IFOW

Institute for the Future of Work

Equality through transition

A discussion paper

Putting people first

"

We tend to overestimate the effect of technology in the short run and underestimate the effect in the long run.

Roy Charles Amara[1]

Contents

Authors
Dr Anne-Marie Imafidon MBE, Trustee of IFOW
Anna Thomas, Director of IFOW
Caitlin Doherty, IFOW Principal Researcher
Joshua Simons, IFOW Research Fellow

This paper is a team effort and we are grateful to a wide range of individuals and organisations for their input. Particular thanks goes to Helen Mountfield QC; Professor Michael A Osborne; Eric King; the Oxford Internet Institute; Dr Victoria Nash; Professor Jeremias Prassl; Professor Vili Ledonvirta; Aiha Nguyen and everyone who attended IFOW's scoping workshop on 5 July 2018.

Introduction

Equality is at the core of workplace rights, movements and democracy. But what does it really mean? Equality suggests the equal treatment, and respect, of people across different groups. Our aspiration for equality is built into the foundations of our legal and democratic systems and has found expression in different ways, such as the Equality Act and the right of each citizen to a vote. This aspiration has driven the most important constitutional moments in recent history. In the work space, where technology is driving a multi-dimensional transformation and stakes are high, the drive to promote meaningful equality is sharpening.

This discussion paper will identify key positive and negative implications of technology on equality in the work space, offer a new framework for thinking about impacts on different aspects of work, and initiate new activities for our 'Equality Through Transition' theme. A focus on the human experience of technology through work as it transforms is at the centre of our approach. We think this offers a grittier and more personal perspective on some of the most significant challenges facing modern industrial society. Our overarching goal is to make work fairer and better through the Fourth Industrial Revolution (4IR), with a focus on our most vulnerable communities. But the central role that work plays in people's lives, the way in which power dynamics in society are played out through work, and the fact that new technologies are often created in the work place, means that this approach should also contribute to the broader debate on Technology and Ethics.

Technology has the potential to be a real driver for inclusion, as well as growth. But to achieve this potential, we must think harder about what we mean by and expect from 'equality,' and tackle some thorny issues.

Technology impacts work place equality in a number of distinct, but related, ways. We use the term 'technology' broadly to include AI-related technology, robotics, big data analysis, the internet and internet of things, digital technologies and combining these technologies in diverse ways. The focus of this paper, however, is on technologies enabled by and developed with artificial intelligence (AI). This is for two reasons: first, AI is one of the most significant drivers of change in our transforming labour markets; second, it presents very specific challenges to our current frameworks for promoting equality. The tension between the search for a meaningful form of *legal* equality in a landscape of increasing *social and economic* inequality is felt most acutely at work. As we discuss below, this tension has been exposed by recent innovations in machine learning. In the work place, the need to understand and resolve the tension is acute. Technology has the potential to be a real driver for inclusion, as well as growth. But to achieve this potential, we must think harder about what we mean by and expect from 'equality' and tackle some thorny issues.

“

The debate on AI at work is both exciting – bordering on overhyped – and surprisingly narrow.

Anne-Marie Imafidon MBE

Our equalities framework

At each level, AI-related technology has the potential to either exacerbate or alleviate inequality. What happens will depend on the choices we make to develop, integrate and control technology at each level.

As an introductory framework to help us think about the implications of technology on equality, this paper identifies three levels at which technology – particularly technologies built on systems of AI and machine learning (ML)[2] – transform the experience of work. We have designed this framework to distinguish the impacts of technology on formal and structural equality and examine their connections. The time is right for this discussion. We are at a cross-roads: the way we embed technological systems through our institutions will determine whether or not technology is used to promote meaningful equality, reflected in positive experience.

In this paper, we start with the impact of AI on formal, or legal, equality. This is our keystone, as it is for democracy, although we think afresh about what it means, what we need, and whether legal mechanisms are achieving their underlying purpose. At each level, AI-related technology has the potential to either exacerbate or alleviate inequality. What happens will depend on the choices we make to develop, integrate and control technology at each level.

Technology and work: three levels of transformation

First, at the individual level: AI is increasingly used to make or support decisions, undertake functions and change processes in ways which are profoundly and immediately transforming the experience and conditions of work.[3] AI has the potential to open up access to the labour market, to support remote and flexible working, to target under-represented groups for jobs and to monitor diversity more closely.[4] But instances such as the high-profile collapse of the Amazon hiring tool, which down-graded candidate applications containing references to 'woman', demonstrate that unreflective use of ML can exacerbate entrenched structural inequalities.[5]

Second, at the level of firms: AI, the internet, large data sets and increasingly powerful processing, are combining to drive the creation of new business models, reshaping the relationships and power dynamics between workers and employers. This is disrupting the jobs we carry out, how we do business and the infrastructure within which work takes place.[6] Specific impacts on different types of equality at a firm level include: asymmetry of information, control by algorithmic management and the institutionalising of distinctions between groups by platforms.[7] We note that board-level decisions about technology use and adaption of business models are overwhelmingly made by non-representative groups.[8]

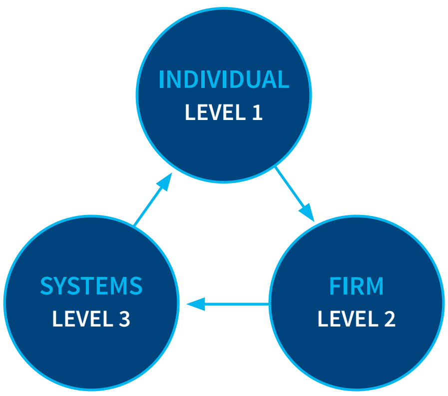

Third, at a systems level, the concentrated control and ownership of data-driven technologies amongst corporate 'giants' is contributing to wider shifts in the distribution of economic resources and power ('structural transformation').[9] Here, new technology has significant potential to increase our productivity, reduce costs and generate more wealth for distribution. But recent research from the International Monetary Fund (IMF)[10], the Organisation for Economic Co-Operation and Development (OECD)[11] and World Bank[12] suggests that, on our current trajectory, the rise of 'labour-light' digital platforms, combined with the internationalisation of big business, is driving market concentration and may be intensifying structural social and economic inequalities across regions. As things stand, the adverse effects of disruption are not spread evenly between groups of different gender, race or socio-economic disadvantage.

Technology at work

Machine learning

One area of AI stands out in terms of its capabilities and implications on equality at work: machine learning algorithms. ML is a form of AI which makes predictions and classifications based on the structures and patterns found within data sets. This may amount to a decision in itself, which an AI system acts upon, or it can inform a decision made by a human. At work, this can include decisions on such fundamental topics as hiring, firing, performance or pay. ML has the potential to address thorny problems, open access and achieve unprecedented accuracy on an unprecedented scale.

As with AI, ML is often presented in a way that suggests neutrality: an algorithm of independent capability with more processing power and less emotion than human actors. The reality is far messier. ML is a set of techniques designed by a human which addresses a problem defined by a human, and is trained on data-sets which usually encode the structures, opportunities and disadvantages of a very human landscape. Decisions made by, or with the help of ML, are 'socio-technical.' This is a particular challenge to our narrow way of looking at legal equality, and to legal equality itself. We expect AI to treat us as equals. But if the data sets involved in 'teaching' AI and ML systems how to operate are based on the outcomes of social inequalities, the systems themselves will replicate these outcomes.[13] Used without careful regard to impacts on equality at each level, ML often learns in a way that compounds disadvantage. This means that decisions or practice relevant to Level 1 may reflect the structural inequalities of Level 3, leading to multiple disadvantages experienced at an individual level.

Machine learning is a set of techniques designed by a human which addresses a problem defined by a human, and is trained on data-sets which usually encode the structures, opportunities and disadvantages of a very human landscape.

Technology at work *continued*

What do we know?

Our scoping evidence review and workshop hosted with the Oxford Internet Institute (OII) has revealed trends and gaps in our understanding to date of the impact of AI on equality, as both principle and practice, within the workplace. At the outset of this programme, our preliminary findings are that:

Use of data-centric AI-related tools is growing fast in the work space because of ease of access and computer processing. Small and Medium Enterprises' (SMEs) use of technology is now increasing, with indirect use of tools is possible via agents and platforms.[14] We know ML applications are increasingly common and pervasive in business, but we know little about the extent of use in the workplace by UK firms, which varies significantly across the EU. This means that we are dependent on individual case studies and reporting, coverage of which is dominated by industry products and research.[15]

Use of generic AI-related tools is uneven and tends to be associated with insecure and low-paid work, in particular workers classified as contractors undertaking platform work and sectors already vulnerable to disruption, including retail and transport.[16] By contrast, the overarching trend towards digitization has widened the gap within sectors and among companies between early adopters and others.[17]

The size and diversity of data-sets for aggregation to feed ML is increasing. At work, this means that data sets on social media use are frequently combined with data sets on web searches, leading to a consolidation of individuals' information without proper mechanisms for oversight or review. ML is powerful because of the quantity, breadth and granularity of data on which models are trained. So, although it is hard to quantify, the power of ML and those who own and control it, is shifting the power dynamics between employers and workers.[18] Academic research has started into aspects of these shifts and led to OII's welcome FairWork initiative.[19]

We know machine learning applications are increasingly common and pervasive in business, but we know little about the extent of use in the workplace by UK firms, which varies significantly across the EU.

A very small number of companies are shaping use and experience of ML in the work space with their combined technical expertise capable of designing ML tools, control over data sources and provision of internet-based services. This has led to a unique concentration of economic, political and cultural power which informs our mission to promote equality at work.[20] This asymmetry of information means that it is hard to achieve critical and informed positions on how ML affects human experience at work.

Academic analytical research by computer scientists is addressing the implications of ML for different mathematical characterisations of fairness, technical explanation and transparency. This fantastic body of work has focused on interpretability and technical challenges to individual rights, including worker and privacy rights, and has significantly contributed to our understanding of new challenges. In particular, progress has been made in understanding how AI can be rendered interpretable.[21] This research is relevant to our interests but is cross-domain and has not focused on work.

Technology at work *continued*

Our programme, which focuses on the issue of workplace equality during the transition aims to draw from the strengths from both the AI and work debates and fill the current gaps. A sharper focus on the implications of AI on equality at work is one of the best ways to make work better and fairer.

The debate on AI and Ethics has entered the mainstream. We have seen a string of initiatives world-wide, led mostly by technologists, focusing on the importance of future ethical conduct. In the past few years we have seen a proliferation of new codes[22] purporting to regulate the ethical conduct of AI use. Moreover, all four major AI companies are developing tools aimed at mitigating bias[23] and the study of ethics is being added to existing curriculums within STEM subjects. The debate in the UK is rooted by the welcome establishment of the Centre for Data Ethics and Innovation and the UK AI Council.[24] These bodies and developments are relevant to work in general terms, but they have not yet been applied to examining the various impacts on work with precision.

Insofar as current initiatives extend to work, they are constrained by a narrow focus on individual rights. We think that the three levels of transformation which we outline in this paper exist in a symbiotic state, where adjustments and disruptions at one level lead to consequences throughout. We must therefore take a more holistic approach to viewing the impact and effect of ML and AI systems on human experience.

Some key topics are almost entirely absent from current initiatives. We are seeing important academic and union work beginning on digital standards,[25] algorithmic management and control[26] and surveillance. But there is a marked absence of work on targeting and hiring tools beyond internal ad hoc corporate initiatives.[27] These are welcome, but facts reported are controlled by those responsible.[28]

There is very little robust research on discrimination and diversity at work and indeed through the AI cycle of innovation, notwithstanding some high-profile campaigning. Increasing concern has not yet led to robust research on impacts on gender, racial or socio-economic diversity, even limited to the level of individual rights.

Non-governmental bodies dedicated to human rights are not taking up the issue of 'automated' discrimination in practice. A recent LSE review of automated discrimination in data-driven systems through the lens of 28 civil society organisations active in the field of human rights and social justice across 9 EU countries noted a low uptake in the issue of automated discrimination.[29] A cross-disciplinary approach is needed but still unusual.

This analysis indicates some marked gaps in the AI and work zone. In particular: first, granular and intimate; second, broader and holistic look at the different impacts of AI on the individual experience of work is missing. Our programme, which focuses on the issue of workplace equality during the transition, aims to draw from the strengths from both the AI and work debates and fill the current gaps. A sharper focus on the implications of AI on equality at work is one of the best ways to make work better and fairer. Conversely, due to the contractual and historical power structures that are peculiar to work and the fact that work sits at the centre of individual lives, communities and the economy, we believe that work is a lens to view the multi-dimensional impacts of technology on society as a whole.

Three challenges

We have identified three key challenges, one for each level of transformation.

Key Challenge 1:
Discrimination and the rights of the individual

Machine learning may pose a challenge to our existing approach towards individual equality in the workplace. It also exposes its limitations. The root of the challenge lies in the combination of immense technical and computational power offered by ML, the socio-technical nature of ML systems, and the decisions at which they arrive.

The Equality Act 2010 is the central pillar that currently exists to protect against inequalities at work for the individual, and so is the obvious place to start to examine the impact of ML on equality at an individual level. It originates in and codifies legislation written well before ML as we know it: the Race Relations Act written in 1976, the Sex Discrimination Act 1975, Disability Discrimination Act 1995.

Direct discrimination is the most blatant form of unequal treatment and occurs when one person treats another less favourably 'on the grounds of' a protected characteristic such as race, gender, religion. Under the law as it stands, individuals are protected from discrimination based on one protected characteristic, not a combination of characteristics.

Indirect discrimination recognises that inequality may result from acts that are not overt or intentional. Instead, it may operate more subtly by way of practices or criteria that appear to be neutral but in fact operate so as to discriminate against one group when compared to the treatment of another.[30] Here, the employee must establish a causal relationship between the practice and the disadvantage, but not between the disadvantage and the protected characteristic (so the employee does not have to offer an explanation of 'why'). This type of discrimination allows the employer to explain or justify the practice for a reason that is unrelated to a characteristic. Here, the tribunal will carry out an objective balancing exercise, weighing up the discriminatory effect of the practice and the reasonable needs of the employer. Each element required to establish indirect discrimination is subject to a significant body of caselaw. We are not aware of any case law which, as yet, deals with ML.

We have seen that ML systems are not inherently objective. The goal of machine learning is to make predictions based on statistical pattern analysis: to classify and so treat individuals or different groups 'differently' because of characteristics, as well as behaviour, which have been found to be reliable predictors. This means that actions informed by the logic of absolute neutrality – our aspiration for formal equality – is not achievable through ML. The factors taken into account and complexity of methods by which ML systems combine different types of datasets to 'discriminate' between groups *may* not be caught by our existing set of protected characteristics. We know that training ML systems on existing datasets may replicate or compound existing inequalities, particularly at Levels 2 and 3, in ways that are difficult for humans to understand or trace. For example, data patterns found in use of social media, location or website history on a computer in a public library revealing socio-economic status (which is not protected) are taken into account in a way previously not possible.

The factors taken into account and complexity of methods by which machine learning systems combine different types of datasets to 'discriminate' between groups may not be caught by our existing set of protected characteristics.

Three challenges *continued*

Indirect discrimination recognises that inequality may result from acts that are not overt or intentional. Instead, it may operate more subtly by way of practices or criteria that appear to be neutral but in fact operate so as to discriminate against one group when compared to the treatment of another.

The most astute and relevant work on the implications of ML for the equal treatment of individuals has been by computer scientists, who have developed a range of 'computational' definitions of fairness.[31] These almost all aim at some sort of statistical relationship between groups. They include: anti-classification, which can require omitting protected attributes (and their proxies) from training sets; demographic parity, when predictive performance is equalised across groups that are defined by protected characteristics; and calibration across subgroups, which focuses on the outcome and retrospectively aims to make sure that the outcome does not depend on a protected characteristic.

As ML is used increasingly widely, and affects several fundamental components of work, it is becoming increasingly important to make sure that ML is not only mathematically fair but that decisions involving ML do not end up reinforcing social and other inequalities via a back route which may not be caught by our equality law as it is currently understood or enforced.

Three challenges *continued*

Key Challenge 2: Meaningful explanation and justification

An explanation of how an machine learning system arrives at its decision can offer the subject a map of the programmatic pathway the system has taken with the available data, but cannot justify why the decision was based on that particular set of data.

The aspiration for a formal right for an explanation to underpin formal equality (as it exists) has evolved in a number of subject-specific contexts for example the Employment Rights Act, where it can be seen in the right to a statement of reasons for dismissal. More recently, the concept of a 'right to an explanation' has developed in the context of data protection. It is expressed in the General Data Protection Regulation ('GDPR'): "[the data subject should have] the right ... to obtain an explanation of the decision reached". Given the (AI) focus of this paper, GDPR is our starting point to examine explainability of firms to individuals with regards to automated processing.[32]

As discussed above, the debate on this 'right' has to date been dominated by the need for technical explanations for ML-guided decision-making (the 'black box' issue). This reveals a narrow focus on a narrow interpretation of Level 1: how the ML system can be explained to the individual.

The technical problem of understanding *how* a machine learning model arrives at its predictions or classifications is far from simple. ML, which is not limited to human approaches to analysis, perspective on causation or solving problems, may discover surprising or unintuitive relationships within data. There is increasing consensus that, despite three decades of work, the best approaches to the training machine learning algorithms are still not satisfactory for a legal or policy context.

Even meaningful explainability cannot justify the use of ML systems to determine aspects of human lives that would ordinarily be subject to legislation on equality and discrimination. An explanation of how an ML system arrives at its decision can offer the subject a map of the programmatic pathway the system has taken with the available data, but cannot justify why the decision was based on that particular set of data. Corporate explanations, now required under the GDPR, are to be welcomed as they increase the level of transparency in ML and AI decision making. But explanations alone will not suffice if we are to develop a framework for the use of ML systems within the workplace which places meaningful equality at the core of its design and use. Looking at the use of ML at work, we can see clearly that the implementation of ML should be explained and justified (to some extent) at the level of individuals, firms and social structures.

Further, specific limitations to the GDPR need testing. These may include the requirement that relevant data must be capable of identifying an individual and the limitation of the 'right to an explanation' being limited to decisions made by 'solely' automated systems resulting in a decision which has 'legal' or other significant effects.[33]

Three challenges *continued*

Key Challenge 3: Ensuring wider firm accountability

A full explanation of the decisions arrived at via AI and ML is pre-requisite to meet our third challenge: wider accountability at a corporate level. Increased use of ML in the workplace generate new questions about the power, responsibility and accountability of institutions that direct ML to achieve their goals in a systematic way, with systematic consequences. We hope our three-tier framework will help to facilitate a conversation about this too.

Our framework shows that – if we are to have a meaningful conversation about equality even just within the sphere of individual rights – we must consider equality at an institutional level too. We need a discussion about the implications of technology for the responsibilities of corporates within society as a whole, especially as market power (and data power along with it) becomes more concentrated in the digital age. This means putting people first at a collective and social level, not just as individuals. What are the consequences of our first and second challenges for our social contract? We should note that even at this systems level, the direction of ML is almost always guided by humans: it is humans who integrate the technology, determine the purpose of their business and operating model and who create the policy which oversees this innovation.

Our programmes under the 'Equality Through Transition' theme are being set up to respond to the particular Future of Work challenges we have identified. We will be working with partners and champions to develop and test these initiatives over the next two years.

We need a discussion about the implications of technology for the responsibilities of corporates within society as a whole, especially as market power (and data power along with it) becomes more concentrated in the digital age. This means putting people first at a collective and social level, not just as individuals.

Three programmes

In response to Key Challenge 1: A discrimination task-force

Our preliminary assessment has pinpointed a number of ways in which our existing formal legal framework may be tested by the nature of ML. In response to these challenges, IFOW is establishing *a task-force* chaired by Principal of Mansfield College, Oxford, Helen Mountfield QC, to examine how current direct and indirect anti-discrimination legislation applies to case studies, some real and some hypothetical. We will start by looking at the application of current law to our case studies, which will include cases of hiring and targeting algorithms used in the work space.

Task-force remit

The precise remit of the task-force will be determined at its first working meeting, but it is likely to cover:

(i) What amounts to ‘provision, criterion or practice’ and ‘particular disadvantage’ in the context of ML models?[34]

(ii) How can we assess ‘proportionality’ in the context of ML models?

(iii) How does the burden of proof operate in the context of ML?

(iv) How do Intellectual Property caveats interact with EA in the context of ML?

(v) Is there a case for adding socioeconomic disadvantage, or any other protected characteristic or combination of protected characteristics, to existing protection?

(vi) Are any new or other measures needed to achieve meaningful formal equality at an individual level?

The overall ambition of this task-force is to test our current legal framework for equality, its remit and operation to ensure it is achieving its purpose in the age of AI.

IFOW is establishing a task-force chaired by Principal of Mansfield College, Oxford, Helen Mountfield QC, to examine how current direct and indirect anti-discrimination legislation applies to case studies, some real and some hypothetical.

Three programmes *continued*

In response to Key Challenge 2: A fairness impact assessment

The main aim of the framework is to ensure that governments and businesses think through the multidimensional implications for equality when they chose to adopt machine learning.

IFOW will also develop with partners a simple fairness impact assessment tool, aimed at embedding and promoting equality from design onwards. Principals will be developed, applications explored and series of protocols will be created to promote best practice when designing and adopting a machine learning model in the workplace.

The main aim of the framework is to ensure that governments and businesses think through the multidimensional implications for equality when they choose to adopt ML. It also supports work under Key Challenge 3. In most cases, an explanation of the working of a machine learning model itself will not be sufficient. In our view, what is more important is why a model has been used, *what* are full implications of use and *who* is responsible for such use.

So our proposed framework will support best practice and provide meaningful information about choice, impact and responsibility in the following ways:

(i) encourage thinking about purpose for introducing the technology and different means and choices to achieve goals set

(ii) encourage consideration of implications for equality at each level

(iii) encourage documentation of decisions that impact on equality at each level

(iv) encourage especially high levels of practice where worker relationship involved

(v) create workplace culture which values and promotes equality – this will have a bearing on the design of tools themselves

(vi) drive and inform discussion about institutional accountability more broadly, with work as a lens through which we can focus on new challenges to equality – as discussed in the next challenge.

Our consultation is ongoing but scoping to date suggests the following principles will feature in the framework. In sketching the framework out, we have drawn from work on probing amendments tabled under the Data Protection Act.[35]

(i) Reasoning	(ii) Disclosing	(iii) Testing	(iv) Assessing	(v) Explaining
Employers and others should write and share a *'statement of purpose'* when using AI to support or make decisions which are relevant to the terms and conditions of work.	Employers and others should share the existence and type of AI model they use, along with its statement of purpose.	Large employers and others should devise and document a reasonable level of advance testing of AI-related systems when proposed for use in the work domain.[36]	Employers and others should devise, undertake and record regular 'algorithmic impact assessments' of AI-related systems to evaluate the potential for direct and indirect discrimination and other impacts on bias or fairness.	Workers should have a meaningful and actionable explanation of the *what, why and how* the system has impacted on the individual and their interests.

Three programmes *continued*

In response to Key Challenge 3: Accountability of the firm

Finally, we will explore the implications of our research and the programmes we have had outlined above for wider firm accountability. In particular, we will consider the extent of corporate employer responsibility to champion responsible use of AI in the work place, where stakes are particularly high. This responsibility may extend beyond the minimum currently required by legislation, and beyond the soft codes which exist more broadly for AI use too.

IFOW will comment on and ask the ICO, Data and Ethics Centre, and Equality and Human Rights Commission to consider our first and second programmes and their interests in them. We will also review their remits, responsibilities and means to respond to the challenges we have identified. This should support the important work of the Data and Ethics Centre and inform a wider debate about whether there is merit in establishing a new AI regulator or not.

In particular, we will consider the extent of corporate employer responsibility to champion responsible use of AI in the work place, where stakes are particularly high.

Conclusion

In his article 'Constitutional democracy in the age of artificial intelligence', Paul Nemitz[37] advises us to think about how we might create a culture which incorporates the principles behind human rights and discrimination legislation into new technologies by design.

In this paper we discuss how we draw from this philosophy of change to effect the policy, practice and legislation needed to make sure that AI promotes rather than diminishes equality. We judge this by the changing experience of work. Stepping outside the legal framework of equality, to the guiding principles behind it, is the first move. We can be guided by our values in our interpretation of equality in the new and exciting realm of artificial intelligence. If we can prioritise and implement principles of equality by creating frameworks and regulation for new technology, then our expectations and trust in these new systems will begin to change at a cultural level. Rather than see a tension between social equality and the introduction of AI, we can choose to see an opportunity and use human experience as a framework to think about the nature of equality and how to protect this fundamental principle as work is transformed. If we can embed these values at each level, we can build a culture of use (as well as adaptation and design) that has equality at its core, in the present and the future.

Rather than see a tension between social equality and the introduction of AI, we can choose to see an opportunity and use human experience as a framework to think about the nature of equality and how to protect this fundamental principle as work is transformed.

Endnotes

1 This principle is often referred to as 'Amara's Law', and the saying attributed to the Stanford Computer Scientist Roy Charles Amara (1925–2007). A discussion of the difficulty of finding the source of the quote, and its subsequent misinterpretations, can be found here: Matt Ridley, "Don't Write off the next Big Thing Too Soon | Comment | The Times," *The Times*, November 6, 2017, https://www.thetimes.co.uk/article/dont-write-off-the-next-big-thing-too-soon-rbf2q9sck.

2 A fuller account of the relationship between AI and ML, as well as ML's value in discussions of the future of work, can be found in the following section of this paper.

3 In July 2018, IFOW held a workshop in collaboration with the Oxford Internet Institute on the use of AI in work spaces. The briefing note for this workshop is available here: https://static1.squarespace.com/static/5aa269bbd274cb0df1e696c8/t/5b484fcd1ae6cf8d84dd5274/1531465677579/IFOW+OII+workshop+background+note.pdf, accessed 23 January, 2019.

4 For further discussion on the opportunities offered by AI and machine learning within work see: "You and AI – The Future of Work | Royal Society," accessed January 8, 2019, https://royalsociety.org/science-events-and-lectures/2018/09/you-and-ai/.

And Sanna Ojanperä, Neave O'Clery, and Mark Graham, "Data Science, Artificial Intelligence and the Futures of Work" (London, England: The Alan Turing Institute, October 2018), https://doi.org/10.5281/zenodo.1475162.

5 "Amazon Scraps Secret AI Recruiting Tool That Showed Bias against Women," *Reuters*, October 10, 2018, https://www.reuters.com/article/us-amazon-com-jobs-automation-insight-idUSKCN1MK08G.

6 Jeremias Prassl, Humans as a Service: *The Promise and Perils of Work in the Gig Economy*, 1 edition (Oxford, United Kingdom ; New York, NY: OUP Oxford, 2018)., pp. 136–137.

7 Julia Ticona and Alexandra Mateescu, "Trusted Strangers: Carework Platforms' Cultural Entrepreneurship in the on-Demand Economy:," *New Media & Society*, May 15, 2018, https://doi.org/10.1177/1461444818773727.

8 "Quick Take: Women on Corporate Boards,", Catalyst, November 27, 2012, https://www.catalyst.org/knowledge/women-corporate-boards.

And "Gender Diversity in Senior Management and Boards | Equality and Human Rights Commission," accessed January 30, 2019, https://www.equalityhumanrights.com/en/our-research/reading-lists/gender-diversity-senior-management-and-boards.

9 Paul Nemitz, "Constitutional Democracy and Technology in the Age of Artificial Intelligence," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 376, no. 2133 (November 28, 2018): 20180089, https://doi.org/10.1098/rsta.2018.0089., p. 3.

10 Weicheng Lian, Mai Chi Dao 'The Drivers of Declining Labour Share of Income' IMF (2017)
Martin Mühleisen, "The Impact of Digital Technology on Society and Economic Growth," *IMF F&D Magazine*, 06–18, https://www.imf.org/external/pubs/ft/fandd/2018/06/impact-of-digital-technology-on-economic-growth/muhleisen.htm.

11 "OECD Jobs Strategy - OECD Jobs Strategy," accessed January 30, 2019, http://www.oecd.org/employment/jobs-strategy/.

12 "World Development Report 2016: Digital Dividends," Text/HTML, World Bank, accessed January 23, 2019, http://www.worldbank.org/en/publication/wdr2016.

13 See "Artificial Intelligence Can Entrench Disparities—Here's What We Must Do – The Cancer Letter Publications," accessed January 14, 2019, https://cancerletter.com/articles/20181116_1/.
This article outlines the 'GIGO' (garbage in, garbage out) problem in data harvesting for AI systems.

14 See Ed Mackenzie, "Tight Labour Markets and Self-Service Beer: Is the Productivity Slowdown about to Reverse?," *Bank Underground* (blog), July 6, 2018, https://bankunderground.co.uk/2018/07/06/tight-labour-markets-and-self-service-beer-is-the-productivity-slowdown-about-to-reverse/.

15 J Scott Brennen, Philip N Howard, and Rasmus Kleis Nielsen, "An Industry-Led Debate: How UK Media Cover Artificial Intelligence" (Royal Society, 2019), https://reutersinstitute.politics.ox.ac.uk/sites/default/files/2018-12/Brennen_UK_Media_Coverage_of_AI_FINAL.pdf.

16 Alex J Wood et al., "Good Gig, Bad Gig: Autonomy and Algorithmic Control in the Global Gig Economy," *Work, Employment and Society*, August 8, 2018, 0950017018785616, https://doi.org/10.1177/0950017018785616.

Endnotes *continued*

17 Jacques Bughin and Jonathan Woetzel, "Global Trends: Navigating a World of Disruption | McKinsey," accessed January 30, 2019, https://www.mckinsey.com/featured-insights/innovation-and-growth/navigating-a-world-of-disruption.

18 See "Can We Teach Morality to Machines? Three Perspectives on Ethics for Artificial Intelligence — Oxford Internet Institute," accessed January 23, 2019, https://www.oii.ox.ac.uk/blog/can-we-teach-morality-to-machines-three-perspectives-on-ethics-for-artificial-intelligence/.

This was the topic of a conference held in New York City by Data & Society and partners in 2018 "Manufacturing Truth: Machine Learning and Bias," accessed January 23, 2019, http://www.newmuseum.org/calendar/view/1210/manufacturing-truth-machine-learning-and-bias.

19 "A Fairwork Foundation: Towards Fair Work in the Platform Economy — Oxford Internet Institute," accessed January 23, 2019, https://www.oii.ox.ac.uk/research/projects/a-fairwork-foundation-towards-fair-work-in-the-platform-economy/.

20 For discussion of the democratic and constitutional risks posed by the market monopoly held by the 'big five' (Apple, Facebook, Google, Microsoft and Amazon), see Paul Nemitz, op. cit.

21 L. Edwards and M. Veale, "Slave to the Algorithm? Why a 'right to an Explanation' Is Probably Not the Remedy You Are Looking For," *Duke Law & Technology Review* 16, no. 1 (2017): 18–84.

22 Alan Winfield, "A Round Up of Robotics and AI Ethics," *Alan Winfield's Web Log* (blog), December 23, 2017, http://alanwinfield.blogspot.com/2017/12/a-round-up-of-robotics-and-ai-ethics.html.

23 AI Now Institute. "AI Now 2018 Report." NYC, New York, (2018). https://ainowinstitute.org/AI_Now_2018_Report.pdf, pp. 28–29.

24 As well as the UK AI Council, the House of Lords has established an AI Select Committee: "Artificial Intelligence Committee," UK Parliament, accessed January 23, 2019, https://www.parliament.uk/ai-committee.

25 OII, 'Fairwork', op. cit.

26 "University Researchers Present to Commission on Future of Work — University of Leicester," Blog Entry, accessed January 30, 2019, https://www2.le.ac.uk/staff/announcements/university-researchers-present-to-commission-on-future-of-work.

27 AI Now Institute, "AI Now 2018 Report" (NYC, New York, 2018), https://ainowinstitute.org/AI_Now_2018_Report.pdf., p. 29

28 Brennen, Howard, and Nielsen, "An Industry-Led Debate: How UK Media Cover Artificial Intelligence." Op. cit.

29 Jędrzej Niklas and Seeta Peña Gangadharan. "Data-Driven Discrimination: A New Challenge for Civil Society." *Impact of Social Sciences* (blog), July 10, 2018. http://blogs.lse.ac.uk/impactofsocialsciences/2018/07/10/data-driven-discrimination-a-new-challenge-for-civil-society/.

Tom Simonite, "AI Is the Future—But Where Are the Women? | WIRED," *WIRED*, August 17, 2018, https://www.wired.com/story/artificial-intelligence-researchers-gender-imbalance/.

30 See "Your Rights under the Equality Act 2010 | Equality and Human Rights Commission," accessed January 30, 2019, https://www.equalityhumanrights.com/en/advice-and-guidance/your-rights-under-equality-act-2010#h3.

31 AI Now Institute, "AI Now 2018 Report.", p. 26.

32 Like discrimination, this 'right' exists across domains but is more keenly experienced in the higher-stakes environment of the work place.

33 Veale and Edwards, op. cit.

34 See Section 19 of the Equality Act (2010), Expert Participation, "Equality Act 2010," Text, accessed January 30, 2019, https://www.legislation.gov.uk/ukpga/2010/15/section/19.

35 See The Data Protection Bill, House of Commons, Probing Amendments, 'Employer's duty to undertake an Algorithmic Impact Assessment', March 2018.

36 We suggest comparing application of the Charter principals before and with use of the system. Testing should encourage thoughtful introduction and mitigate risk where impact may engage rights or interests concerning fundamental aspects of working life: recruitment, terms and conditions of employment, access to opportunities for promotion, transfer or training, and dismissal.

37 Nemitz is the Principal Advisor to the Directorate General for Justice and Consumers of the European Commission.

Institute for the Future of Work
Somerset House, The Strand
London WC2R 1LA
T +44 (0)20 7123 4567

www.ifow.org
@_futureofwork