THE NETGAIN PARTNERSHIP

AUTOMATION *& the* QUANTIFIED SOCIETY

A report from

Upturn

By David G. Robinson and Miranda Bogen

The primary research and drafting of this report took place over the spring and summer of 2017.

Cover photo: Alex Berger / Flickr

About NetGain

In February 2015, a group of leading public interest funders made a public commitment to jointly address the challenges and opportunities of the digital age. This new NetGain partnership was supported by five founding partners: the Ford Foundation, John D. and Catherine T. MacArthur Foundation, John S. and James L. Knight Foundation, Open Society Foundations, and Mozilla Foundation. This year, for the first time, a sixth funder—the Omidyar Network—has joined with the original five in an expanded coalition.

NetGain deepens the commitment to digital rights from the five founding partners, each of which has long worked on Internet and technological issues, spending approximately $65 million collectively each year. By collaborating, the funders aim to make their investments more efficient and better coordinated, allowing the group to address the Internet's most pressing challenges—those that are too large for any one organization to tackle alone.

Each year, NetGain members select and work together on a broad, shared theme. In 2015, the theme was the "pipeline" of tech-related talent into public interest roles, in government and the social sector. In 2016, the coalition focused on the Internet of Things. This year, the coalition's work focuses on automated decision-making and the quantified society, the subject of the framing paper you are reading now.

NetGain is guided by six core principles. Member philanthropies commit to:

- Work to make the Internet an open, secure, and equitable space for free expression, economic opportunity, knowledge exchange, and civic engagement that everyone can access and afford.
- Support the opportunities created by the networked public sphere: new modes of civic and social participation, lowered barriers to engagement, and innovative ways to organize for positive change and social justice; and guard against potential harm: censorship, self-segregation, the spread of misinformation, and polarization.
- Transform learning and ensure that young people have the skills they need to succeed in a connected, complex world: digital literacy, critical thinking, problem solving, curiosity, empathy, understanding, persistence, and more.
- Cultivate leaders in business, government, and civil society to understand and fulfill the promise of the Internet, and support cross-sector alliances to ensure technology and data are used to advance the public good.
- Contribute to the design of Internet and information technology policies, practices, and products that enhance data security and protection of individual privacy.
- Ensure that philanthropy leads in digital security and data ethics in its own practices.

About Upturn

Upturn is a team of technology and policy experts based in Washington DC, working to ensure that technology serves the dignity and well-being of all people.

Acknowledgements

This report would not have been possible without the time and support of the many interviewees whose input and reflection guided the authors' analysis. Participants in a working session at the Data & Society Research Institute, in New York, provided extremely helpful feedback on an earlier draft of this report. Additionally, program staff from several foundations, including the Ford Foundation, the John D. and Catherine T. MacArthur Foundation, the John S. and James L. Knight Foundation, the Open Society Foundations, the Mozilla Foundation, and Omidyar Network provided input that helped to shape and guide this research.

CONTENTS

EXECUTIVE SUMMARY

THROUGHOUT THE WORLD, POWERFUL INSTITUTIONS are increasingly using computers to automate decisions that matter to people's lives. Artificial intelligence and other advances in software development, pervasive data collection, and vanishingly cheap computing power are working together to drive considerable change, creating what some have called the "quantified society."

NetGain funders and the social sector at large are deeply committed to protecting civil and human rights, defending the vulnerable, and building a more humane and equitable world. Meanwhile, technological change is sweeping the landscape, creating both new opportunities and new needs. The current moment demands careful thought and sustained attention, informed by both deeply held values and technical insight.

This framing paper is designed to spark and structure an urgently needed conversation about what automation in a quantified society means for core human values, across the full spectrum of issues that motivate NetGain funders and their allies—and how we can work to amplify its benefits while mitigating its harms. This report is concerned with decisions made by software, that shape the lives of vulnerable people and groups. It reflects extensive research, interviews with civil society stakeholders across the world, and our own sustained work over years of consulting projects, coalition advocacy, and scholarship and teaching on these issues.

Artificial intelligence is just one technique for automating decisions. Other, older software tools are also being deployed in new ways that touch vulnerable groups.

We begin by briefly describing what's new in computing that is driving so much social, economic, and political change. Reading those few pages, you will learn what "machine learning" is all about. This technique, the dominant form of artificial intelligence in use today, harnesses computing power to find patterns in historical data—patterns that can then be used as the basis for predictions and decisions. At the same time, AI is just one technique for automating decisions, and other, older software tools are also being deployed in new ways that touch vulnerable groups.

In the second section, we share a few key insights about *how* social and political systems tend to change at times like this, when computers take on new responsibilities. The benefits of these new technologies are substantial and widespread. At the same time, research on human-computer interaction bears out what many people feel intuitively: Giving computers more power inside a complex, important institution doesn't just make things faster or cheaper—the institution itself, and the people inside it, often redefine their goals and change the way they act. When people place too much trust in computers, they risk deferring to automated judgments that may be wrong or ill-in-

formed. This trust, coupled with problematic training data, algorithms based on incorrect assumptions, and embedded social bias, can often leave minority populations as what engineers call "edge cases"—people and situations the computer does not handle well.

Finally, the heart of our report offers five core themes at the intersection of social change, automation, and the quantified society. With each theme, we explain what is happening and why, anchoring a high-level discussion with real-world examples. We then describe trends to watch for, which are important pathways of benefit and risk that we see inside each area, and potential approaches to addressing these developments. We do not propose or assess specific investments, but instead suggest directions that may warrant further exploration. Each theme is here because it brings a cluster of related developments into focus.

Core Themes: Automation and the Quantified Society

1. Corporate Power, Information, and the Attention Economy: A handful of major companies, most based in the U.S., operate globally dominant Internet platforms. These firms are constantly linked to their users via smartphones and other technology. They have direct, intimate, and immediate knowledge of each of their users, and have a historically unprecedented role in informing, influencing, and motivating the behavior of billions of people. Voting, civic mobilization, and intellectual exchange, around the world, are among the many facets of life these firms increasingly enable, mediate, and shape through increasingly automated processes. Beyond the direct impact of platforms themselves, philanthropy by the major platforms and their founders plays a growing role in the social sector, where it could shape public debates by highlighting the best impacts of AI and big data while downplaying structural risks. Understanding and responding to the impact of major online platforms is a crucial challenge for the social sector in the years ahead.
2. Patterns, Discrimination, and Justice: When automated decisions are based on historical data, they risk entrenching unjust social patterns and projecting such patterns into the future. Much of the data that exists today about crime, health, commerce and other vital domains was gathered to suit the needs of well-resourced organizations, sometimes to the detriment of the marginalized. From a computer's perspective, minorities are often "edge cases," exceptions that a system may not be well designed to handle. Systems that reinforce existing patterns often work against the goals of human rights and civil rights advocacy, insofar as advocates seek to *change* longstanding social patterns. Emergent efforts aim to develop more inclusive data, systematically monitor bias and risk, and rectify discriminatory patterns.
3. Jobs, Work, and Meaning in an Era of Automated Decisions: Automated hiring, workplace monitoring, employee assessment, and supply chain and logistics systems are shifting the landscape of non-discrimination law, workplace protection, and labor organizing. Other forms of automation threaten to knock out or redefine entire categories of work. Whether technology displaces workers or simply changes the nature of their

work, this shift risks deepening the economic inequality and social instability that is already emerging across the world. The specifics of changing labor markets will be different in developed and developing economies, in countries with and without strong social safety nets, and in geographies with relatively robust or relatively weak worker protections—but across the board, civil society will need to be attentive to shifting currents to determine appropriate interventions.

4. Automated Decisions in the Public Sector: Governments around the world increasingly use automation to make important decisions about people's lives, often without broad public consultation or careful assessment of new systems' impact. When attempting to adapt these technologies for public ends, governments have struggled to access needed expertise and to navigate normative and legal concerns related to equal treatment, privacy, and other ethical challenges. To leverage new technologies in ways that prioritize the public good, technical expertise, and knowledge is vital. But in their efforts to catch up with private sector innovation, governments often rely on exogenous sources of expertise—such as partnerships with leading technology platforms—that may interpose different values into the exercise of public authority. It remains uncertain to what extent the progress of digital technology will be reflected in public services—and where the technology does arrive, how it may reshape how those services operate—for example, by changing the way that judges perceive and interact with defendants in the courtroom.

5. Freedom, Transparency, and State Power: Automated profiling and predictive analytics make mass surveillance both less expensive and more powerful. That puts more people at risk of being swept up in government scrutiny, targeted by preemptive enforcement, subjected to state prejudice, or manipulated by disinformation. State and state-supported exploitation of automated tools to manipulate public opinion is developing rapidly, while efforts to map and constrain such activities lag behind. Prompt and coordinated philanthropic activity may be able to shift this dynamic, including by building or strengthening walls between government and corporate surveillance. Approaches in this area align with broader privacy and digital rights goals.

Automation and the quantified society are transforming areas of vital concern to NetGain and its allies. We hope this paper provides a clear frame and a useful starting point for the high-impact discussions that lie ahead. ■

INTRODUCTION

THROUGHOUT THE WORLD, POWERFUL INSTITUTIONS are increasingly using computers to automate decisions that shape people's lives. New techniques for developing software without human instruction, pervasive data collection, and vanishingly cheap data storage and processing power are contributing to this trend.

Civil society is experimenting with different labels that highlight facets of these related developments: automated decisions, algorithms, artificial intelligence (AI), the quantified society, big data, and data at scale, among others. These terms point to fast-evolving structural challenges that face NetGain funders and their peers—not only in their tech-related grants, but also in their need to anticipate and shape the broader landscape within which all of their efforts unfold.

"Machine learning" is one part of the story, where computers are programmed to detect patterns in the data they are given. Artificial intelligence (AI) is an umbrella term for machines taking over cognitive tasks once done by humans, and machine learning is the primary engine behind AI today. The rapid changes machine learning brings—coupled with longer-standing challenges posed by other decision-making software—are remaking the landscape that philanthropies work to improve. And the scale and variety of personal data stored and analyzed by major institutions is an inescapable part of the same narrative.

What do these developments mean for philanthropy? How will they impact social equities and human rights for those on the margins of society? How will NetGain's member philanthropies, and ultimately the broader community of funders, grantees, and other stakeholders working toward its common goals, need to adjust their thoughts and plans in response to these trends?

This framing paper is designed to spark, organize, and motivate a shared exploration of those questions.

Our thinking here reflects an extensive literature review regarding automation, the quantified society, and civil and human rights, including perspectives drawn from current affairs and public policy, law, computer science, and anthropology. We conducted semi-structured interviews with a global cohort of civil society and academic stakeholders to gather their perspectives on these developments. Additionally, the questions explored in this paper are ones we've grappled with extensively over the last several years, through a blend of consulting and research projects, coalition advocacy efforts, and writing and teaching.

The rapid changes machine learning brings—coupled with longer-standing challenges posed by other decision-making software—are remaking the landscape that philanthropies work to improve.

Finally, the text you are reading has itself benefited from feedback by a diverse range of reviewers. At the same time, we are a small, U.S.-based team, mindful of our own inherent limits. We fully expect that other perspectives, beyond the ones described here—and other frames, beyond the ones proposed in this paper—will also play vitally important roles as philanthropy grapples with the transformative impact of automated decisions in a variety of social and geographic contexts. We hope this paper will stimulate conversation and reflection within and beyond the NetGain community. ■

II UNDERSTANDING TODAY'S AUTOMATION

THIS REPORT IS CONCERNED WITH LIFE-ALTERING decisions, made by software, that impact vulnerable people and groups. Computers have long played at least some role in the decision-making processes of most large and powerful institutions, and in many contexts, that role is rapidly growing. As the extent of automation in human and civil rights contexts increases, the challenges and opportunities for promoting and protecting civil and human rights will often change apace.

Software's role in decisions impacting human and civil rights is widely varied, as the many examples throughout this report illustrate. It ranges from simply matching names in a database, to awarding public benefits, to granting visas, to weighing dozens or hundreds of factors to determine which job candidates are best suited for an interview or which city blocks are most in need of a police patrol. Advanced computer logic also enables the automation of activity in the physical world, in contexts like manufacturing, shipping, driving, and robotics.

Algorithms are basic building blocks in mathematics and computer science. Systems that rely on algorithms don't always replace human judgment.

It is useful to think of automation as a matter of degree.[1] An "algorithm," in the public imagination, is something mystical, mysterious, and inscrutable. In fact, however, algorithms are simply a sequence of steps used to accomplish some task. Algorithms are basic building blocks in mathematics and computer science. Systems that rely on algorithms don't always replace human judgment, and they aren't always hard to understand. On the other hand, some systems—after their initial design and ignition—do operate entirely automatically without human intervention, and some are so complex that even experts struggle to understand how they operate.

Remarkable Benefits and Rapid Change

Our report naturally focuses on unsolved problems, rather than on the many areas in which automation and the quantified society improve the daily lives of people on the margins of society, and of people generally—improvements that are substantial, often transformative, and that will continue without philanthropic intervention. A vital part of the background for this conversation is to appreciate the many ways in which the rapidly advancing technology of digitally automated decisions brings benefits to everyone.

These benefits include real-time automated translation among many languages; cloud-based services that automate rote office tasks and reduce the capital costs of launching a business; recommendation engines that enhance people's experience of music, video, and other media; and innumerable "enterprise" innovations inside large

businesses and supply chains, introducing efficiencies and driving down the costs of consumer goods and services.

A Newly Powerful Technology: Machine Learning

One driver of growing automation is an approach called machine learning, in which the computer detects patterns or correlations in existing data (often called its "training data") to use as models in anticipating future outcomes. This is a different approach than following pre-defined instructions from a human engineer. Machine learning systems are, as Britain's Royal Society put it, "computers that learn by example."[2]

Machine learning is advancing by leaps and bounds, making computers far better able than ever before to undertake a range of formerly human responsibilities, from identifying the faces in surveillance footage to making subtle, consequential predictions about how people will behave. Indeed, machine learning is the dominant approach to artificial intelligence, the delegation to machines of choices formerly requiring human capabilities and effort.

The simplest machine learning models, such as those used in credit scoring, use a few well-understood input factors that are closely related to the outcome being predicted. Their outputs can be simple, linear combinations of a small number of factors.

On the other hand, "deep learning" systems find correlations among potentially vast numbers of input variables, and go on to predict outcomes or make choices whose causes cannot be readily intuited by a human observer. Because the systems learn these patterns independent of human guidance, it can be difficult to identify problematic logic, let alone to correct it. In response to this development, some argue that systems must be "explainable" and "accountable"—but technical, regulatory, and advocacy communities and other interested parties lack consensus on precisely what form those concepts ought to take, from both a technical and regulatory perspective.[3] Academic researchers are actively working on techniques to better understand how these systems reach their conclusions, but as new machine learning methods are introduced across a wider range of contexts, the lag in our ability to understand and tweak the behavior of these complex digital decision systems is understandably concerning.

Although machine learning is widespread, it is also important to note that many critical decisions made by software, that impact civil and human rights, do not presently rely on advanced machine learning methods. These decisions depend heavily on traditionally programmed rules—where a software engineer explicitly specifies the rules a computer will follow—and on the humans who receive and interpret the decisions. ■

Machine Learning in Practice: Job Applications

The first uses of software to screen job applicants were relatively basic. For instance, recruiters would identify keywords tied to the skills and competencies necessary for a job, and a software program would flag applicants whose resumé or cover letter included those words.

New machine learning systems for hiring use a different approach. Rather than relying on human recruiters to identify which words in a resumé make a candidate seem promising, a machine learning system can analyze the job applications that were submitted by past and current employees, and compare each employee's job application to his or her subsequent performance. (The employer might define performance by productivity, scores in performance reviews, or something else that the employer measures and cares about.) The machine "learns" by identifying features in job applications that statistically correlate with high job performance. Such features might include the type of education a candidate has, or the number of years' experience he or she has in a similar job. The data that gets analyzed for patterns—in this case, the job applications and subsequent performance data—is called *training data*. The algorithm that finds these patterns is a *learner*.

Using patterns in the training data, the learner builds a *model* that relates job applicants' traits to job performance. New job applications can then be input into the model, and the model can output predictions, usually as a numerical score: how likely is this applicant to succeed, if hired? The software can rank candidates according to those predictions, and hiring managers can interview those with the highest scores.

Machine learning can help hiring managers more quickly find qualified candidates, especially those who traditional hiring software might overlook. Sometimes, the model contains surprises: for instance, some employers have discovered that college degrees are statistically less important to job performance than they thought. But such systems can also find correlations that may be inappropriate to use. For example, if opportunities have been denied to women in the past, then features that are more typical of male job applicants might emerge as key statistical predictors of success. Human review of a learned model, informed by both engineering knowledge and domain expertise, remains vital in practice.

III CONCEPTUAL TOOLS FOR THE NEW LANDSCAPE

THIS SECTION PROVIDES IMPORTANT BACKGROUND for consideration of the civil and human rights impacts of automation and the quantified society, describing the common threads of opportunity and challenge that arise across many areas of philanthropic and public concern in this context.

The same technologies that turn the smartphone into a quasi-magical device, with its real-time traffic and transit updates, instant translation among dozens of languages, and ability to recognize and recommend nearly any recorded music or book, are also helping doctors and technicians diagnose disease, recruiters find and hire job candidates who lack conventional credentials, humanitarian aid workers track population movement following natural disasters, and small businesses predict and meet market demands. Further gains, such as self-driving cars that empower seniors to remain mobile, are on the horizon. For these and other reasons, many people around the world believe strongly that technological progress is mostly good.

Automated decisions offer significantly greater efficiency and consistency, and reduced cost, relative to the human methods that they replace.

The key driver of the major changes described in this report is that automated decisions offer significantly greater efficiency and consistency, and reduced cost, relative to the human methods that they replace, in a startlingly wide and growing range of contexts. Market forces are driving rapid adoption by business. Sometimes, though, speedy adoption of these new tools can happen without careful thought about potential drawbacks, and ways to address those drawbacks.[4]

Automation Can Trigger Cascading, Unanticipated Change

When computers are used to automate a given task, it can be tempting to assume that the institutions and processes surrounding that task will remain the same—that automation will make things faster, but not otherwise different. This assumption—sometimes called the "substitution myth"—is often a mistake.

Instead, it is reasonable to expect and helpful to predict that automation will change the nature of the institutions, relationships, and activities into which it is introduced—for better or worse. For example:

- **Automated CV screening can redefine the hiring process.** In the domain of employment, software that can assess CVs initially seemed likely to help human recruiters sort

through the huge piles of job applications that many listings receive. But, increasingly, these same technologies now make it cheap for employers to proactively identify promising candidates for a job—in effect, mass, automated headhunting. Such tools may, in some instances, obviate the problem they were initially intended to solve, as a strategy of individualized outreach replaces the massive pile of applications that the recruiting software first promised to help sort through.

- **Electronic medical records can lead doctors to "treat the chart."** The introduction of automation "into medicine . . . is influencing the way [physicians] learn, the way they make decisions, and even their bedside manner."[5] Many participants in the healthcare debate argued that huge savings could be achieved by digitizing patient medical records, which would automate the laborious paperwork and billing activities that consume so much attention in medical organizations. But automation of patient recordkeeping did not generally yield the promised savings. Instead, in many instances the auto-generated documentation gave hospitals a chance to bill for innumerable small charges that formerly weren't worth tracking. Meanwhile, patient charts came to be laden with verbose boilerplate, making them harder to read and less useful for clinicians.[6]

In short, automation and quantification change the things they touch.[7] The argument that James C. Scott makes in *Seeing Like a State*—that bureaucracies redefine the world around them in order to make it "legible"—applies with similar force to computer software. People come to be represented as data, and that data becomes a simulacrum for who they are. Especially for minority populations, whose situations often differ from the statistical or cultural "norm" of the population as a whole, this translation into numbers often leaves them as "edge cases" that the model may not correctly predict.

People Often Trust Computers Too Much

Automated systems deliver enormous benefits across a huge range of areas, among them the promise of correcting flawed or biased human decision-making. At the same time, relying heavily on automated systems may lead people to repeat certain predictable mistakes. Such errors, often called "automation bias," are explored in depth in the academic literature on human factors engineering.[8]

In short, people often trust software too much, and end up deferring to computers even when the machines are mistaken. Commercial aviation today, for example, is safer than ever, but the incidents that still do occur are often a result of pilots' excessive trust and reliance on their aircraft's autopilot and other automated systems.[9] Human skills and judgment often atrophy as software is introduced. And the people who design and build software seldom have a full appreciation for the gamut of real-world challenges that can and will arise in the setting where the computer is being added. As Nicholas Carr has written, automation bias "creeps in when people give undue weight to the information coming through their monitors."[10] (Note this bias—in favor of machines— is different from social biases that can be reflected in automated decisions, which we discuss in Section 4.)

This problem is particularly acute where computers are used to make import-

ant decisions. Because the systems are usually reliable—and the analog alternatives so much less so—they can inspire "learned carelessness" in the humans who operate them.[11] For example, whenever machine learning is trained on patterns from current or historical data in order to make future predictions, the system can—and likely will—reproduce whatever biases the training data reflects. But the predictions, emerging from a computer, may still naturally seem trustworthy and lead people to accept their validity despite flawed assumptions.

The best outcomes will likely come from solutions that recognize the vulnerabilities inherent in both human and automated behavior.

Governance Regimes May Need to Be Updated

Most of the governance processes that are in place today to protect civil and human rights were developed before the current wave of automation began. Many will need to be updated to match new realities.

When computers replace humans, an activity that was formerly inscrutable—inside of someone's head—is transformed into something explicit, a system of inputs and outputs whose workings might, at least in principle, be easier to understand, monitor, and hold accountable. For example, an algorithm that assesses a consumer's credit risk will have concrete factors underlying that decision, such as payment history and amount of money owed, which can be shared with the affected individual or with regulators. And perhaps more easily than human thoughts, software code and selection of underlying data is controlled, and can be altered, through explicit regulatory choices.

Many of today's civil and human rights regimes long predate the current wave of automation.

But while human legal and political systems have centuries of practice understanding and regulating human decision-making, relatively little experience is available to guide our approach to regulating decisions that emerge from computer systems—particularly those powered by machine learning.

This is not an empty field or a brand new challenge, but rather a vibrant space with a fast-evolving, if not yet mature, range of governance approaches[12]:

- Governance strategies designed for rigidly mechanical bureaucracy may apply naturally to software-based automation. These include mandates both for disclosure of information relevant to a decision, and descriptions of the process by which a decision was reached. For example, in India, the Right to Information (RTI) regime requires government authorities to "provide reasons for [their] administrative or quasi-judicial decisions to affected persons."[13] There is an open question about whether or not these rules apply to technologically automated, as well as traditional policy-driven, decisions.
- Some purpose-built legal mandates—some longstanding, and others new—apply directly to automated choices. The U.S. Fair Credit Reporting Act, enacted in 1970 amidst a wave of concern about the human impact of databases, requires in some cases

that consumers be given "adverse action notices" that explain why they were denied a job or loan. The EU's new General Data Protection Regulation (GDPR) includes a "right to explanation" for significant decisions reached on a purely automated basis, though that right has existed in some form in previous data protection laws without its intended effect, and its true reach and meaning under the stronger legal framework of the GDPR remain unclear.[14]

- In certain contexts, law and policy may hold the person or entity that builds or operates an automated system accountable for that system's autonomous choices, an approach that Madeline Elish has described as treating human beings as "moral crumple zones,"[15] and which may effectively preserve some existing legal protections in the context of newly automated systems. On the other hand, debates are ongoing about how the law might inversely hold autonomous systems themselves responsible for their actions, whether machine liability would actually lead to meaningful remedy, and what "punishment" would deter future problematic behavior.

Debates over how best to govern automation, using tools both old and new, are far from resolved, and will continue to evolve in response to the emergent forces we describe below.

Human Rights Frameworks Could Be a Useful Lens

In recent years, the digital rights community has focused heavily on the two areas of privacy and freedom of expression, both enshrined in many domestic laws as well as in the Universal Declaration of Human Rights. As attention has turned to automation and all its related concepts, civil society groups have used these lenses to express concern with the growing influence of algorithms and automation on people's lives. Automated decisions challenge privacy because they rely on the widespread collection and processing of personal data, advocates have argued, while automated ranking systems and filtering algorithms that online platforms rely on to sort and moderate content challenge people's ability to express their opinions online. Somewhat separately, groups more traditionally focused on civil rights, particularly in the U.S., are beginning to engage on issues of fairness and discrimination in automated systems—values also included in human rights frameworks.

Nevertheless, human rights frameworks remain an underutilized lens through which to understand the impact of an increasingly quantified and automated society. Article 2 of the Universal Declaration of Human Rights enshrines non-discrimination as an inalienable right of all people, and Article 6 holds that everyone has the right to recognition before the law. Article 7 reiterates the value of equal treatment and protection, Article 21 provides that everyone have equal access to public services, and Article 23 reminds us that everyone has "the right to work, to free choice of employment, to just and favourable conditions of work and to protection against unemployment."[16] All of these rights are implicated by automated processes that, without care, risk invisibly enshrining inequity, quietly making decisions in the courtroom and in social service

agencies that are difficult to see or challenge, and subversively steering people away from dignified work—or automating away jobs altogether.

In the ongoing debates about automation and algorithmic decisions, moving beyond privacy, data protection-centric conversations and domestic legal frameworks to invoke a greater range of universal human rights that are facing new challenges in new contexts could provide a powerful and unifying framework for stakeholders concerned about these issues across the world.

The Private Sector Plays a Central Role

Private companies have long exerted influence on global affairs, but have taken on new, expanded roles as globalization and the Internet have left traditional public authorities with a reduced role in mediating public life or administering certain rights.[17] As a result, the allocation of power and influence is in flux:[18] Government and regulatory agencies continue to play a primary role in maintaining public spaces and protecting the vulnerable, but large technology companies are increasingly involved in formerly public, civic functions like building infrastructure,[19] validating and arbitrating identity,[20] and facilitating public discourse.[21] At the same time, these companies are racing ahead to develop even more advanced machine learning techniques in service of their products, outpacing the capacity of government to constrain them. As corporations become more powerful in relation to government, they become both harder to regulate and more valuable as potential allies to constrain unwanted actions by government. ■

Corporations are becoming harder to regulate, and may become more powerful allies in efforts to constrain unwanted actions by government.