ends by discussing solutions to algorithmic bias, explores the concept of algorithmic greenlining and provides recommendations on how to update our laws to address this growing problem.

What Are Algorithms?

An algorithm is a set of rules or instructions used to solve a problem or perform a task. This report looks at computer algorithms designed to make predictions and answer questions. These algorithms, otherwise known as automated decision systems, use statistical patterns and data analytics to make decisions that impact everyday life.

Algorithms are used to:

- Sort resumes for employers
- Decide access to and eligibility for social services
- Determine who sees advertisements for jobs, housing or loans
- Choose which employees to hire, fire and promote
- Determine access to housing and credit
- Predict crime and risk of recidivism
- Decide hospital treatment plans and access to insurance

How Do They Work?

Simply put, decision-making algorithms work by taking the characteristics of an individual, like the age, income and zip code of a loan applicant, and reporting back a prediction of that person's outcome—for example, the likelihood they will default on the loan—according to a set of rules. That prediction is then used to make a decision—in this case, to approve or deny the loan.[2]

Algorithms often learn the rules for making predictions by first analyzing what's known as "training data" to discover useful patterns and relationships between variables. The patterns or algorithmic insights gleaned from the training data become the basis for rules governing future decisions and predictions.

However, if the training data is biased then the algorithm can pick up on that pattern of discrimination and replicate it in future decisions. For example, a bank's historical lending data may show that it routinely and unfairly gave higher interest rates to residents in a majority Black ZIP code. A banking algorithm trained on that biased data could pick up on that pattern of discrimination and learn to charge residents in that ZIP code more for their loans, even if they don't know the race of the applicant.

What is Algorithmic Bias and Why Does it Matter?

Algorithmic bias occurs when an algorithmic decision creates unfair outcomes that unjustifiably and arbitrarily privilege certain groups over others. This matters because algorithms act as gatekeepers to economic opportunity. Companies and our public institutions use algorithms to decide who gets access to affordable credit, jobs, education, government resources, health care and investment. Addressing algorithmic bias, particularly in critical areas like employment, education, housing and credit, is critical to closing the racial wealth gap.

AI Harms[3]

Individual Harms Illegal Discrimination | Unfair Practices	Collective Social Harms
Hiring Employment Insurance and Social Benefits Housing Education	**Loss of Opportunity**
Credit Differential Prices of Goods	**Economic Loss**
Loss of Liberty Increased Surveillance Stereotype Reinforcement Dignitary Harms	**Social Stigmatization**

Is Algorithmic Bias Illegal?

There is no general law in the United States that specifically prohibits algorithmic bias. However, various state, federal and local laws generally prohibit discrimination, whether by algorithms or by humans, in some contexts, such as credit, education, housing and employment. Two key legal concepts that animate anti-discrimination laws are: (1) disparate treatment and (2) disparate impact.[4]

Disparate Treatment: A decision-maker is liable for disparate treatment discrimination when they have a discriminatory intent or motive, or use protected attributes (such as gender or race) as the basis for a decision. A hiring algorithm commits disparate treatment if it rejects female job candidates because they are female.

Disparate Impact: A decision-maker commits disparate impact discrimination if a facially neutral policy or decision has an unjustifiably disproportionate adverse impact on a protected class of individuals.[5] A policy is facially neutral if it does not use any protected attribute such as race or gender as the basis for the decision. So a hiring algorithm that rejects all candidates below six feet tall may create a disparate impact on female applicants, but there is no disparate treatment because height is not a protected characteristic.

Existing anti-discrimination laws do not work well in the algorithmic context as most were written before the internet was even invented. Furthermore, proving bias is difficult as people negatively affected by discriminatory algorithms may not know why or how that decision was made—or even that an algorithm was behind the decision that harmed them. Our anti-discrimination laws must be updated and adapted to properly regulate algorithmic bias and discrimination.[6] As discussed further below, one way to prevent or reduce unfair bias is to pass algorithmic transparency and accountability laws.

The Value of Algorithmic Transparency: Proving Disparate Treatment and Impact

It's hard to prove discrimination. Take a disparate treatment claim in hiring: An employer can lie about the motivations behind their choice not to hire any Black job applicants. Even if we have statistical evidence that there is a disparate impact in hiring, a court will generally look at the motivations of the decision-maker and whether there was a legitimate (i.e. non-race based) reason that explains the disparate impact and racial disparity in hiring. From this example, we can see that humans can get away with discrimination if they simply lie about their motivations or if they can point to any alternate justification for why they made a discriminatory decision. In some cases, implicit bias means people may not even realize the biased motivations that drive their decisions.

In comparison to humans, algorithms can be much more honest because their choices are based on math and data rather than fuzzy human logic and emotions. With proper algorithmic transparency and record-keeping practices, a court or regulator can know the factors an algorithm considered when making a decision and how the algorithm and its designers approached tradeoffs between relevant values like accuracy, profit maximization and non-discrimination. In order to do this a court needs access to three things: (1) the algorithm, (2) the training data and (3) the objective function of the algorithm, which is the outcome that the algorithm is attempting to optimize or predict.[7]

Where Does Algorithmic Bias Come From?

Broadly speaking, algorithmic bias arises from the choices developers make in creating the algorithm, rather than an explicit discriminatory motive.[8]

Choosing an Outcome

One of the critical choices in building an algorithm is deciding what outcome it is designed to predict. Choosing an outcome requires subjective value judgments about how to define amorphous concepts like productivity or creditworthiness in measurable and quantifiable ways. This is a key choice because it shapes what the algorithm is optimized to do and the types of data that it considers when making a decision. Bias can enter into the algorithm when designers choose outcomes that favor certain groups over others.[9]

Consider a simple worker performance algorithm designed to predict which employees are most productive. If a designer chose to measure productivity by number of hours worked, as opposed to other measures such as number of emails sent, it could create a disparate impact and disadvantage women who face higher childcare burdens.

Choosing the Data Inputs

After choosing an outcome, designers select the inputs or predictor variables an algorithm can consider when trying to predict that outcome. The choice around what training data and what variables an algorithm has access to can also introduce bias if the designers only give the algorithm data that is more favorable to one group than another or uses subjective data that is biased or mismeasures objective reality (such as performance reviews by a biased supervisor).

Another issue is that data inputs which are relevant to making accurate decisions may be proxies for protected attributes like race or gender due to structural and historical bias. For example, income data can show men earn more than women, and arrest data may reveal that Black men are arrested at higher rates than White men. This data is not inaccurate in the sense that the data mismeasures reality but rather they fail to account for the systemic biases that gave rise to the differences in the data. Algorithms blindly using this data without understanding the context around particular statistical disparities can reinforce and replicate patterns of discrimination in the decisions they make.[10]

Predictive Policing Algorithm Directs Police to Already Overpoliced Communities

The criminal justice system in the United States is often unjust. Black males receive sentences that are on average 19.1% longer than White males that commit similar crimes, even when controlling for past violence in their criminal history.[11] In Oakland, California researchers found that Black neighborhoods have 200 times more drug arrests than other Oakland neighborhoods—even though drug use rates are generally the same across racial lines.[12] A policing algorithm only trained on arrest data might then infer that Black neighborhoods use more drugs even though the arrest data says more about the over-policing of Black communities than actual rates of criminality and drug use.

Consider PredPol, a predictive policing algorithm intended to forecast the location of future crimes. When researchers applied the algorithm to Oakland, they found that the algorithm targeted Black neighborhoods at twice the rate of White ones. This is because the algorithm relied on Oakland's racially biased arrest data to make its predictions. When the researchers compared PredPol data to police data, they found that the neighborhoods PredPol identified as crime hotspots were the same neighborhoods that were already disproportionately targeted by the Oakland police for drug arrests. This example shows us how algorithms can reinforce existing patterns of over-policing and bias and the need to disentangle the effect of systemic discrimination from the data that powers decision-making algorithms.

Predictive Policing in Oakland vs. Actual Drug Use[13]

The chart on the left shows the demographic breakdown of people targeted for policing based on a simulation of PredPol in Oakland. The chart on the right shows actual estimated use of illicit drugs.

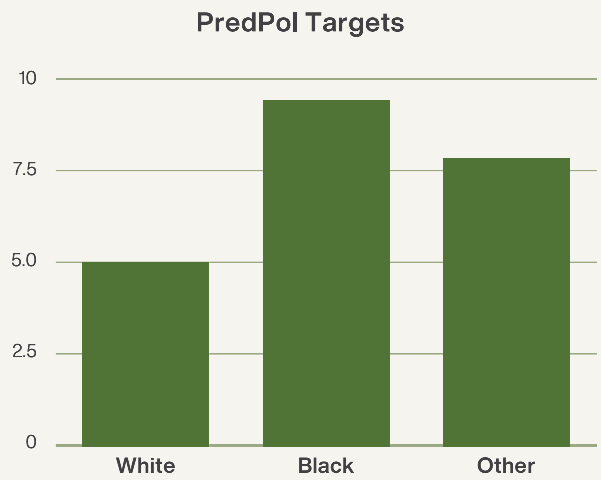

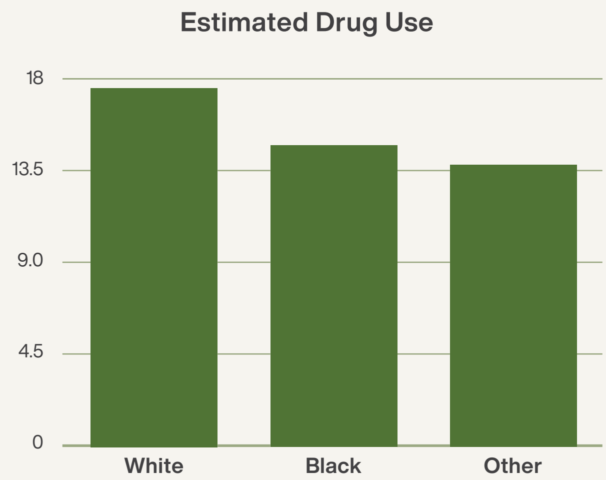

Choosing the Training Data

Advanced algorithms learn how to make decisions by learning from "training data," data that often reflects the decisions humans made in the past. Decisions based on incorrect, partial or non-representative training data have the potential to create unfair outcomes because the algorithm will draw the wrong conclusions from the data, resulting in incorrect predictions or decisions that are less accurate for groups that are underrepresented or misrepresented in the training dataset.

Non-Representative Training Data Can Lead to Inaccurate Results

CheXNet is an algorithm designed to predict whether a patient has pneumonia and other lung diseases based on their chest X-Rays. The algorithm learned how to diagnose the disease by analyzing over 100,000 chest X-Ray images that had diagnoses or "labels" attached to them by practicing radiologists.

Researchers examining the CheXNet algorithm and training dataset found that it performed significantly worse in diagnosing chest diseases in women when it was trained on a dataset where women were underrepresented and only made up 25% of the training images.[14] When the algorithm was trained on gender balanced data, this disparity disappeared.

The CheXNet Medical Diagnosis Algorithm [15]

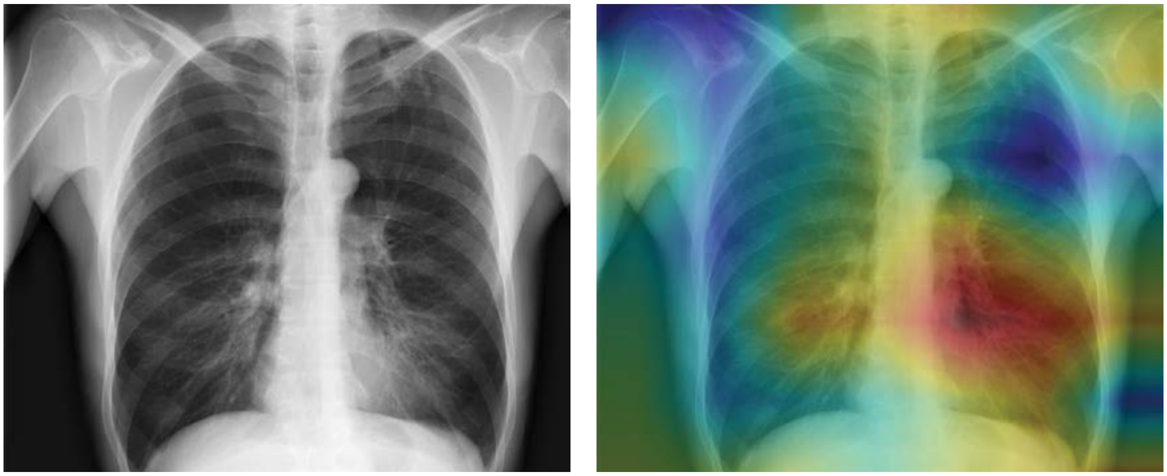

Input
Chest X-Ray Image

Output
Pneumonia Positive (85%)

CheXNet is an advanced algorithm that takes an X-ray image as an input, and outputs a probability of a particular disease or condition. In this example, CheXNet, correctly detects pneumonia and indicates areas in the body that are most likely affected by the disease.

Algorithmic Bias in Health Care

Algorithmic Bias in Health Care

The health care sector uses algorithms to process data from billions of medical records, insurance claims, clinical drug trials and other sources, including data from health devices like FitBit or the Apple Watch. The algorithms use this data to personalize treatment plans or guide the development of drugs and new therapies. Health care algorithms can lower health care costs and help doctors make more accurate decisions and diagnoses.

The Problem

Health care algorithms, if built improperly or used incorrectly, can lead hospitals to give substandard care to people of color or to make decisions that work for wealthier patients while harming lower-income ones.

Health Care Algorithms Can Discriminate Against Patients That Spend Less on Care

Bias can arise even if unbiased algorithms are applied inappropriately in health care. One example of that is the how one hospital system misused a health care analytic tool that gives health plans and health systems access to patient data and over 1,700 rules and algorithms that characterize a patient's care history and unmet needs. This tool is used to help identify which patients are likely to need individualized care plans and additional medical attention. Researchers found that one hospital system applied the tool in a way that would have recommended White patients for these resources while effectively denying those same resources for equally sick Black patients.[16] Through this use of the tool, Black patients effectively had to be sicker than White patients in order to be identified as patients who might benefit from additional care.

Where Things Went Wrong:

One of the 1,700 rules and algorithms in the analytics tool was specifically designed to predict future costs for patients – this model is not biased as a predictor of future cost. However, bias arose because the health system only used the cost prediction algorithm to characterize patient's health needs. The cost prediction model was not designed for this purpose. Cost is not a meaningful proxy for health. Future cost can be used as an "ingredient" to help identify certain populations who might benefit from additional support, but it must be used in the context of patients' other medical history and unmet needs. Unfortunately, the health system only used future cost to identify people for additional medical attention. In doing so, they identified relatively more White patients because White patients typically utilize more health care services than Black patients – which makes future predicted costs for White patients higher than Black patients. A key contributor to this disparity is due to a racial wealth gap that has given rise to structural differences in income and access to health care.

The developer of the algorithm noted that this use of the tool was completely inconsistent with the design of the tool and the training they offered to the hospital.This example shows us how properly defining outcomes and what to predict is a key consideration in building and applying fair and unbiased algorithms. The potential discriminatory impacts that could result from an incorrect use of an algorithm also reveals that proper training is necessary to ensure that these powerful decision-making tools are used properly and not in ways that cause unintended consequences.

Algorithmic Bias in Employment

Algorithmic Bias in Employment

Technology continues to transform the work environment as algorithms replace the traditional hiring and management processes. Companies small and large, such as Goldman Sachs, Amazon and staffing agencies, are using algorithms to recruit, promote and hire employees.

The Problem

Biased hiring algorithms replicate existing workplace dynamics that favor particular groups of employees over others.

Amazon's Recruiting Algorithm Learned to Favor Men over Women

In 2014, Amazon attempted to develop a recruiting algorithm to rate the resumes of job candidates and predict which applicants would do well at Amazon. The algorithm learned to favor men while penalizing women. Resumes that included the word "women's" or included the names of all-women's colleges were downgraded while resumes containing masculine language associated with male engineers were favored. Researchers scrapped the experiment because they were unsuccessful in fixing the issue.[18]

Where Things Went Wrong

This bias likely arose because of female underrepresentation in Amazon's training data. The algorithm scored applicants by observing patterns in resumes submitted to the company over a 10-year period. However, the tech industry is predominantly male and Amazon's dataset and hiring practices likely reflected that male dominance. So even though Amazon did not use gender as a variable in its model, the algorithm picked up on gender-related characteristics in the dataset. Amazon's hiring algorithm provides a clear example of how non-representative datasets can skew decisions in ways that harm underrepresented groups and how structural bias in a dataset (i.e. the lack of women in STEM fields and in tech company leadership) can affect decision-making algorithms.

Algorithmic Bias in Government Programs

Algorithmic Bias in Government Programs

Federal, state and local governments increasingly turn to algorithms to cut costs and streamline bureaucratic processes. Government algorithms are pervasive; they're used to guide municipal investment, assess the risk of releasing someone on parole, assess fraud or determine eligibility for social programs. They've also become an integral part of "smart cities" like Austin, New York and Atlanta, where they identify infrastructure problems and control traffic, street lighting and other services.

The Problem

Poorly designed government algorithms make bad decisions that can have serious consequences for struggling residents and communities that need access to government services and investments.

Michigan's Fraud Detection Algorithm Wrongly Accused 40,000 People of Unemployment Fraud

Between 2013 and 2015, the state of Michigan implemented an algorithm called the Michigan Data Automated System that accused more than 40,000 people of fraudulently claiming unemployment benefits.[19] Many accused of fraud were forced to pay heavy fines, declared bankruptcy or had their houses foreclosed upon. Many victims attempting to appeal the claims found out that there was no human oversight over the system and that fraud claims were based on incorrect or mismatched data. After several lawsuits and an audit of the MiDAS system, auditors found that MiDAS fraud charges were only affirmed 8% of the time on appeal.[20]

Arkansas Medicaid Algorithm Wrongly Denies Medical Care and Benefits

When Arkansas implemented a Medicaid access algorithm, hundreds of people saw their benefits cut—losing access to home care, nursing visits and medical treatments. Arkansas Legal Aid filed a federal lawsuit in 2016, arguing that the state failed to notify those affected, and that there was also no way to effectively challenge the system, as those denied benefits couldn't understand what information factored into the algorithm's decisions. The process for appealing these decisions was described as "effectively worthless" as less than 5% of appeals were successful. During the court case, the company that created the algorithm found multiple errors due to miscoding and incorrect calculations. An estimated 19% of Medicaid beneficiaries in the state were negatively affected by one error alone.[21]

Where Things Went Wrong

It is difficult to call Arkansas' Medicaid algorithm and Michigan's unemployment algorithm biased, since there is no evidence that their decisions were more likely to be wrong for men compared to women or for Black residents compared to White ones. Rather, these algorithms were poorly designed and simply bad at making decisions. However, the impact of these bad predictions fell disproportionately on residents who lost their jobs and needed access to health care. The difficulties victims faced in trying to appeal these algorithmic decisions highlight the need for technological due process, or the right to have a meaningful opportunity to understand and contest algorithmic decisions.[22]

Algorithmic Bias in Education

Algorithmic Bias in Education

Schools use algorithms to grade student's essays or to check for plagiarism. Colleges use algorithms to identify at-risk students or to determine the likelihood of a student accepting an admission offer.

The Problem

Algorithms that predict student achievement can punish students at low-performing schools and reinforce educational disparities.

Grading Algorithms in the UK Gives Lower Grades to Lower-Income Students

In 2020, the COVID-19 pandemic forced schools in England to cancel final exams nationwide, making it difficult to give out final grades and determine college placements. As a result, England's Office of Qualifications and Examinations Regulation (Ofqual) turned to an algorithm to calculate student grades. To calculate grades, the algorithm relied on teachers' prediction of that student's final grades, their academic performance and, critically, a school's historical performance data. The algorithm lowered 40% of teacher-provided grades in calculating the final results. An analysis of the algorithm found that the algorithm was more likely to lower grades for lower-income students and those who did not attend smaller private schools.[23] After a large public outcry, Ofqual scrapped the algorithmic grades and students received their teacher-assigned grades.

Where Things Went Wrong

The Ofqual algorithm is another example of a mismatch between the outcome an algorithm is supposed to predict and what it actually predicts. Ofqual's algorithm did not really determine a student's actual achievement throughout the year but rather predicted how well students in a particular school "should" do. The algorithm's focus on historical school performance as a predictor meant high achieving students in poorly performing schools were more likely to have their grades lowered. Ofqual's grading algorithm also raises questions about the ethics of assigning grades to students based on their school quality rather than more personal measures of achievement. In addition, the algorithm gave greater weight to teacher grades in schools with small class sizes, giving students at private schools an unfair leg up.

Algorithmic Bias in Credit and Finance

Algorithmic Bias in Credit and Finance

Banks and the fintech industry have eagerly replaced loan officers with algorithms that are more complex and use more sources of data than ever before to make more precise, and profitable, decisions about creditworthiness and loan applications. Other financial algorithms detect fraudulent credit card transactions or help companies comply with changing regulations and lending requirements.

The Problem

Access to credit is a key element of the American dream, enabling families to start small businesses and purchase homes, important vehicles for wealth and economic opportunity. However, biased banking algorithms can make it more expensive for Black and Brown borrowers to access credit and loans compared to similarly situated White borrowers.

Advertising Algorithms Allow Lenders to Target Vulnerable Consumers with Predatory Products[24]

Facebook, Google and third party data brokers use data collection algorithms to profile customers, placing them into buckets known as customer segments based on personal characteristics, behaviors and preferences. Advertisers use these segments to target ads to customers likely to buy their products. A Senate report warned that these tools could also allow lenders to target financially vulnerable populations with subprime loans that have predatory interest rates. In response to these concerns, Facebook and Google banned predatory lending ads on their platforms in 2015 but many lenders have found ways to evade these bans.[25]

Sample List of Targeting Products Identifying Financially Vulnerable Populations [26]		
Credit Reliant	Burdened by Debt: Singles	Tough Start: Young Single Parents
Rocky Road	Mid-Life Strugglers: Families	Living on Loans: Young Urban Single Parents
Very Elderly	Credit Crunched: City Families	Rough Retirement: Small Town and Rural Seniors
Rolling the Dice	Relying on Aid: Retired Singles	Ethnic Second-City Strugglers
Fragile Families	Struggling Elders: Singles	Enduring Hardships
Very Spartan	Retiring on Empty: Singles	Humble Beginnings
X-tra Needy	Small Town Shallow Pockets	Financial Challenges
Hard Times	Rural and Barely Making It	Meager Metro Means

This is a table listing several customer "segments" or categories that advertisers can use to specifically target vulnerable customers with ads. For example, customers in the "Hard Times" segment were described as "older, down-scale and ethnically-diverse singles typically concentrated in inner-city apartments. . . the bottom of the socioeconomic ladder, the poorest lifestyle segment in the nation . . . this is an underclass of the working poor and destitute seniors without family support."

Mortgage Lending Algorithms Continue to Overcharge Black and Brown Borrowers

Online banking algorithms emerged as a way to combat racial discrimination present in traditional, face-to-face lending.[27] Despite those claims, a UC Berkeley study showed that both traditional and online lenders overcharge Black and Brown borrowers for mortgage loans to the tune of $765 million a year compared to equally qualified White borrowers.[28] The study found that online algorithmic lenders charged similarly situated Black and Brown borrowers more than White ones, but the upshot is that they overcharged these borrowers less than traditional lenders.[29]

Where Things Went Wrong

The targeted marketing and lending algorithms described in this section illustrate the potential and problems of having so much personal and demographic data. Targeted marketing algorithms can connect vulnerable customers to the resources they need, but unscrupulous lenders can also abuse those tools and trap them in a cycle of poverty. Lending algorithms discriminate less than traditional lenders but still give White borrowers better rates and loans than Black ones. UC Berkeley researchers suggest that this bias is due to geographic and behavioral pricing strategies that charge more in financial deserts or if a customer is unlikely to shop around at competing lenders.[30] This raises serious questions about the fairness and legality of using data unrelated to credit repayment risk, such as shopping behavior, to make decisions about loan terms and rates.

Algorithmic Bias in Housing and Development

Algorithmic Bias in Housing and Development

Landlords and property managers are turning to algorithms to screen out bad renters while homeowners use algorithms to help buy and sell their homes. On a larger scale, cities and homebuilders use planning algorithms to guide development and determine which neighborhoods receive investment.

The Problem

Algorithms not only limit access to home loans for communities of color (see above), cities can misuse algorithms in ways that perpetuate housing disinvestment and redlining in low-income communities. Others screen potential renters in ways that unfairly disadvantage people of color.

Rental Screening Algorithms Can Deny Housing to Over-Policed Communities

Over 90% of landlords now use tenant screening algorithms to review and score potential renters.[31] One such algorithm called CrimSAFE was developed by Corelogic, a property and data analytics company. The algorithm prevented a mother from moving in with her disabled 14-year-old son because he had a shoplifting arrest from 2014 that was later dismissed. Corelogic was recently sued due to accusations that its algorithm disproportionately and illegally screened out Black and Latino applicants based on their criminal records.[32]

Detroit's "Market Value Analysis" Algorithm Targets Poorest, Blackest Neighborhoods for Disinvestment

Over 25 cities use a tool called the Market Value Analysis algorithm to classify and map neighborhoods by market strength and investment value. Cities use MVA maps to craft tailored urban development plans for each type of neighborhood. These plans determine which neighborhoods receive housing subsidies, tax breaks, upgraded transit or greater code enforcement. Cities using the MVA are encouraged by its developer to prioritize investments and public subsidies first in stronger markets before investing in weaker, distressed areas as a way to maximize the return on investment for public development dollars.[33]

In Detroit, city officials used the MVA to justify the reduction and disconnection of water and sewage utilities, as well as the withholding of federal, state and local redevelopment dollars in Detroit's "weak markets," which happened to be its Blackest and poorest neighborhoods.[34] The recommendations from Indianapolis' MVA meant small business support, home repair and rehabilitation, homebuyer assistance and foreclosure prevention programs were unavailable to the most distressed and isolated neighborhoods in the city.[35]

Where Things Went Wrong

The MVA uses variables like average home prices, vacancy rates, foreclosures and homeownership to determine neighborhood "value," while CrimSAFE uses criminal records to determine who is a good tenant. But these are not ahistorical, objective predictors of market value or criminality; rather they are data points that reflect systemic bias. Redlining accounts for 30% of the gap in homeownership and 40% of the gap in home values for Black Americans between 1950 and 1980.[36] Similarly Black and Brown Americans are much more likely to be racially profiled, stopped and targeted by police for arrest.[37] The data used in these algorithms therefore takes the "longstanding and deeply embedded link between race, property devaluation and anti-black violence" and essentially repackages it in terms of the risks that certain people and communities pose to landlords and public investment.[38] In doing so, these algorithms risk reinforcing patterns of disinvestment, redlining and housing discrimination.

Algorithmic Bias in Everything Else:

Price Optimization Algorithms

Algorithmic Bias in Everything Else: Price Optimization Algorithms

These days, it's common for two different people to pay different prices for the same exact product. Companies like Amazon, Expedia, Steam, Hilton, Safeway and countless more use algorithms to set prices on the fly. While dynamic pricing existed before algorithms, modern pricing algorithms have access to much more data about consumers. These algorithms can adjust prices by crunching information such as inventory levels, your location, income, emotional state or even your phone's battery level.

The Problem

While pricing algorithms can grow sales and increase profits, they can also result in price gouging and unfair price discrimination.

Princeton Review's Pricing Algorithm Charged Asian-Americans More for Online Tutoring

The Princeton Review, an educational company, used a pricing algorithm that charged between $6,600 and $8,400 for its Premier SAT course that offered online tutoring. A ProPublica analysis found that ZIP codes with a high median income or a large Asian population were more likely to be quoted the highest prices. Customers in predominantly Asian neighborhoods were nearly twice as likely to see higher prices compared to the general population, even if it was a low-income area.[39]

Insurance Algorithms Charge POC Neighborhoods More for Car Insurance than White Ones, Irrespective of Accident Risks

A study looking at insurance rates and payouts in California, Illinois, Texas and Missouri found that drivers in neighborhoods of color paid higher auto insurance rates than White ones even when the accident risk in those neighborhoods was nearly the same.[40] Overall, predominantly Black ZIP codes pay premiums that are 60% higher than predominantly White ones despite having similar risk profiles.[41] Research on Allstate's price-adjustment algorithm in Maryland found that it was optimized to increase profits by charging "big spenders" more for car insurance.[42]

Ride-Hailing Algorithms Charge Riders More for Traveling to Majority POC Neighborhoods

Thanks to a Chicago law requiring ride-hailing companies to disclose their fares, researchers discovered that Uber and Lyft's pricing algorithm charges riders more if they are dropped off in neighborhoods with a high non-White population.[43] The study found evidence that riders were charged more based on attributes such as age, education, home price and ethnicity.

Where Things Went Wrong

Algorithmic pricing models focus on maximizing profit, which they do well, but they are not designed to account for the equity and fairness implications of those decisions. Companies using algorithms to engage in dynamic and personalized pricing are able to maximize profit by studying our data, estimating the maximum we're willing to pay and charging that price. While these practices are generally legal, the power imbalance between customers and sellers means that these practices can be unfair and exploitative. Without the proper safeguards to limit pricing algorithms, they can continue to exploit and overcharge communities of color and produce a broad array of unfair effects.

Price Discrimination: Is it Illegal?

Surprisingly, price discrimination on the basis of race or other protected classes isn't illegal under federal price discrimination law.[44] Even if race-based price discrimination was illegal under state or federal law, these pricing algorithms generally don't change prices on your race or gender but rather based on your willingness to pay, how much competition is in your area, or even your emotional state. Uber, for example, discovered that riders were more willing to pay surge fares when their phone battery was low.[45] However, price discrimination can also be a good thing, like senior pricing or hardship discounts. This suggests that we need to update our price discrimination laws to limit personalized pricing practices that are overly manipulative and to prevent algorithmic price gouging.

Recommendations for Fixing Algorithmic Bias

Recommendations for Fixing Algorithmic Bias

There is a growing body of research outlining the solutions we need to end algorithmic discrimination and build more equitable automated decision systems. This section will focus on three types of solutions as a starting point: (1) algorithmic transparency and accountability, (2) race-aware algorithms and (3) algorithmic greenlining.

Algorithmic Transparency and Accountability

Algorithmic transparency enables people who use, regulate and are affected by automated decision systems to understand how the algorithm is designed, how it is used to make decisions, and the factors behind that decision. Without transparency it is difficult to achieve algorithmic accountability.

Algorithmic accountability refers to "the assignment of responsibility for how an algorithm is created and its impact on society; if harm occurs, accountable systems include a mechanism for redress."[46] In other words, algorithmic accountability laws allow us to identify and fix algorithmic harms and to enforce our existing laws against anti-discrimination. Algorithmic transparency and accountability measures include algorithmic impact assessments, data audits to test for bias, and critically, a set of laws that penalize algorithmic bias, particularly in essential areas like housing, employment and credit.

Algorithmic Impact Assessments and Data Audits

Governments and the community members they serve lack the necessary information to assess how algorithms are working and how they affect individuals, families and communities. Requiring companies and public agencies to conduct "algorithmic impact assessments" can help solve this problem. An impact assessment would require public agencies and other algorithmic operators to evaluate their automated decision systems and their impacts on fairness, justice, bias and other community concerns, and to consult with affected communities.

Data audits involve having third parties examine algorithms and the underlying data for bias and to see if the algorithm is transparent, fair and its decisions are explainable. Data audits help developers ensure their algorithm works and that it complies with applicable laws and ethical norms.

Audits and impact assessments can help stop companies from taking shortcuts in developing and implementing their algorithms and ensure that they confront the implications of their tool before it goes live. These practices can also build trust and public confidence that the algorithm is fair and works as intended.

Key Barriers

- Disagreement around the scope of algorithmic transparency requirements.[47]
- Public agencies and private companies are not required by law to conduct algorithmic impact assessments or data audits.
- Regulators and government agencies may lack the technical expertise and resources to meaningfully audit and review algorithmic systems.
- Community members and non-government experts are often left out of the accountability process. Algorithmic assessment must be interdisciplinary to account for gaps in individual knowledge and perspectives.[48]

Recommendations

- Enact state, local and federal laws that require detailed algorithmic impact assessments and data audits for automated decision systems that involve employment, education, insurance, banking, public agencies, healthcare and housing. Build systems for community input and engagement.
- Increase budgets for technical capacity-building and hiring within our government institutions so that regulators, attorneys general and other enforcement agencies have the ability to meaningfully investigate, document and audit algorithms for bias and discrimination. Effective transparency and accountability requires effective data regulators who are empowered to conduct algorithmic bias assessments, have the technical capacity to analyze and audit decision-making algorithms and can penalize companies for unfair and illegal practices.

Race-Conscious Algorithms

Many companies choose to avoid using protected characteristics like race and gender in their algorithms to avoid liability for discrimination. However, this veil of ignorance around sensitive attributes, like gender or race, isn't all that helpful because algorithms can still figure out who you are based on related data like your internet history or they can discriminate using proxy variables like ZIP code or educational background. Moreover, if we start trying to eliminate proxy variables, then the accuracy of algorithmic decisions begins to suffer—so this is not an effective solution either. In fact, research shows that allowing algorithms to access data on race and other characteristics can actually help prevent discrimination and allow for more accurate decisions.[49]

While this may appear counterintuitive, awareness of race or other protected characteristics allows algorithms and their developers to correct for situations where the predictive power of a particular variable is not the same across different races, ethnicities or genders.[50]

Race-Conscious Educational Algorithms Can Improve Fairness and Equity in Decision-Making

Predictors of educational success like extracurricular activities, high school grades and standardized test scores are less predictive of college success and GPA for Black students than they are for White ones. An experiment revealed that a race-aware college admissions algorithm could admit twice as many Black students without lowering the predicted GPA of the class when compared to a race-blind algorithm. Allowing algorithms to adjust their scoring and evaluation criteria to account for these types of differences can yield more accurate, equitable results that mitigate risks of algorithmic redlining.[51]

Key Barriers

- Legal limitations and the threat of lawsuits discourage companies from explicitly using data on race and other sensitive attributes to make algorithmic decisions.[52]
- Resistance to collecting and using race data from an ethical and moral perspective.
- Disagreement on whether to limit race-awareness to the training stage of the algorithm or to allow its use in both the training and decision-making contexts.

Recommendation

- Require companies to document and measure algorithmic outcomes across protected classes. Build consensus through research, regulatory rulemakings and community engagement around the contexts and conditions where algorithms can have access to sensitive characteristics and develop legal protections, or safe harbors for those situations.

Algorithmic Greenlining

> "Big Data processes codify the past. They do not invent the future. Doing that requires moral imagination, and that's something only humans can provide. We have to explicitly embed better values into our algorithms, creating Big Data models that follow our ethical lead. Sometimes that will mean putting fairness ahead of profit."
>
> —Cathy O'Neil, author, "Weapons of Math Destruction"

The final strategy we discuss to prevent algorithmic bias is to focus on algorithmic equity. Instead of a defensive strategy aimed at limiting discrimination and preventing disparate impacts, algorithmic greenlining emphasizes using automated decision systems in ways that promote equity and help close the racial wealth gap. In this vision, algorithms are not only fair, but go beyond simply not causing harm to address systemic barriers to economic opportunity. Optimizing for equity requires a shift away from designing algorithms that maximize profit and return on investment and towards algorithms that optimize for shared prosperity, racial justice and resilient ecosystems.[53]

Optimizing for Equity: California's Climate Justice Algorithm

In 2012, the Greenlining Institute successfully advocated for SB 535 (De Leon, 2012), a bill directing 25% of the dollars raised by California's Greenhouse Gas Reduction Fund to communities that are the most burdened by pollution exposure and poor socioeconomic outcomes.[54] Rather than manually identifying these communities on a case by case basis, the law directed California to develop a scientific tool to classify and decide which communities in California are "disadvantaged" under the statute and would therefore qualify for priority funding and investment.

After a robust community engagement process, the California EPA developed "CalEnviroScreen," an algorithm that identifies disadvantaged communities across 8,000 census tracts. CalEnviroScreen works by examining multiple indicators such as exposure to pollution and unemployment rates. After processing this data, the algorithm outputs a CalEnviroScreen score that quantifies the environmental and socioeconomic burdens within a community and determines whether it is eligible for targeted investments.[55]

CalEnviroScreen Equity Indicators

Pollution Burden	Population Characteristics
Exposures · Ozone · PM2.5 · Traffic · Toxic Release from Facilities · Pesitcide Use · Diesel PM · Drinking Water Contaminants **Environmental Effects** · Solid Waste Sites and Facilities · Cleanup Sites · Groundwater Threats · Impaired Water Bodies · Hazardous Waste Generators and Facilities	**Sensitive Populations** · Asthma · Cardiovascular Disease · Low Birth-Weight Infants **Socioeconomic Factors** · Poverty · Educational Attainment · Unemployment · Linguistic Isolation · Housing Burdened Low Income Households

CalEnviroScreen

Pollution Burden × **Population Characteristics** = **CalEnviroScreen Score**

Pollution Burden: Average of Exposures and Environmental Effects

Population Characteristics: Average of Sensitive Populations and Socioeconomic Factors

While this algorithm is relatively simple, it demonstrates how algorithms can help make data-driven decisions and automate complex and time consuming tasks. Since implementation of SB 535 and the CalEnviroScreen algorithm, over $3 billion in state funding has gone towards investments in disadvantaged communities, with a larger share of that funding going to low-income Black and Latino communities.[56] Beyond this, CalEnviroScreen provided invaluable and irrefutable data to residents and policymakers showing that exposure to pollution and poverty is most pronounced in Black and Brown communities. In this way, CalEnviroScreen demonstrates the potential of algorithms to drive equity outcomes and improve economic opportunities for communities of color.

Key Barriers

- Extractive, profit-seeking behavior by private companies and institutions.
- Legal prohibitions on consideration of race and other protected attributes in decision-making.
- Difficulty defining and measuring equity outcomes.
- Limited community engagement and technical capacity.

Recommendation

- Develop laws that integrate equity metrics into public decision-making algorithms, particularly for those that govern access to community investment and resource allocation. Engage and fund equity mapping and community-centered data gathering efforts to track the cumulative impacts of exposure to health disparities, economic inequality and pollution as a way to inform equitable policy implementation.[57]

Conclusion

Algorithms and automated decisions are powerful, pervasive and, as this report shows, often unfair, inaccurate and discriminatory. Even tech giants like Facebook, Microsoft and Google have joined privacy and consumer advocates to ask Congress and state legislatures to establish new rules and regulations for algorithms and A.I.[58] This push for legislative action presents an opportunity to not only develop policies that minimize unfair algorithmic discrimination but also to create a system where decision-makers optimize algorithms for equity and inclusion, and design them in ways that drive investments to the most vulnerable communities and use them to build a better and more equal society.

Acknowledgements

The Greenlining Institute

Founded in 1993, The Greenlining Institute envisions a nation where communities of color thrive and race is never a barrier to economic opportunity. Because people of color will be the majority of our population by 2044, America will prosper only if communities of color prosper. Greenlining advances economic opportunity and empowerment for people of color through advocacy, community and coalition building, research and leadership development. We work on a variety of major policy issues because economic opportunity doesn't operate in a vacuum. Rather than seeing these issues as being in separate silos, Greenlining views them as interconnected threads in a web of opportunity.

Technology Equity

The Technology Equity team's mission is to ensure equal access to technology, eliminate algorithmic bias and discrimination while encouraging the use of data-driven solutions that build economic opportunity in communities of color.

Authors

Gissela Moya is Greenlining's 2019-20 Technology Equity Fellow and works with the Tech Equity team on issues including broadband access, privacy and algorithmic equity. She holds a B.A. in political science from UC Berkeley and strongly believes we must deliberately invest in communities of color to ensure an equitable future. She is the co-author of Greenlining's recent report, On the Wrong Side of the Digital Divide: Life Without Internet Access, and Why We Must Fix it in the Age of COVID-19. Her most recent work regarding broadband access and the impacts of the 2020 digital census on communities of color has been featured in the San Francisco Chronicle, The Progressive, Forbes and other media outlets.

Vinhcent Le is a Technology Equity Attorney at The Greenlining Institute, where he develops Greenlining's strategy to protect consumer privacy, prevent algorithmic bias in economic opportunity and to close the digital divide. In this role, Vinhcent helped secure multi-million-dollar commitments to increase broadband access in California, modernization of the Lifeline program and the development of a program providing laptops to low-income students across the state.

Editorial

Bruce Mirken, Media Relations Director at The Greenlining Institute

Design

Ashley Johnson, Digital Strategy Manager at The Greenlining Institute

Nadia Kim, Communication Coordinator at The Greenlining Institute

Endnotes

1 Onuoha, M., & Nucera, D. (2018). A People's Guide to A.I., https://alliedmedia.org/wp-content/uploads/2020/09/peoples-guide-ai.pdf, 4.

2 Kleinberg, J., Ludwig, J., Mullainathan, S., & Sunstein, C. R. (2018). Discrimination in the Age of Algorithms. Journal of Legal Analysis, 10, 113-174. doi:10.1093/jla/laz001,132-33.

3 Source: The Algorithmic Justice League and Megan Smith (former Chief Technology officer of the USA)

4 Kleinberg, J., Ludwig, J., Mullainathan, S., & Sunstein, C. R. (2018). Discrimination in the Age of Algorithms. Journal of Legal Analysis, 10, 113-174. doi:10.1093/jla/laz001, 121-122.

5 Protected classes include sex, race, age, disability, color, creed, national origin, religion and genetic information.

6 Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. Calif. L. Rev., 104, 671, 675, Kroll, J. A., Barocas, S., Felten, E. W., Reidenberg, J. R., Robinson, D. G., & Yu, H. (2016). Accountable algorithms. U. Pa. L. Rev., 165, 633.

7 Kleinberg, J., Ludwig, J., Mullainathan, S., & Sunstein, C. R. (2018). Discrimination in the Age of Algorithms. Journal of Legal Analysis, 10, 113-174. doi:10.1093/jla/laz001.

8 Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. Calif. L. Rev., 104, 671, 677; Kroll, J. A., Barocas, S., Felten, E. W., Reidenberg, J. R., Robinson, D. G., & Yu, H. (2016). Accountable algorithms. U. Pa. L. Rev., 165, 633, 128-132.

9 Kleinberg, J., Ludwig, J., Mullainathan, S., & Sunstein, C. R. (2018). Discrimination in the Age of Algorithms. Journal of Legal Analysis, 10, 113-174. doi:10.1093/jla/laz001.

10 Barocas, S., & Selbst, A. D. (2016). Big data's disparate impact. Calif. L. Rev., 104, 671, 691-692.

11 United States Sentencing Commission. (2018). Demographic Differences in Sentencing: An Update to the 2012 Booker Report. https://www.ussc.gov/sites/default/files/pdf/research-and-publications/research-publications/2017/20171114_Demographics.pdf, 2

12 Lum, K., & Isaac, W. (2016). To predict and serve?. Significance, 13(5), 14-19. https://doi.org/10.1111/j.1740-9713.2016.00960.x; Smith, J. (2016, October 09). Crime-prediction tool PredPol amplifies racially biased policing, study shows, Mic. https://www.mic.com/articles/156286/crime-prediction-tool-pred-pol-only-amplifies-racially-biased-policing-study-shows.

13 source: National Survey on Drug Use and Health, Human Rights Data Analysis Group and Mic

14 Larrazabal, A. J., Nieto, N., Peterson, V., Milone, D. H., & Ferrante, E. (2020). Gender imbalance in medical imaging datasets produces biased classifiers for computer-aided diagnosis. Proceedings of the National Academy of Sciences, 117(23), 12592-12594. https://www.pnas.org/content/117/23/12592.

15 Rajpurkar, P., Irvin, J., Zhu, K., Yang, B., Mehta, H., Duan, T., ... & Lungren, M. P. (2017). Chexnet: Radiologist-level pneumonia detection on chest x-rays with deep learning. arXiv preprint arXiv:1711.05225.

16 Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. Science, 366(6464), 447-453. https://sendhil.org/wp-content/uploads/2020/01/Publication-67.pdf.

17 Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. Science, 366(6464), 447-453. https://sendhil.org/wp-content/uploads/2020/01/Publication-67.pdf.

18 Dastin, J. (2018, October 10). Amazon Scraps Secret AI Recruiting Tool That Showed Bias Against Women. Reuters. https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G

19 Wykstra, S., Motluk, A., & Kaplan, A. (2020, May 29). Government's Use of Algorithm Serves Up False Fraud Charges. Undark. https://undark.org/2020/06/01/michigan-unemployment-fraud-algorithm/; Garza, A. (2020, May 28). Automated Systems Trapping Citizens in Bureaucratic Limbo. Time. https://time.com/5840609/algorithm-unemployment/.

20 Ringler, D. (2016). Office of the Auditor General Performance Audit Report: Michigan Integrated Data Automated System (MiDAS). State of Michigan Auditor General, https://audgen.michigan.gov/finalpdfs/15_16/r641059315.pdf, 28.

21 Lecher, C. (2018, March 21). What happens when an algorithm cuts your health care. The Verge. https://www.theverge.com/2018/3/21/17144260/healthcare-medicaid-algorithm-arkansas-cerebral-palsy.

22 Citron, D. K. (2007). Technological due process. Wash. UL Rev., 85, 1249, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1012360.

23 Bedingfield, W. (2020, August 21). Everything that went wrong with the botched A-Levels algorithm. Wired. https://www.wired.co.uk/article/alevel-exam-algorithm; Duncan, P., McIntyre, N., Storer, R., & Levett, C. (2020, August 13). Who won and who lost: When A-levels meet the algorithm. The Guardian. https://www.theguardian.com/education/2020/aug/13/who-won-and-who-lost-when-a-levels-meet-the-algorithm; Lee, M. W., & Walter, M. (2020). Equality Impact Assessment: Literature Review. Ofqual, https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/879605/Equality_impact_assessment_literature_review_15_April_2020.pdf.

24 United States Senate Committee on Commerce, Science, and Transportation (2013), A Review of the Data Broker Industry: Collection, Use, and Sale of Consumer Data for Marketing Purposes, http://www.commerce.senate.gov/public/?a=Files.Serve&File_id=0d2b3642-6221-4888-a631-08f2f255b577.

25 Jones, C., Eaglesham, J., & Andriotis, A. (2020, June 03). How Payday Lenders Target Consumers Hurt by Coronavirus. The Wall Street Journal. https://www.wsj.com/articles/how-payday-lenders-target-consumers-hurt-by-coronavirus-11591176601.

26 United States Senate Committee on Commerce, Science, and Transportation (2013), A Review of the Data Broker Industry: Collection, Use, and Sale of Consumer Data for Marketing Purposes, http://www.commerce.senate.gov/public/?a=Files.Serve&File_id=0d2b3642-6221-4888-a631-08f2f255b577.

27 Miller, J. (2020, September 18). Is an Algorithm Less Racist Than a Loan Officer? The New York Times. https://www.nytimes.com/2020/09/18/business/digital-mortgages.html; Klein, A. (2020), Reducing bias in AI-based financial services. Brookings. https://www.brookings.edu/research/reducing-bias-in-ai-based-financial-services/

28 Bartlett, R., Morse, A., Stanton, R., & Wallace, N. (2019). Consumer-lending discrimination in the FinTech era (No. w25943). National Bureau of Economic Research. https://faculty.haas.berkeley.edu/morse/research/papers/discrim.pdf;

29 Bartlett, R., Morse, A., Stanton, R., & Wallace, N. (2019). Consumer-lending discrimination in the Fintech era (No. w25943). National Bureau of Economic Research..29; But see Bhutta, N., & Hizmo, A. (2020). Do minorities pay more for mortgages? https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3352876, arguing that interest rate differentials by race and ethnicity can be explained by other pricing options and rebates.

30 Bartlett, R., Morse, A., Stanton, R., & Wallace, N. (2019). Consumer-lending discrimination in the Fintech era (No. w25943). National Bureau of Economic Research, https://faculty.haas.berkeley.edu/morse/research/papers/discrim.pdf;

31 Kirchner, L. (2020, September 24). Can Algorithms Violate Fair Housing Laws? The Markup. https://themarkup.org/locked-out/2020/09/24/fair-housing-laws-algorithms-tenant-screenings.

32 CONNECTICUT FAIR HOUSING CENTER v. CORELOGIC RENTAL PROPERTY SOLUTIONS, LLC, No. 3: 18-CV-705 (VLB) (D. Conn. Aug. 7, 2020)., See also Houston, S. (2018, August 24). Center Files Federal Lawsuit Against National Tenant Screening Company. https://www.ctfairhousing.org/corelogic/; Lecher, C. (2019, February 01). Automated background checks are deciding who's fit for a home. https://www.theverge.com/2019/2/1/18205174/automation-background-check-criminal-records-corelogic

33 Safransky, S. (2020). Geographies of algorithmic violence: Redlining the smart city. International Journal of Urban and Regional Research, 44(2), 200-218. https://doi.org/10.1111/1468-2427.12833.

34 Safransky, S. (2020). Geographies of algorithmic violence: Redlining the smart city. International Journal of Urban and Regional Research, 44(2), 200-218. https://doi.org/10.1111/1468-2427.12833.

35 The MVA recommends people-based strategies like jobs training and social services in distressed areas while reserving place-making strategies for wealthier, more up-and-coming areas. Indianapolis' development plan appears to follow the MVA recommendations, as Type 02 "Address Underlying Issues" neighborhoods which combine high poverty and high percentages of Black residents are not targeted for place-making interventions. City of Indianapolis (2017), "Neighborhood Investment Strategy." Indianapolis Metropolitan Development Commission, 140-60,182-83. https://citybase-cms-prod.s3.amazonaws.com/c8f0ca24f2e54b1292b5d84b3e454575.pdf.

36 Aaronson, D., Hartley, D., & Mazumder, B. (2017). The effects of the 1930s HOLC" redlining" maps (No. 2017-12). Working Paper. http://hdl.handle.net/10419/20056 at 28-33.

37 Despite making up only 13% of the general population Black men and women account for 21-28% of people arrested in 2017. Sawyer, W., & Jones, A. (2019). Arrest, Release, Repeat: How Police and Jails are Misused to Respond to Social Problems. Prison Policy Initiative. https://www.prisonpolicy.org/reports/repeatarrests.html.

38 Safransky, S. (2020). Geographies of algorithmic violence: Redlining the smart city. International Journal of Urban and Regional Research, 44(2), 200-218. https://doi.org/10.1111/1468-2427.12833.

39 Angwin, J., Mattu, S., Larson, J., (2015, September 1). The Tiger Mom Tax: Asians Are Nearly Twice as Likely to Get a Higher Price from Princeton Review. Propublica. https://www.propublica.org/article/asians-nearly-twice-as-likely-to-get-higher-price-from-princeton-review.

40 Angwin, J., Mattu, S., Larson, J., & Kirchner, L. (2017, April 05). Minority Neighborhoods Pay Higher Car Insurance Premiums Than White Areas With the Same Risk. Propublica. https://www.propublica.org/article/minority-neighborhoods-higher-car-insurance-premiums-white-areas-same-risk.

41 Feltner, T., & Heller, D. (2015). High Price of Mandatory Auto Insurance in Predominantly African American Communities. 16. Washington, DC: Consumer Federation of America. https://consumerfed.org/wp-content/uploads/2015/11/151118_insuranceinpredominantlyafricanamericancommunities_CFA.pdf; Consumer Federation of America (2020, June 17). Systemic Racism in Auto Insurance Exists and Must Be Addressed By Insurance Commissioners and Lawmakers. https://consumerfed.org/press_release/systemic-racism-in-auto-insurance-exists-and-must-be-addressed-by-insurance-commissioners-and-lawmakers/.

42 Varner, M., & Sankin, A. (2020, February 25). Suckers List: How Allstate's Secret Auto Insurance Algorithm Squeezes Big Spenders, The Markup. https://themarkup.org/allstates-algorithm/2020/02/25/car-insurance-suckers-list.

43 Pandey, A., & Caliskan, A. (2020). Iterative Effect-Size Bias in Ridehailing: Measuring Social Bias in Dynamic Pricing of 100 Million Rides. arXiv preprint arXiv:2006.04599. https://arxiv.org/pdf/2006.04599.pdf

44 Edwards, M. A. (2006). Price and Prejudice: The Case Against Consumer Equality in the Information Age. Lewis & Clark L. Rev., 10, 559.

45 Chowdhry, A. (2016, May 26). Uber: Users Are More Likely To Pay Surge Pricing If Their Phone Battery Is Low. Forbes. https://www.forbes.com/sites/amitchowdhry/2016/05/25/uber-low-battery/?sh=3ce65eea74b3

46 Caplan, R., Donovan, J., Hanson, L., & Matthews, J. (2018). Algorithmic Accountability: A Primer, Data & Society. https://datasociety.net/wp-content/uploads/2019/09/DandS_Algorithmic_Accountability.pdf.

47 Kroll, J. A., Barocas, S., Felten, E. W., Reidenberg, J. R., Robinson, D. G., & Yu, H. (2016). Accountable Algorithms. U. Pa. L. Rev., 165, 633.

48 Richardson, R. (2019). Confronting Black Boxes: A Shadow Report of the New York City Automated Decision System Task Force, AI Now Institute. https:// ainowinstitute.org/ads-shadowreport-2019.html.

49 Kleinberg, J., Ludwig, J., Mullainathan, S., & Sunstein, C. R. (2018). Discrimination in the Age of Algorithms. Journal of Legal Analysis, 10, 113-174. doi:10.1093/jla/laz001.; Hardt, M., Price, E., & Srebro, N. (2016). Equality of Opportunity in Supervised Learning. Advances in Neural Information Processing Systems, 29, 3315-3323.

50 Kleinberg, J., Ludwig, J., Mullainathan, S., & Sunstein, C. R. (2018). Discrimination in the Age of Algorithms. Journal of Legal Analysis, 10, 113-174. doi:10.1093/jla/laz001.

51 Kleinberg, J., Ludwig, J., Mullainathan, S., & Sunstein, C. R. (2018). Discrimination in the Age of Algorithms. Journal of Legal Analysis, 10, 113-174. doi:10.1093/jla/laz001.

52 The Supreme Court reaffirmed a commitment to race-neutral remedies to racial differences in a 2014 decision stating: "Remedial orders in disparate-impact cases should concentrate on the elimination of the offending practice that 'arbitrarily. . . operates invidiously to discriminate on the basis of race.' If additional measures are adopted, courts should strive to design them to eliminate racial disparities through race-neutral means." Texas Department of Housing and Community Affairs v. Inclusive Communities Project, Inc., 135 S. Ct. 2507 (2015).

53 Cooper, S., Sanchez, A. (2020). Greenlined Economy Guidebook, The Greenlining Institute. https://greenlining.org/wp-content/uploads/2020/09/Greenlined-Economy-Guidebook-2020.pdf

54 California Global Warming Solutions Act of 2006: Greenhouse Gas Reduction Fund, S.B. 535, 2012 Reg. Sess. (Cal. 2012). https://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill_id=201120120SB535

55 Rodriquez, M., & Zeise, L. (2017). CalEnviroScreen 3.0, California Environmental Protection Agency and Office of Environmental Health Hazard Assessment. https://oehha.ca.gov/media/downloads/calenviroscreen/report/ces3report.pdf

56 California Environmental Protection Agency (2018). Analysis of Race/Ethnicity, Age, and CalEnviroScreen 3.0 Scores. https://oehha.ca.gov/media/downloads/calenviroscreen/document-calenviroscreen/raceageces3analysis.pdf;

57 Evergreen Collaborative (2020). Designing a New National Equity Mapping Program: Identifying Communities that Face Environmental Injustice, Using Lessons Learned from State Equity Mapping Programs. https://collaborative.evergreenaction.com/policy-hub/EquityMapping.pdf.

58 Herrera, S. (2020, January 27). Tech Giants' New Appeal to Governments: Please Regulate Us. The Wall Street Journal. https://www.wsj.com/articles/tech-giants-new-appeal-to-governments-please-regulate-us-11580126502.

Photography

Cover:
Markus Spiske

Page 3:
Branko Stancevic

Page 20:
whoislimos

Page 25:
Campaign Creators

Page 32:
seventyfourimages