



The Work Couch

NAVIGATING TODAY'S TRICKY PEOPLE CHALLENGES TO
CREATE TOMORROW'S SUSTAINABLE WORKPLACES

Episode 14 – AI (Part 2): Privacy, bias and discrimination with Patrick Brodie and Jake Wall

Content note:

Please note that all information in this episode is correct at the time of recording on 21 May 2024

- Ellie:** Hi and welcome to the Work Couch podcast, your fortnightly deep dive into all things employment. Brought to you by the award-winning employment team at law firm RPC, we discuss the whole spectrum of employment law with the emphasis firmly on people. My name is Ellie Gelder. I'm a senior editor in the employment equality and engagement team here at RPC, and I will be your host as we explore the constantly evolving and consistently challenging world of employment law and all the curveballs that it brings to businesses today. We hope by the end of the podcast, you'll feel better prepared to respond to these people challenges in a practical, commercial and inclusive way. And to make sure you don't miss any of our fortnightly episodes, please do hit the like and follow button and share with a colleague.
- To mark [London Tech Week](#), running from 10-14th June this year, we are thrilled to bring you our 3-part mini-series on the ever-evolving and exciting topic which is Artificial Intelligence.
- Last time, in part 1, we focused on the impact of AI on litigation, how to best use it responsibly and how to regulate AI, including a discussion of the EU's AI Act.
- Today, in part 2, we're going to explore how AI interacts with and affects employment law, specifically the issues of privacy, bias and discrimination.
- With me to share their expert insights, I am thrilled to welcome to the Work Couch Jake Wall, Policy Manager in Skills and Future of Work at [techUK](#), and [Patrick Brodie](#), RPC partner and head of Employment, Engagement and Equality.
- Hi both, thank you so much for joining us today!
-
- Patrick:** Thanks, Ellie, it's great to be with you.
-
- Jake:** Thank you, Ellie.
-
- Ellie:** Jake, if I can start with you, can you just tell us about your role at techUK and why you chose to go into this sphere?
-
- Jake:** That's a really good question. My role centres on skills and the future of work at techUK. So that means I'm focused on how the UK can equip people and businesses in every part of the country with the skills and conditions they need to thrive amidst the opportunities of the future workplace. In practice, that means I speak to stakeholders from government, industry, academia, trade unions and beyond about technologies like AI and employment trends like flexible work and the gig economy and what they mean for policymakers. I chose to go into this area looking at skills and the future of work because I really am passionate about the diverse routes into tech jobs. As a teenager, I dropped out of secondary school with no qualifications, and it was through technology-enabled distance learning and the Open University that I was actually able to re-engage with learning and return to education. So having access to those distance learning opportunities with no prerequisites really did change my life. And there is a Malcolm X quote, I think, that says, "Education is a passport to the future". And when I was able to take up that OU course, it really did feel like that. You can see a world of opportunity ahead of you. And after that foray into distance learning through the OU, I was actually able to use that experience to get into a brick-and-mortar university and ultimately ended up with a master's degree. Both of my degrees were focused on international politics, and I was always more of a social scientist. And although my role now is centred on technology, it's also very much about people. It's people that develop, deploy, and use the technology. So I think I'm fortunate that I'm in a role where I can really talk

about the benefits of technology and speak to the human element too, informed by my own somewhat atypical background. And interestingly, what we're seeing now in the context of quite stark skills gaps, particularly in technology, is that distance learning, retraining, upskilling, and taking different paths into jobs are not just becoming more common but encouraged. So to be able to champion the many routes into tech careers, whether it be a skills bootcamp, online courses, apprenticeship, degree apprenticeship, traineeship and ensure that people have access to those life changing opportunities is something really quite close to my heart.

Ellie: That's fantastic, Jake. Thank you for sharing your experiences and that really speaks to how we like to cover topics on the Work Couch. We like to look at the people element as well. So can you tell us a bit more about techUK's involvement with the Trades Union Congress, TUC, AI at Work advisory panel on platform workers and decision-making?

Jake: Yes, absolutely. So, the Trades Union Congress created the [AI at Work Bill Advisory Task force](#) to inform their work on drafting an AI at Work Bill with Cloisters Chambers. The Bill is a TUC product that aims to fill perceived gaps in UK employment law as it pertains to the use of AI at work. So techUK was really pleased to join the task force, which held only an advisory role to bring a strong business voice to those discussions and we joined the task force, noting that we did not intend to endorse a final Bill. And our advice was focused on the workability of any potential legislation. So we were really eager to see that the draft Bill did not result in onerous compliance burdens for businesses, particularly SMEs, that definitions around key concepts like AI were appropriate, and that the Bill remained focused on employment rights and didn't burgeon into a kind of wider AI Bill. To that end, we were really happy to see that some of our recommendations were taken on board by the TUC around definitions and overall focus, and that the Bill does not seek to regulate AI technologies more broadly. But techUK does hold a number of concerns with the published Bill. AI is set to transform the workplace not all uses of AI at work will have the same implications, and it will be key to ensure that any proposed, legislation recognises this and doesn't capture all applications with the same requirements. The aim of the TUC Bill is to encourage employers to take greater account for those more significant decisions, such as relating to hiring, dismissal, pay award progression, appraisals and disciplinary action. So we therefore recommended that the definitions of higher risk applications, of AI, were narrowed to focus only on those significant employment decisions and exempting routine functions of AI or algorithmic systems and removing the onus on employers to prove the sentence on a high risk. And we think that risk assessments for businesses underpinned by strong guidance could reduce compliance burdens where the use of workplace AI is really low and low risk. It was really great to be part of this conversation with the TUC and partners on the Bill, as it's a really important topic and a fast-moving space. Getting the balance right to foster innovation while ensuring good standards of work. It's going to be a key challenge for any future government. Moreover, keeping on track with developments in AI and their implications for work and workers is going to require a collaborative approach that brings in perspectives from across the piece, including businesses and workers. So this forum was a really welcome step in that direction.

Ellie: Thank you for that background, Jake. Patrick, can you tell us a bit more about what the Bill means for employment law and the UK's workforce?

Patrick: Yeah, of course, Ellie. So the AI Bill is in many ways simply a starting point. It was created and published to act as a catalyst for people to think about how the adoption and reliance on AI in the workplace, which often affects people's rights and working conditions, should be properly regulated. At its core, the Bill seeks to ensure that workers are protected with also the right to understand and challenge decisions that affect them. So we should treat it as an early contribution to the development of policy on effective regulation of AI. Safeguards afforded to workers are the foundation blocks of the Bill, but equally the Bill seeks to recognise that these safeguards don't have to be a trade-off to innovation and technological evolution. That's because it's recognised in a competitive global marketplace, businesses must succeed, but the Bill proposes that this success, and this is a balancing act, should be in line with the protection of employees and civil liberties in a democratic society. The cornerstone of the Bill is the meaning of AI. Indeed, all of the proposed protections and safeguards offered to employees, flow from this definition, which is intended to embrace a modern meaning of what is often a debated term. From this definition, the Bill then looks to translate into rights and obligations broadly four key principles and values: transparency, explainability, equality and worker consultation. The Bill recognises that existing legislation provides workers with some safeguards in circumstances where AI decision-making output leads to unjustified detriment, inequality, or harm. For example, a decision adversely affecting an employee might engage the Equality Act, Employment Rights Act or GDPR. However, the research group supported by the work of two brilliant barristers, again, Jake mentioned them from Cloisters, suggests that the existing legal framework needs to be bolstered to address legal gaps. The key gaps relate to a lack of transparency and explainability in AI, an imbalance in power over data with employers holding the principal right to data, the lack of a work voice in consultation on the adoption

of AI, and a weakness in the protections against discriminatory algorithmic output. But as said, the Bill is a starting point, not the end destination. And there will clearly be much more consultation and work.

Ellie: Thanks Patrick, that's a really useful overview of some of the key gaps there to think about. And Jake, I know that techUK has also recently published a report called [Making AI Work for Britain](#). Can you just tell us a little bit more about this report?

Jake: Of course. So our Making AI Work for Britain report was published at the end of 2023. And it explored the impact of AI on jobs and skills across the UK. The intention of the report was to cut through some of the maybe more well-known or entrenched arguments around AI replacing jobs to inform a better understanding of how jobs may change as AI is adopted, including a growth in AI augmented work undertaken by humans, how we drive adoption to capitalise on the productivity benefits of the technology and set out what action must be taken to ensure our regulatory environment and education ecosystem is fit for an AI enabled future. But to give a brief overview of a fairly comprehensive report, So, firstly, adoption, we found the degree to which AI has been adopted varies depending on industry, company size and region. Most notably, SMEs are less likely to be digital adopters and the not insignificant minority still do not use basic digital tools. So as AI is increasingly integrated into everyday digital tools, usage is accelerating. But those businesses behind on digital adoption remain behind the curve, and tackling that adoption gap will therefore be key for business success and competitiveness, but also reducing labour market disruptions too. Secondly, this adoption will change skills and jobs. So according to Experis research, half of British companies say AI will boost headcount over the next two years, and the commercialisation and widespread adoption of AI will require multi-skilled teams. And that means not just AI and machine learning experts, but also strong soft skills, domain specific and business knowledge and expertise in complementary and enabling technologies like cyber security and cloud computing. In the report, we speak about the era of AI competence, soft skills and human expertise, meaning people will need the AI competence to use AI effectively. The soft skills that are hard to automate that will be critical in working with AI outputs and interactions with colleagues, clients and customers. And speaking to some law firms, we recognise that the use AI in legal has a number of applications, such as helping to automate forms and contracts and analysing large volumes of data and documents. According to Raconteur, 63% of legal work consists of repetitive, fact-based decisions and this is prime for automation. The ability to use AI to more quickly analyse documents and case law can free up considerable time for law professionals but as in other sectors, the most effective use of AI in law is likely to be found when deployed in conjunction with strong human expertise. Present tools are probably unlikely to give due consideration to legal principles like a human would and equally, lawyers have an obligation to provide clear information and not mislead their clients. So that would mean astute legal acumen alongside strong AI and soft skills will remain critical and it's important to remember that AI will develop alongside other emerging technologies like virtual reality and augmented reality, changing skills and jobs in other ways. Thirdly, the report includes a section dedicated to good management practices. If it was as simple as technology adoption equals boosting productivity, the productivity puzzle would look like less of a puzzle. But it is people adopting, developing, supervising, and using the technology that will ultimately determine its impact. So good management practices, including training, but also transparency and information sharing on the use of technology at work, are therefore key to driving productivity gains. Workers that can collaborate effectively with AI report higher levels of job satisfaction and engagement. And a study by PWC actually found that 75% of employees know of systems that would help them produce higher quality work. I think that really goes to show how empowering workers' voices can enable innovative solutions to common problems. And finally, I will touch on regulation. the UK's approach to AI regulation sets out five key principles, to be overseen and applied by existing regulators: safety and security, transparency, fairness, accountability and contestability.

Ellie: So Patrick, we've heard about Tech UK's involvement in the sphere of AI and the TUC AI Bill. So I wanted to dive into talking about how the evolution of AI more generally might impact various elements of employment law. So to start off, when it comes to decision-making processes being increasingly influenced by AI algorithms, what are the key considerations to ensure transparency and accountability, especially in sectors with significant employment implications?

Patrick: Gosh, that's a big question. But I'll try and break it down. First, you're right. Transparency, or put another way, explainability alongside privacy and bias are the three main concerns held by many about AI.

I'm seeing an increasing focus on AI ethics policies. Those policies will be unique to an organisation with the language and emphasis reflecting organisation's own personality and character. Some broad themes and principles do emerge. Firstly, senior cross-disciplinary leadership groups are coming together. Typically, they'll cover technology, risk regulation, and human resources. And that group will have AI oversight and accountability. Now, that oversight might engage various areas of focus, whether ensuring AI is trustworthy and ethical. They'll look at protecting and mitigating against the risk of unlawful bias being replicated or advanced, consideration around the fair and lawful collection, retention, and processing of data. And then as a foundation principle, AI being transparent, fair, and then the language of explainability and equally ensuring its robustness. Commonly organisations will undertake impact assessments to determine the risk and potential adverse consequence of AI systems adopted by the company. And that will be with a particular focus on high-risk AI which might, for example, include, an algorithm that nudges users to a particular outcome or preference, or it might draw inferences that might be said to be invasive, they could be counterintuitive, indeed they might be non-verifiable or from multiple sources. The impact assessment will review and assess and address the ethical consequences of the adopted AI, including in relation to its potential failure. And for high-risk AI, there should always be some form of human intervention or review to verify the process or the decision. To remove systemic unlawful bias that might subsist in the AI systems, organisations are becoming ever more focused on reviewing, assessing, and understanding the data used, the variables adopted. The decision-making process by which AI achieves an output should be explainable. This may include how and why an output was reached, including if the output was wrong. This explanation then allows a person, whether the employer or the worker, to understand the decision, the sensitivity of the model, and what factors can be changed for a different decision.

Ellie: How do you foresee that interplay between evolving AI technologies and jobs for the future?

Patrick: Yeah, I'm acutely conscious that the fear of job losses because of technology and automation, including AI, has been with us since the 60s. Studies have suggested that in the near term, about 50%, could extend up to 80%, of workers could see their roles affected by AI. That's a significant number. Now a sobering statistic, genuinely sobering statistic, is that with the effective adoption of convergent technologies and tools, about 50% of all workers' tasks could be delivered significantly faster without reducing quality. Now, with that in mind, our thoughts typically turn to tasks that are routine and repetitive. That's because it's better for technology to absorb this work, freeing people up for those really interesting, challenging, and rewarding tasks. However, with the rise of generative AI and increasingly sophisticated machine learning, a lot of non-routine, creative, and knowledge-based tasks, which until recently were seen as a preserve of us humans will be capable of AI replication. And I'll feel this as much as any in relation to being a lawyer in the legal sector. At its most basic, I'm a knowledge-based worker. And of course, the clouds have been gathering for a few years. And I've recognised that my task, which at some point, probably up until the 2020s, felt out of reach of the machines because it's language-based, and it relies on multiple sources of information and data. I realise that those tasks will no longer be delivered by me. I recognise again, reflecting on what Jake said, my work will be augmented. It will be augmented by deep learning in particular generative AI with its language recognition and reconstruction capabilities. And that's then going to be, it'll be supercharged with different forms of supervised, unsupervised and reinforcement learning. And, I simply think that that synthesis of computational power, data volume, provided the data is effectively curated a really important element and predictive capabilities, are supported by human intervention. It will change the landscape of the workplace.

However, I remain hopeful that it's not going to lead to an existential societal crisis. My comfort remains founded, I think, in the belief, and it's as with prior industrial revolutions, that new roles and skills will be created and emerge. For example, a few years ago, did we have prompt engineers, AI trainers, or machine managers? Equally, again, something Jake touched on, intrinsic human qualities will remain. Indeed, they'll thrive. So really simple example, an Uber driver who provides their passenger with a service that's warm, they'll be rated more highly. But it's not to say that my horizon doesn't carry small grey clouds, it does. And I worry about the risk of automation, displacing large numbers of roles and doing so disproportionately impacting those more economically vulnerable or in precarious roles or roles that would ordinarily be those taken on by school leavers or those new to the workplace. I think organisations will have to think really carefully about how AI will impact their workforces and equally in turn their responsibilities to local communities and the wider society, especially those in underrepresented groups or who in economically vulnerable employment might be most impacted.

-
- Ellie:** And Jake, I just wanted to touch on the topic of bias in AI systems, which is a topic that's recently attracted attention, particularly concerning gender, race, neurodiversity, and other factors. How do you think we can address and mitigate those biases? And what challenges would you say persist in effectively combating those?
-
- Jake:** So I think this is a really important issue, and one that's probably at the heart of the broader public trust debate around AI. But what is significant is that AI is not innately biased. Biases arise because AI has been created and developed by humans and trained on human data. The biases reside in humans, but risk perhaps becoming entrenched by technology these biases already exist in traditional or non-AI processes, so a growing awareness and impetus to address them in the context of AI actually creates an opportunity to develop more equitable outcomes or processes. And these algorithms biases may ultimately be held to stronger account than years of human biases. And systems in their data as they pertain to decision-making can be more easily interrogated than humans, and maybe even more easily improved and rectified as well. Recruitment is a really big part of the conversation in terms of AI bias in decision-making. What's interesting is that there are organisations that are using AI in recruitment for the sole purpose of mitigating bias in traditional processes and ensuring fairness through all stages of the recruitment process. So MeVitae as a techUK member focussed on fair hiring. They have solutions for blind recruiting that use augmented intelligence to anonymise personal identifying information in CVs and for talent screening, using the technology to pass past, present and future data and predict candidates' suitability without algorithmic or cognitive bias. And there's some really great results for broadening talent pools, saving time and cost, and boosting gender diversity, too. So I think that shows there is really an opportunity as well as a risk. But humans will need the tech and AI literacy and ethical reasoning skills to supervise and monitor for issues like bias and they'll need to have a good understanding of the role that AI should play, how it should synergize with their own human judgement and expertise and an awareness of its limitations. So that will need to be supported by good internal policies and processes and good management practices. We know that includes training. A good AI governance is definitely part of that too. However, there is no one size fits all approach to AI assurance and some organisations are only going to have the resources to focus on key areas of highest risk. So there will need to be proper support for these smaller businesses, especially to help them with their AI governance and manage risks in a way that's most appropriate and effective for them.
-
- Ellie:** Thank you, Jake. It's really interesting to hear the opportunities as well as the risks that AI presents in terms of bias. And Patrick, with AI's increasing role in recruitment and other HR processes, if you can outline for us the legal challenges that might arise concerning fairness, discrimination, and the protection of employee rights in this sphere.
-
- Patrick:** Of course. Jake speaks to the sunny uplands of the benefits that AI will bring and as a lawyer I take you down to the sort of the gloomier valleys of the risks. So my apologies. We know that organisations, employers are increasingly aware of risks and weaknesses in AI, especially as its capabilities we know are only as good as the data on which it's trained and tested, we've learned of the risk that if the training data is flawed, then the outputs will be similarly weak and imperfect. The outputs then risk maintaining or amplifying the bias present in the data, or indeed even in the prompts on which the model was built. For example, I accept that this somewhat historic example, but it's to illustrate a point. AI models that had been previously built, which is seeking to predict the best candidate for leadership role, had a data bias towards middle-aged white males. And that was because after all the predominant characteristic, historic characteristic based on that historic data was that people in corporate leadership roles were predominantly white men. And that you can see that that creates the implicit bias within the data. Now, a narrow answer to this is to increase the data in relation to the underrepresented group. However, breaking down bias also requires those creating the models to be aware of the problem. We see biases in data across all characteristics, race, gender, sexual orientation, religion. And sometimes that's about who's in the room building the models and that access or the absence of access to diverse thought. In a setting outside of employment but in my mind, nonetheless, it perfectly illustrates the point.

There was a study in the US that analysed the accuracy and bias of AI algorithm that determined and apportioned medical funding for patients. When writing the model, one of the statistical data variables was how much individuals spent financially on health care against available medical data. Individuals were given an equal weight in relation to this data. Instinctively, this equality suggests also an equality of treatment. However, it failed to recognise that the fact that black patients spent less on health care in the early stages of illness. Now, the effect of this unanticipated input flaw was that the AI model determined wrongly that black patients needed less future health funding. Now, the bias flowed from the algorithmic modelling and the incredibly weighted variables. So without intervention, the model would have amplified an existing bias. Turning back to employment, we're told, for example, that the recruitment process will be quicker and more objective and accurate if we use AI. And indeed, there are lots of examples of that being right. Then there are AI systems that can be introduced to improve pooling and selection in large -scale redundancy exercise. The model is built by being trained and tested on historic workforce data before it's then applied to the existing exercise at hand. Now, anyone adopting an AI model should have confidence that the model's predictive output doesn't risk bias or unfair or unexplained decision making. Otherwise, the inevitable risks of unfair dismissal, the impermissible interference with trust and confidence, or indirect discrimination arise. Or indeed, the risk of direct discrimination if there's human intervention to resolve algorithmic bias by rebalancing the output to address the underrepresentation. Now, inevitably, AI systems involving the collation and processing of data will also need to be GDPR compliant.

The legal risks, many of which may not be immediately obvious at the point of the adoption of the AI, will have to be carefully worked through.

Ellie: Well thank you so much Patrick and Jake for joining me today and explaining the interplay between AI and employment law and a raft of important factors for businesses to take into account now as they embark on their AI journeys.

Jake: Thank you very much Ellie.

Patrick: Thank you Ellie.

Ellie: And I'm delighted to say that Patrick and Jake will be joining us again next time on the concluding part to our mini-series on AI when we'll be looking at the role of emotional intelligence and AI's impact on employees' wellbeing so we hope you'll join us then for what promises to be another fascinating discussion.

If you would like to revisit anything we discussed today, you can access transcripts of every episode of The Work Couch podcast by going to our website: www.rpc.co.uk/theworkcouch. Or, if you have questions for me or any of our speakers, or perhaps suggestions of topics you would like us to cover on a future episode of The Work Couch, please get in touch by emailing us at theworkcouch@rpc.co.uk – we would love to hear from you.

Thank you all for listening and we hope you'll join us again in two weeks.



RPC is a modern, progressive and commercially focused City law firm. We are based in London, Hong Kong, Singapore and Bristol. We put our clients and our people at the heart of what we do.