

Algorithmic Management: Assessing the Impacts of AI at Work

Aislinn Kelly-Lyth and Anna Thomas*

Abstract:

Algorithmic outputs are increasingly shaping the employee experience, presenting a host of risks and impacts with far-reaching consequences. This contribution considers how algorithmic impact assessments should complement, as well as inform, an overarching ‘top-down’ framework for the governance of algorithmic management systems. While generalised obligations are crucial, identifying risk mitigations on a case-by-case basis can provide significant added value by (i) evaluating impacts and facilitating *context-specific* responses to risks and impacts identified; (ii) striking a *balance* between generalised requirements and complete self-regulation; and (iii) ensuring that due regard to anticipated impacts and risk mitigation is *built in* from the design and development stages, through to deployment in the workplace. The criteria for an effective impact assessment obligation in the algorithmic management are identified, including the appropriate stages, actors, and procedure. The Good Work Charter, which operates as a synthesis of legal principles, rights, and obligations, as well as ethical principles as they apply to the workplace, is proposed as an assessment framework. Finally, the article compares the proposed model with the existing obligation to carry out data protection impact assessments for high-risk data processing. The shortcomings of the latter obligation are explored, and a legislative approach to avoid duplication is proposed.

Keywords: employment, algorithmic management, technology, impact assessments, General Data Protection Regulation

1. Introduction

Algorithmic outputs are increasingly shaping the employee experience.¹ The concept of using ‘algorithmic management’ systems to instruct, evaluate, discipline, and reward workers has quickly spread from the gig economy into more ‘traditional’ workplaces.² By 2018, 75 percent of surveyed HR professionals were already using data to ‘understand . . . workforce performance and productivity’, with 14 percent using machine learning to develop ‘people

* Researcher, Bonavero Institute of Human Rights, University of Oxford; Director, Institute for the Future of Work (IFOW). Aislinn acknowledges funding from the European Research Council under the European Union's Horizon 2020 research and innovation programme (grant agreement No 947806). We are grateful to Jeremias Adams-Prassl, Halefom Abraha, Six Silberman, James Hayton, Abigail Gilbert, Stephanie Sheir, Thomas Serban Von Davier, Lu Zhang, Lara Groves, Aidan Peppin, Phoebe Moore, Aude Cefaliello, and Lukas Hondrich for their thoughts on earlier versions of this work.

¹ Abigail Gilbert and Anna Thomas, ‘The Amazonian Era—The Gigification of Work’ (Institute for the Future of Work 2021). Institute for the Future of Work, ‘Introducing the Pissarides Review into the Future of Work and Wellbeing’ (Institute for the Future of Work 2022) 6.

² Cansu Safak and James Farrar, ‘Managed by Bots: Data-Driven Exploitation in the Gig Economy’ (Worker Info Exchange 2021); Wolfie Christl, ‘Digitale Überwachung und Kontrolle am Arbeitsplatz’ (2021); Gilbert and Thomas (n 1).

reports'.³ By July 2021, 66 percent of firms reported new adoption of digital management practices since March 2020.⁴ Some of the most widely deployed operational software packages now provide granular insights into worker activities and behaviour, shaping decision-making at the group and individual level.⁵

The use of technology to evaluate and manage workers is not a new phenomenon. Frederick Taylor was already arguing for a 'scientific management' approach to labour by the late 1800s, seeking to maximise efficiency by reducing the time taken for individual tasks.⁶ What *is* new is the scale, ubiquity and impact of these practices.⁷ The 'rapid erosion of technological and economic constraints on employee monitoring' has facilitated the collection of worker data on a hitherto unimaginable scale,⁸ and employers are increasingly able to amalgamate and analyse this data in order to automate or augment managerial decisions.⁹ In addition, as performance and capabilities are anticipated, rather than assessed, people analytics tools offer new advisory and predictive managerial functions, alongside traditional ones, to select or nudge behaviours.¹⁰ In these circumstances, algorithmic management is now widely recognised as a socio-technical concept, in the same way that algorithmic systems are a 'dynamic arrangement of people and code',¹¹ reflecting the technological and organisational infrastructures¹² and human choices involved.

It is clear that algorithmic management practices impact workers.¹³ Some of these harms are material: in Amazon warehouses, for example, where workers' movements are tracked down to the second, injury rates are reported to be 80 percent higher than the average.¹⁴ Research has also begun to document subtler and less visible impacts, such as a reduction in people's autonomy at work as choices are restricted or undermined.¹⁵ Meanwhile, threats to privacy and data protection have grown with the intensity of collection of personal data,¹⁶ alongside impacts on equality of opportunity and outcome: automated systems process huge swathes of data encoding past patterns of behaviour to produce outputs which may compound

³ CIPD and Workday, 'People Analytics: Driving Business Performance with People Data' (2018) 30, 33 <https://www.cipd.co.uk/Images/people-analytics-report_tcm18-43755.pdf> accessed 7 March 2022.

⁴ Anna Valero, Capucine Riom, and Juliana Oliveira-Cunha, 'The Business Response to Covid-19 One Year on: Findings from the Second Wave of the CEP-CBI Survey on Technology Adoption' (Centre for Economic Performance 2021).

⁵ Christl (n 2).

⁶ Matthew T Bodie and others, 'The Law and Policy of People Analytics' (2017) 88 University of Colorado Law Review 961, 964–968; Ifeoma Ajunwa, Kate Crawford, and Jason Schultz, 'Limitless Worker Surveillance' (2017) 105 California Law Review 735, 740–747.

⁷ Recent years have seen a qualitative shift in the use of algorithmic management systems: Gilbert and Thomas (n 1) 9; Jeremias Adams-Prassl, Six Silberman, Halefom Abraha, and Aislinn Kelly-Lyth, 'Regulating Algorithmic Management: A Blueprint' (elsewhere in this issue).

⁸ Ajunwa, Crawford and Schultz (n 6) 743.

⁹ Institute for the Future of Work (n 2) 6.

¹⁰ Gilbert and Thomas (n 1).

¹¹ Nick Seaver, 'Captivating Algorithms: Recommender Systems as Traps' (2019) 24 Journal of Material Culture 421.

¹² Mohammad Hossein Jarrahi and others, 'Algorithmic Management in a Work Context' (2021) 8 Big Data & Society 20539517211020332.

¹³ Jeremias Adams-Prassl, 'What If Your Boss Was an Algorithm? The Rise of Artificial Intelligence at Work' (2019) 41 Comparative Labor Law & Policy Journal 123. Research suggests that impacts on work span access, conditions and quality of work: Gilbert and Thomas (n 1); Global Partnership on Artificial Intelligence, 'AI for Fair Work: AI for Fair Work Report' (November 2022) <<https://perma.cc/99DL-LJNL>>.

¹⁴ Strategic Organizing Center, 'Primed for Pain: Amazon's Epidemic of Workplace Injuries' (2021) <<https://thesoc.org/wp-content/uploads/2021/02/PrimedForPain.pdf>> accessed 7 March 2022.

¹⁵ Gilbert and Thomas (n 1) 32.

¹⁶ Alessandro Mantelero, *Beyond Data: Human Rights, Ethical and Social Impact Assessment in AI* (Springer 2022).

different types of inequality at a group and individual level without intervention.¹⁷ Some deployments of algorithmic management pose a direct threat to the realisation of rights, such as attempts to combat unionisation by analysing worker interactions.¹⁸ In short, algorithmic management presents a host of risks and impacts with far-reaching consequences.

This contribution will consider how algorithmic impact assessments should complement, as well as inform, an overarching ‘top-down’ framework for the governance of algorithmic management systems. The model proposed combines different approaches to risk and impact assessment across the dimensions which shape the employment experience. It does so by deploying a framework for the evaluation of impacts on the conditions and quality of work: the Good Work Charter.¹⁹ The Good Work Charter maps onto the broad span of the impacts of algorithmic management, from access to conditions and quality of work. Our approach enables context-specific evaluation from legal, ethical, and societal perspectives and obligations synthesised in the framework proposed, thus broadening the scope of potential harm mitigations to include those that can only be identified at the more granular level.

We begin in section 2 by defining the scope of the idea, identifying the algorithmic tools which would be subject to assessment. Section 3 considers how a case-by-case consideration of potential harms can complement top-down regulation which sets out obligations which apply in all cases. A concrete proposal for a new legislative duty is set out in section 4. Section 5 considers the extent to which the data protection impact assessments obligation imposed by the General Data Protection Regulation (GDPR) already effectively realises this proposal. Our model is aimed at new primary legislation, but it is also suitable for adoption in statutory codes or guidance developed by regulators, as proposed by the UK’s emerging AI Strategy.²⁰ It therefore doubles as an ‘assurance’ model to encourage early intervention and socially responsible behaviours, building understanding and capacity as part of the ‘paradigm shift’²¹ beyond compliance and technical audits.²²

2. Algorithmic management tools

The regulatory issue considered in this contribution is the use of technology by employers *qua* employers, with potential harmful impacts on workers. The purpose of the regulatory response should be to address or minimise those impacts in a proportionate manner. This purposive focus

¹⁷ Solon Barocas and Andrew D Selbst, ‘Big Data’s Disparate Impact’ (2016) 104 California Law Review 671; Anna Thomas and others, ‘Mind the Gap: How to Fill the Equality and AI Accountability Gap in an Automated World’ (Institute for the Future of Work 2020); Jeremias Adams-Prassl, Reuben Binns, and Aislinn Kelly-Lyth, ‘Directly Discriminatory Algorithms’ [2022] Modern Law Review (Early View).

¹⁸ See, for example, Jason Del Rey and Shirin Ghaffary, ‘Leaked: Confidential Amazon memo reveals new software to track unions’ (Vox/Recode, 6 October 2020) <<https://perma.cc/WQF2-A9Y4>>.

¹⁹ This approach is based on the premise that changes to the nature, conditions and quality of work, and its organisation, are shaped by algorithmic systems: Gilbert and Thomas (n 1). On the value of sector-specific algorithmic impact assessments, see Alessandro Mantelero, ‘AI and Big Data: A Blueprint for a Human Rights, Social and Ethical Impact Assessment’ (2018) 34 Computer Law & Security Review 754, 755. The Institute for the Future of Work’s Good Work Charter can be found at <<https://www.ifow.org/publications/the-ifow-good-work-charter>> accessed 4 September 2022.

²⁰ Department for Business, Energy & Industrial Strategy and others, *Establishing a pro-innovation approach to regulating AI* (Command Paper CP 728, July 2022), available at <<https://www.gov.uk/government/publications/establishing-a-pro-innovation-approach-to-regulating-ai>> accessed 4 September 2022.

²¹ Mantelero (n 16) describes this ‘paradigm’ shift as a response to understanding the limitations of a rights-based approach that hinges on individual control and consent over the use of personal data.

²² We note that the Information Commissioner’s Office is supporting voluntary guidance on undertaking ‘good work audits’ produced by the IFOW (forthcoming).

implies a crucial limitation: our concern is with the impact of tools which affect in-work experiences, with a particular focus on work conditions and quality, rather than on technologies which displace labour.²³ This focus maps a recent shift in regulatory discourses on technology and labour, with increasing recognition of algorithmic management's impacts on working conditions and quality now informing the policy response.²⁴

Clarity of purpose enables clarity of definition: we are concerned with impacts of technologies which affect in-work experiences, and those technologies can be defined by reference to their functions. A function-centric approach is often taken when defining 'AI': definitions generally centre on technological capabilities. In their proposal for a summer research project on artificial intelligence in 1955, McCarthy et al expressed an aim to 'find how to make machines . . . solve kinds of problems now reserved for humans'.²⁵ More recently, the OECD has defined 'AI systems' by reference to functionalities: a machine-based system is 'AI' if it 'can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments'.²⁶

At the sectoral level, function-centric definitions can be even more specific. The EU's proposed AI Act, for example, describes as 'high risk' systems which are 'intended to be used for making decisions on promotion and termination of work-related contractual relationships, for task allocation and for monitoring and evaluating performance and behaviour of persons in such relationships'.²⁷ Its proposed Platform Work Directive, meanwhile, would regulate 'automated monitoring systems which are used to monitor, supervise or evaluate . . . work performance', and 'automated decision-making systems which are used to take or support decisions that significantly affect . . . workers' working conditions'.²⁸ Canada's Directive on Automated Decision-Making, which is subject to consultation at the time of writing with a view to capturing risk and impacts on the employee experience, applies to 'any system, tool, or statistical models used to recommend or make a decision about a client'.²⁹

These definitions reflect the fact that, in the employment context, technology impacts on workers' lives when it exercises, or informs the exercise of, the *managerial prerogative*.³⁰ The 'managerial prerogative' can be understood as the employer's power to control and direct the work process.³¹ It operates at the individual employee level, as where the employer

²³ Stefano Scarpetta and Mark Pearson, 'What Happened to Jobs at High Risk of Automation?' (OECD, January 2021).

²⁴ See the changes in the European Commission's proposals to regulate technology and labour between 2018 and 2021, outlined in Aislinn Kelly-Lyth, 'The AI Act and Algorithmic Management' (2021) Comparative Labor Law & Policy Journal (Dispatch No 39) 1–2; and Anna Thomas, Stephanie Sheir and others, 'Policy Briefing: Algorithmic Impact Assessments—Building a Systematic Framework of Accountability for Algorithmic Decision Making' (Institute for the Future of Work 2021) <<https://perma.cc/3H9W-ZPHU>>.

²⁵ John McCarthy and others, 'A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence, August 31, 1955' (2006) 27 AI Magazine 12.

²⁶ OECD, *Recommendation of the Council on Artificial Intelligence*, OECD/LEGAL/0449 at 7.

²⁷ Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act) and Amending Certain Union Legislative Acts, COM (2021) 206 final (April 22, 2021) (AI Act), Annex III, ¶ 4(b).

²⁸ Proposal for a Directive of the European Parliament and of the Council on improving working conditions in platform work, COM (2021) 762 final (December 9, 2021) (Platform Work Directive) art 6(1).

²⁹ For further examples, see Institute for the Future of Work, 'Policy Tracker' (Stephanie Sheir, 2022) <<https://www.ifow.org/publications/legislation-tracker>> accessed 4 September 2022.

³⁰ The IFOW's proposal would also mandate assessments for systems used to inform 'opportunities for learning, promotion or other benefits', which could be understood to fall outside the managerial prerogative.

³¹ *Management prerogative*, *A Dictionary of Human Resource Management* (2nd edn 2008); Naeema Choudry et al (eds), *Harvey on Industrial Relations and Employment Law* ss 74, 428, 1377, 1862 (issue 295, March 2022). The prerogative is limited in various respects, including by legislation and by contract. This article uses the term 'managerial prerogative' in the legal sense.

schedules shifts, assigns tasks, and awards discretionary bonuses; but also at the organisational level, as where the employer decides to relocate a group of workers, change the size of a shift team, review production or planning processes, or otherwise reorganise the business structure.³²

Concerns about ‘algorithmic management’ tools almost always centre on tools which automate or support the exercise of the managerial prerogative in one dimension or another. Examples can range from fully automated decisions being made on matters such as shift schedules and bonus payments, to intermediate steps such as monitoring or information-gathering being carried out automatically and then informing managerial decisions made by humans. Take, for example, Microsoft’s ‘Workplace Analytics’ tool, which analyses workers’ metadata to provide ‘unprecedented behavioural insights’.³³ Employers have used these algorithmic outputs to inform organisational changes, such as modifying firm-wide reporting structures.³⁴ In 2020, Microsoft added (and subsequently removed) an additional ‘Productivity Score’ functionality, which provided managers with individual employees’ scores for attributes like ‘teamwork’.³⁵

Similar tools sold by other vendors not only inform decision-making, but fully automate some aspects of it; for example, by automatically assigning tasks based on individualised data analysis.³⁶ Institute for the Future of Work (IFOW) research has identified the growing use of ‘connected worker platforms’ which can be readily downloaded from ‘app’ stores and combine information about work and workers from a range of sources. These platforms, located at the centre of the industrial ‘internet of things’, identify potential transformations of the business model, as well as performing new and traditional management functions. The ease, reach, and speed of connected worker platforms combined with nudges made to encourage experimentation and proactive recommendations to alter organisational decisions and processes can be said to drive the ‘gigification’ of work, including management practice.³⁷

Although tools which inform the exercise of the managerial prerogative will generally satisfy definitions of ‘AI’ such as those provided by McCarthy et al and the OECD, sector-specific function-centric definitions such as those proposed by the EU and the IFOW avoid the risks of divergent interpretations of the same concept. For these reasons, the remainder of this contribution refers to ‘algorithmic risk and impact assessments’ (ARIAs) to refer to assessments of algorithmic systems of all types which form part of an ‘algorithmic management tool’.³⁸ ‘Algorithmic management tools’ can, in turn, be defined as those tools which exercise

³² See, for example, *Hollister v National Farmers Union* [1979] ICR 542; *Peter Carnie & Son Ltd v Paton* [1979] IRLR 260, EAT.

³³ The Office 365 Team, ‘Transform Your Organization with Microsoft Workplace Analytics’ (*Microsoft 365 Blog*, 7 May 2017) <<https://perma.cc/RN53-6FYW>>.

³⁴ Microsoft Corporation, ‘Case Study: Rand Mutual Assurance’ (*Workplace Insights*, 17 September 2019) <<https://perma.cc/8A5R-3JRH>>.

³⁵ Alex Hern, ‘Microsoft Productivity Score Feature Criticised as Workplace Surveillance’ *The Guardian* (26 November 2020) <<https://perma.cc/62V6-LZQH>>.

³⁶ Christl (n 2) 38–39.

³⁷ Gilbert and Thomas (n 1).

³⁸ Thomas and others (n 17); Jacob Metcalf and others, ‘Algorithmic Impact Assessments and Accountability: The Co-Construction of Impacts’, *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (ACM 2021) <<https://dl.acm.org/doi/10.1145/3442188.3445935>> accessed 24 February 2022; Margot E Kaminski and Gianclaudio Malgieri, ‘Algorithmic Impact Assessments under the GDPR: Producing Multi-Layered Explanations’ (2021) 11 *International Data Privacy Law* 125.

or inform the use of the managerial prerogative. In other words, the obligation would be technology-agnostic, but use-specific.³⁹

3. The case for algorithmic impact assessments

One approach to mitigating the potential harms of algorithmic management tools is to mandate built-in safeguards. The Platform Work Directive, for example, would, *inter alia*, require human review of algorithmic decisions in the gig economy on matters such as working time and pay.⁴⁰ While generalised obligations are crucial, identifying risk mitigations on a case-by-case basis can provide significant added value, by (i) evaluating impacts and facilitating *context-specific* responses to risks and impacts identified; (ii) striking a *balance* between generalised requirements and complete self-regulation; and (iii) ensuring that due regard to anticipated impacts and risk mitigation is *built in* from the design and development stages, through to deployment in the workplace.

3A. Context-sensitive evaluation and response

Legislative obligations are based on a high-level assessment of risk and prevention and are necessarily generalised: an absence of human oversight, for example, might result in arbitrary decisions standing uncorrected. While legislators can seek to identify and mitigate risks arising across a ‘high-risk sector’ or in all ‘high-risk uses’, some impacts and correlative mitigations will vary across contexts.⁴¹ Even where tools fulfil similar functions, specific design and deployment choices can lead to divergent results.⁴² Participatory efforts to assess and address impacts on a case-by-case basis mean that mitigations or other responses can be more effective to identify impacts—including unforeseen, invisible, and collective harms—and more responsive to them, when those who are exposed or at risk of suffering the harms are involved. For the same reasons, this approach is also more likely to ascertain potentially positive impacts on job quality, where algorithmic systems can either undermine or improve job quality, depending on the approach taken to design and deployment.⁴³

3B. Regulatory balance

Secondly, and relatedly, mandatory impact assessments strike a balance between (i) placing unduly burdensome prohibitions or safeguards on potentially useful technologies and (ii) relying on potentially ineffective voluntary self-regulation. An employer’s use of a system to track and improve supply chain efficiency, for example, may be legitimate in principle but prone to harmful misuse in practice.⁴⁴ Mitigating all risks at the macro level could result in

³⁹ See, by analogy, the AI Act and HR 6580—117th Congress (2021–2022): Algorithmic Accountability Act of 2022 (2 March 2022) (Algorithmic Accountability Act of 2022) <<https://www.congress.gov/bill/117th-congress/house-bill/6580>> (last accessed 7 December 2022), both of which would impose obligation by reference to the deployment of the technology (‘high-risk’ uses at Annex III, and ‘critical decision processes’, respectively).

⁴⁰ Platform Work Directive arts 7 and 8.

⁴¹ Alessandro Mantelero and Maria Samantha Esposito, ‘An Evidence-Based Methodology for Human Rights Impact Assessment (HRIA) in the Development of AI Data-Intensive Systems’ (2021) 41 Computer Law & Security Review 105561, 7; Kaminski and Malgieri (n 38) 139.

⁴² Kaminski and Malgieri (n 38) 139.

⁴³ Sjoerd van den Heuvel and Tanya Bondarouk, ‘The Rise (and Fall?) Of HR Analytics: A Study into the Future Application, Value, Structure, and System Support’ (2017) 4 Journal of Organizational Effectiveness: People and Performance 157; A Michael Froomkin, ‘Regulating Mass Surveillance as Privacy Pollution: Learning from Environmental Impact Statements’ (2015) 2015 University of Illinois Law Review 1713, 1746.

⁴⁴ See, by analogy, Miriam Kullmann and Aude Cefaliello, ‘The Draft Artificial Intelligence Act (AI Act): Offering False Security to Undermine Fundamental Workers’ Rights’ [2021] SSRN Electronic Journal 5 <<https://www.ssrn.com/abstract=3993100>> accessed 13 March 2022.

disproportionately harsh regulation: blunt legal restrictions can become over-determinative. Meanwhile, unguided self-regulation is likely to result in inadequate and differing standards across industry.⁴⁵

By mandating impact assessments, legislators delegate responsibility for contextual risk identification and response to those creating the potential harms: legislators ‘identif[y] the problem, provide[] suggestions of what regulators might consider adequate, and . . . task[] companies with cooperatively coming up with the solutions’.⁴⁶ This is an example of reflexive law, in which legislative action guides and directs participatory self-regulation by dictating the process rather than the outcome.⁴⁷ In this context, guiding the process includes defining the approach taken to ensure appropriate stakeholder participation; and specifying the stages at which assessments are to take place. From an organisational management perspective, collaborative governance which relies on sharing information and expertise is more likely to be effective and lead to changes in planning processes or job design by the human system designers.⁴⁸

Despite the delegation of responsibility for context-specific assessment, some red lines are plainly appropriate. Automated blacklisting of trade unionists, for example, should never be permitted.⁴⁹ Impact assessments thus play a key complementary role, operating *alongside* ‘command and control’-style regulation.⁵⁰ As Kaminski and Malgieri emphasise, algorithmic impact assessments can be a useful part of a ‘much larger system of governance’ but are not a ‘stand-alone mechanism’.⁵¹

A similar ‘collective governance’ approach can be identified in the GDPR. Data protection impact assessments (DPIAs), which must be carried out prior to ‘high risk’ data processing, have been described as a form of ‘meta-regulation’ in which the state makes corporations responsible and accountable for their own self-regulation.⁵² Similarly, the GDPR creates a right not to be subject to solely automated significant decisions,⁵³ but in the exceptional cases where significant automated decision-making is permitted, data controllers are required to proactively identify and provide ‘suitable measures to safeguard the data subject’s rights and freedoms and legitimate interests’.⁵⁴ Article 35(9) provides a further steer, requiring that the data controller ‘shall seek the views of data subjects or their representatives on the intended processing’. A bright-line distinction between the permitted and the prohibited is thus dissolved into a more nuanced procedural obligation in appropriate cases.

3C. In-built consideration of impacts

⁴⁵ Mantelero and Esposito (n 41) 3–4.

⁴⁶ Kaminski and Malgieri (n 38) 129.

⁴⁷ Gunther Teubner, ‘Substantive and Reflexive Elements in Modern Law’ (1983) 17 *Law & Society Review* 239; Ralf Rogowski and ACJM Wilthagen, ‘Reflexive Labour Law: An Introduction’ in Ralf Rogowski and ACJM Wilthagen (eds), *Reflexive Labour Law* (Kluwer Law and Taxation 1994) 7.

⁴⁸ Andrew D Selbst, ‘An Institutional View of Algorithmic Impact Assessments’ (2021) 35 *Harvard Journal of Law & Technology* 117.

⁴⁹ On the need for stronger protections, see Zoe Adams, Abi Adams-Prassl, and Jeremias Adams-Prassl, ‘Online Tribunal Judgments and the Limits of Open Justice’ (2021) 41 *Legal Studies* 53–56.

⁵⁰ Both impact assessments and the concept of ‘command and control’ regulation come from environmental regulation: see *Command and control*, *Jowitt’s Dictionary of English Law* (5th edn 2019) and Froomkin (n 43).

⁵¹ Kaminski and Malgieri (n 38) 138.

⁵² Reuben Binns, ‘Data Protection Impact Assessments: A Meta-Regulatory Approach’ (2017) 7 *International Data Privacy Law* 22, 23; cf Kaminski and Malgieri (n 38) 131.

⁵³ Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (GDPR) art 22(3).

⁵⁴ Kaminski and Malgieri (n 38) 127–128.

Finally, impact assessments embed consideration of potential harm into design and deployment.⁵⁵ Traditional ‘command and control’ legislation, which specifies what regulatees can and cannot do, is prone to encourage compliance with the letter, rather than the spirit, of the law. Laws which prohibit unjustified dismissal for employees with a minimum length of service, for example, have limited power to inspire broader shifts beyond those statutory minima.⁵⁶ Moreover, where legislation is breached, enforcement or amendment is necessarily reactive. By contrast, *ex ante* impact assessments can entail holistic consideration of harms from the outset, and therefore encourage a shift away from minimum compliance standards, signalling a move towards design, development, and deployment of algorithmic systems to promote ‘good’ outcomes. By making the obligation procedural, the primary question is whether reasonable efforts have been made, and whether the approach taken is reasonable in the circumstances, rather than whether all regulatory minima have been met.

4. Concretising the proposal

By targeting tools which affect the exercise of the managerial prerogative, it is possible to create a robust definition of the circumstances in which an impact assessment should be required. This is the first step towards legislative clarity. This section proceeds from that first step by describing what an algorithmic impact assessment would look like *in practice* in the employment context.

It does so by addressing four key questions:

1. *At what stage* should the assessment be carried out?
2. *By whom* should it be conducted?
3. What should it cover *substantively*?
4. What should the *procedural steps* entail? This requires consideration of *who* the assessment should be carried out with, and what *transparency* requirements should apply.

The IFOW has proposed a detailed approach for algorithmic impact assessments, which can be conceptualised as a systematic framework for accountability based on an overarching new, positive duty in the public interest on employers (and some others) to undertake, disclose, and act appropriately upon assessments once a basic threshold is met. This model, which invites the development of sector and other context-specific protocols, guidance, and standardised techniques over time, provides a touchstone for our discussion throughout the section.⁵⁷ Here, we propose use of the Good Work Charter as a framework to underpin an effective, hybrid ARIA in the context of algorithmic management at work.

4A. At what stage and by whom?

⁵⁵ Such requirements incentivise early consideration of impacts and form a basis for conversations with the regulator about appropriate risk mitigation: Froomkin (n 43) 1756.

⁵⁶ Of course, such changes can have broader economic impacts which affect workers: see, for example, Seamus McGuinness and Paul Redmond, ‘The Impact of a Minimum-Wage Increase on Temporary-Contract Workers’ (2019) 40 Fiscal Studies 149.

⁵⁷ Thomas, Sheir and others (n 24). This proposal builds on All-Party Parliamentary Group on the Future of Work, ‘The New Frontier: Artificial Intelligence at Work’ (2021) 12 <www.futureworkappg.org.uk> accessed 12 November 2021, and was drafted by Anna Thomas with input from Helen Mountfield QC, Abigail Gilbert, Stephanie Sheir, and Reuben Binns.

The first question is critical for answering the others. One of the key obstacles to ensuring algorithmic accountability is that responsibility for automated decision-making is diffuse, often spanning multiple organisations—as where a model is trained on data provided by a data broker and then sold to an employer for use.⁵⁸ In recognition of this difficulty, the ARIA should be undertaken by all actors who are developing or deploying algorithmic systems which may be used in the employment context, as well as other key actors across the design cycle and supply chain.⁵⁹ This should be combined with duties for employers to record relevant documentation and to cooperate with vendors where tools have been externally procured.⁶⁰ For example, contracting parties must ensure that the contract provides the assessing party (such as the employer) with the documentation required to produce an algorithmic impact assessment.

There are at least three temporal points at which the most critical human decisions are made, and at which an assessment of impacts should take place: design, development, and deployment.⁶¹ Thus far, regulatory instruments and proposals have generally sought to tackle either design or deployment, overlooking both the development process in between and any interaction between the stages, such as making runtime adjustments. The proposed AI Act, for example, primarily places obligations, including for compliance assessments, on the ‘providers’ (*viz* vendors) of AI systems. By contrast, most obligations under the GDPR, including data protection impact assessments, fall to the ‘data controller’: in this case, the employer.⁶² The proposed Platform Work Directive also includes an obligation akin to an algorithmic impact assessment, in the shape of a requirement to ‘regularly monitor and evaluate the impact’ of algorithmic management tools—but this obligation would apply only to the employer, with no correlative duty imposed on vendors of such tools.⁶³

Despite the apparent inconsistency in the regulatory environment, each of these approaches has its merits. Given that one of the key benefits of requiring an impact assessment is that consideration of harm is embedded into design and development, it makes sense that the obligation should be imposed on designers and developers.⁶⁴ On the other hand, there is a significant degree of abstraction between the developer and the impacts, which translates into a tension between design-stage consideration and context-sensitivity—the latter being another key benefit of impact assessments.⁶⁵ To resolve this tension, tools must *also* be assessed at the deployment stage by the employer and any agents involved in supporting or monitoring use of the system. For larger employers, when impacts or risks are likely to be significant, it would seem preferable to undertake the assessment at all key human decision-making points: project planning, problem formulation, data collection and procurement, feature engineering and

⁵⁸ Thomas and others (n 17) 24–25.

⁵⁹ Thomas and others (n 17).

⁶⁰ The IFOW suggests that such duties should be based, at minimum, on the principles and practices of disclosure in the High Court: Thomas, Sheir and others (n 24).

⁶¹ See similarly Kaminski and Malgieri (n 38) 139; Lorna McGregor, Daragh Murray, and Vivian Ng, ‘International Human Rights Law as a Framework for Algorithmic Accountability’ (2019) 68 *International & Comparative Law Quarterly* 309, 330; Lilian Edwards, ‘Expert Opinion: Regulating AI in Europe—Four Problems and Four Solutions’ (Ada Lovelace Institute 2022) 20 <<https://www.adalovelaceinstitute.org/report/regulating-ai-in-europe/>> accessed 27 April 2022.

⁶² For discussion of this contrast, see Kelly-Lyth (n 24), 5–6.

⁶³ Platform Work Directive, Art 7(1).

⁶⁴ This is particularly true given that many employers outsource their algorithmic management tools, rather than developing them in-house: see, generally, Christl (n 2).

⁶⁵ Martin Ebers, ‘Standardizing AI—The Case of the European Commission’s Proposal for an Artificial Intelligence Act’ in Larry A Di Matteo, Cristina Poncibò, and Michel Cannarsa (eds), *The Cambridge Handbook of Artificial Intelligence: Global Perspectives on Law and Ethics* (Cambridge University Press 2022) s 3.4, noting that setting standards for AI at the macro level is ‘difficult because these systems are used in different industries and sectors’, with impacts which ‘depend[] strongly on their embedding in the respective context.’

variable selection, model selection, training, validation and testing, implementation, and making runtime adjustments.⁶⁶ While there may therefore be multiple impact assessment ‘triggers’, a clear legislative approach would require assessments to build on one another, rather than being carried out in isolation. In all cases, the first assessment must be carried out early on, before the technology is developed or the uses are concretised respectively.⁶⁷

4B. Which impacts?

The next question is substantive: what should an algorithmic impact assessment cover? This raises two sub-questions: what *type* of impacts should be assessed, and *on whom*? Views on both points diverge. Kaminski and Malgieri, for example, suggest that ‘a model [algorithmic impact assessment] could take advantage of the fact that it is conducted on a system-wide level to search for, and mitigate, *social* harms that go *beyond impacted individuals*’.⁶⁸ Charlesworth similarly suggests that focusing impact assessments on ‘legally protected individual or group rights’ can result in a failure to take a ‘holistic view of potential structural inequalities’, meaning that the ‘unfairness of AI outcomes tends to go unchallenged and unaddressed’.⁶⁹

Mantelero and Esposito, by contrast, argue that both legislatures and NGOs consider human rights to be ‘the core of future AI regulation’, and that while ethical considerations are important, they should not form part of the algorithmic impact assessment.⁷⁰ The rationale for focusing the assessment on human rights impacts is that they are ‘clearer, better defined, and [more] stable’ than ‘ethical values’.⁷¹ Elsewhere, Mantelero highlights the limitations of a legal approach in isolation, and his model (a ‘Human Rights Ethical and Social Impact Assessment’) aims to combine ‘the universality of human rights with the local dimension of societal values’.⁷² McGregor, Murray and Ng also suggest that (international) human rights law provides a ‘framework capable of capturing the full algorithmic life cycle from conception to deployment’, since it offers a well-defined and comprehensive conception of harm.⁷³ From an organisational management perspective, a human rights lens can be seen as an important step towards embedding a human-centred approach, developing a culture of decent treatment and preventing and mitigating harm to people, starting with employees but not limited to them.⁷⁴

The primary argument against mandating a wide-ranging ethical assessment is the need for clarity: an obligation to consider a defined list of impacts on a cognate group is more

⁶⁶ Thomas and others (n 17).

⁶⁷ On timing, see Riki Therivel, ‘Literature Review of Impact Assessment in Governments’ (The Scottish Government 2021) 37.

⁶⁸ Kaminski and Malgieri (n 38) 139, emphasis supplied.

⁶⁹ Andrew Charlesworth, ‘Regulating Algorithmic Assemblages: Looking beyond Corporatist AI Ethics’, *Data-Driven Personalisation in Markets, Politics and Law* (Cambridge University Press 2021) 253–255, discussing the Assessment List produced by the EU High-Level Expert Group on Artificial Intelligence.

⁷⁰ Mantelero and Esposito (n 41) 6–7. For human rights as a framework for regulating AI more broadly, see eg Alessandro Mantelero, ‘Regulating AI within the Human Rights Framework: A Roadmapping Methodology’ in Philip Czech et al (eds), *European Yearbook on Human Rights* (Intersentia 2020); Lorna McGregor, Daragh Murray, and Vivian Ng, ‘International Human Rights Law as a Framework for Algorithmic Accountability’ (2019) 68 *International and Comparative Law Quarterly* 309–343; Council of Europe Ad Hoc Committee on Artificial Intelligence (CAHAI), ‘Possible elements of a legal framework on artificial intelligence, based on the Council of Europe’s standards on human rights, democracy and the rule of law’ (3 December 2021, CAHAI(2021)09rev); David Leslie and others, ‘Artificial intelligence, human rights, democracy and the rule of law: a primer’ (2021, The Council of Europe and The Alan Turing Institute).

⁷¹ Mantelero and Esposito (n 41) 5.

⁷² Mantelero (n 19) 755.

⁷³ McGregor, Murray, and Ng (n 61) 313. Note that references to ‘impacts’ in the context of impact assessments are generally proxies for harms: Metcalf and others (n 38).

⁷⁴ For example, see Caroline Rees, ‘To address inequality, companies should put human rights at the core of how they do business’ (WBCSD, 23 June 2022) <<https://perma.cc/N95J-9GKG>>.

obviously concrete than a vague obligation to consider the ‘ethical’ implications of a technology.⁷⁵ However, arguments against consideration of ‘ethics’ risk conflating ethics-washing through unenforced ‘codes of conduct’ with recent work on ethics-based assurance as a structured process which allows employers and employees to assess behaviour for consistency against an agreed normative framework.⁷⁶ Properly done, and applied in a context-sensitive way, an ethical impact assessment should complement and interlock with legal and social impact assessments in the workplace.

The Good Work Charter sets out fundamental principles that define ‘good’ or decent work for policy orientation and practical application, and was developed by a cross-disciplinary group specifically for the employment context. The Charter operates as a synthesis of both legal principles, rights, and obligations,⁷⁷ and ethical principles as they apply to the workplace, with particular regard to the use of artificial intelligence and algorithmic systems. As such, we propose use of the Charter to enable a hybrid ARIA that combines legal, ethical, and social dimensions, evaluated against the Good Work principles. This approach is consistent with, and builds on, the public acknowledgement by the former Chair of the House of Lords Select Committee on Artificial Intelligence and Vice-Chair of the UK’s All Party Parliamentary Group (APPG) on the Future of Work that the Good Work Charter can be used as ‘checklist to consider potential impacts [of algorithmic systems] on work and workers’.⁷⁸ Moreover, since the Charter synthesises international and European legal principles, it provides an appropriate framework for approaches in a wide range of jurisdictions, including the EU.

In light of the discourse around human rights as a framing for algorithmic impact assessments, it is particularly notable that the Good Work Charter reflects both the European Social Charter and the Charter of Fundamental Rights of the European Union.⁷⁹ Three examples are demonstrative:

- The Good Worker Charter establishes the principle of ‘access’, specifying that ‘everyone should have access to good work’.⁸⁰ The European Social Charter similarly establishes that ‘everyone shall have the opportunity to earn his living in an occupation freely entered upon’,⁸¹ while the Charter of Fundamental Rights establishes that ‘[e]veryone has the right to engage in work and to pursue a freely chosen or accepted occupation’.⁸²

⁷⁵ Martin Ebers and others, ‘The European Commission’s Proposal for an Artificial Intelligence Act—A Critical Assessment by Members of the Robotics and AI Law Society (RAILS)’ (2021) 4 J 589, 593, arguing that ‘more detailed classification of risk’ is necessary within the AI in order ‘for the industry to perform the self-assessment of risks associated with their products’.

⁷⁶ Christopher Burr and David Leslie, ‘Ethical Assurance: A Practical Approach to the Responsible Design, Development, and Deployment of Data-Driven Technologies’ [2022] AI and Ethics <<https://doi.org/10.1007/s43681-022-00178-0>> accessed 28 June 2022.

⁷⁷ The legal bases for the Good Work Charter are available at <<https://ifow.webflow.io/publications/the-ifow-good-work-charter>> accessed 29 August 2022.

⁷⁸ All-Party Parliamentary Group on the Future of Work (n 57) 12.

⁷⁹ Good Work Charter (GWC) principle 1 and Charter of Fundamental Rights (CFR) art 15; GWC principle 3 and CFR art 31; GWC principle 4 and CFR Title III; GWC principle 5 and CFR Title I (albeit that the GWC may have a more expansive interpretation of ‘dignity’); GWC principle 6 and CFR arts 10–16 (albeit that the CRF does not impose any obligation to ‘promote autonomy’); GWC principle 7 and CFR art 31; GWC principle 9 and CFR arts 27 and 28; GWC principle 10 and CFR art 14 (although the CFR does not provide for ‘lifelong learning and career guidance’). The only unaddressed GWC principles are principle 2 (fair pay) and 8 (‘everyone should have access to institutions and people who can represent their interests’).

⁸⁰ GWC principle 1.

⁸¹ Council of Europe, European Social Charter, 18 October 1961, ETS 35 (ESC), Part I and Part II, art 1.

⁸² CFR art 15(1). The limitation here is the absence of a requirement for the work to be ‘good’. For discussion of this limitation, see Simon Deakin, ‘Article 15—Freedom to Choose an Occupation and Right to Engage in Work’

- The Good Work Charter holds that ‘[e]veryone should work on fair conditions set out on fair terms’.⁸³ The European Social Charter establishes that ‘all workers should have the right to just conditions of work’,⁸⁴ and the Charter of Fundamental Rights provides that every worker be granted the ‘right to working conditions which respect his or her health, safety and dignity’.⁸⁵ The ethical imperative for fair and decent work is also recognised in the UNESCO recommendation on Ethical AI Regulation,⁸⁶ and the imperative to design, develop, and deploy AI systems in line with social and political rights, including the right to just conditions of work, is recognised by the Council of Europe’s Ad Hoc Committee on Artificial Intelligence (CAHAI).⁸⁷

- The Good Worker Charter posits that ‘[e]veryone should be able to take part in determining and improving working conditions’.⁸⁸ The European Social Charter establishes that ‘workers have the right to take part in the determination and improvement of the working conditions and working environment in the undertaking’.⁸⁹ The Charter of Fundamental Rights provides that ‘[w]orkers or their representatives must . . . be guaranteed information and consultation’,⁹⁰ as well as ‘the right to negotiate and conclude collective agreements’.⁹¹ Similarly, the Human Rights, Democracy and the Rule of Law Assurance Framework for AI co-produced by the UK’s Turing Institute and Council of Europe highlights the ethical as well as legal imperatives for active participation throughout all stages of the AI life cycle.⁹²

This is not to say that the role of the Good Work Charter is equivalent to that of comprehensive international rights instruments: the Charter of Fundamental Rights expressly includes, for example, the right to protection of personal data,⁹³ and its breadth also means that it can go some way to encouraging consideration of structural issues.⁹⁴ The right to family and professional life entails specific considerations of impacts on work-life balance,⁹⁵ for example, while rights to information, consultation, and collective bargaining all link into participatory democracy and workers’ voice.⁹⁶ Here, the Good Work Charter offers an accessible, high level synthesis under the principles of ‘autonomy,’ ‘dignity,’ and ‘participation’. It offers a useful access point, but should not be seen as a substitute for detailed legal analysis in a specific case. The Charter provides a framework which is particularly well-equipped to deal with the impacts of algorithmic management: ‘dignity’ and ‘autonomy’ provide a counterbalance to loss of

in Filip Dorssemont and others (eds), *The Charter of Fundamental Rights of the European Union and the Employment Relation* (Hart Publishing 2019) 346–347.

⁸³ GWC principle 3.

⁸⁴ ESC Part I and Part II art 2.

⁸⁵ CFR art 31(1).

⁸⁶ UNESCO, ‘Recommendation on the ethics of artificial intelligence’, SHS/BIO/REC-AIETHICS/2021, Policy Area 10, Economy and Labour.

⁸⁷ CAHAI (n 70); Isaac Ben-Israel and others, *Towards Regulation of AI Systems: Global perspectives on the development of a legal framework on Artificial Intelligence (AI) systems based on the Council of Europe’s standards on human rights, democracy and the rule of law* (Council of Europe 2020), para 40.

⁸⁸ GWC principle 9.

⁸⁹ ESC, Additional Protocol of 1988 art 3.

⁹⁰ CFR art 27.

⁹¹ CFR art 28.

⁹² Leslie and others (n 70).

⁹³ CFR art 8. Note that two principles from the GWC are not reflected in the CFR, however (n 79).

⁹⁴ The CFR includes rights such as the ‘right to vote and to stand as a candidate at elections to the European Parliament’, for example (art 39). A useful reference point for identifying rights relevant to the employment context is Filip Dorssemont and others (eds), *The Charter of Fundamental Rights of the European Union and the Employment Relation* (Hart Publishing 2019).

⁹⁵ CFR art 33, Csilla Kollonay Lehoczky and Barbara Kresal, ‘Article 33—Family and Professional Life’ in Filip Dorssemont and others (n 94).

⁹⁶ CFR arts 27 and 28.

agency, for example, while ‘participation’ requires information asymmetries to be addressed.⁹⁷ Similarly, ‘autonomy’ demands an ability to make informed choices, thus suggesting a break on the psychological chilling effects on behaviour which have been observed in the context of pervasive monitoring by algorithmic systems in the workplace.⁹⁸

Primary legislation to mandate impact assessments should ideally specify the catalogue of impacts to be considered even (or perhaps especially) where the full remit and extent of these impacts cannot be anticipated in advance of undertaking the assessment itself. Given that the Good Work Charter is a synthesis of relevant provisions from various rights and ethics instruments, it provides a helpful starting point for this exercise.⁹⁹ Secondary legislation, codes and guidance at a regulator and sector level should follow. Legislation could also specify some minimum considerations for each principle—by requiring, for example, that the assessment of impacts on ‘autonomy’ include mandatory consideration of impacts on choice of work assignments and working hours; that any automated calculation of pay be disclosed with variables considered; and that larger employers undertake an equality impact assessment, drawing on the public sector model.¹⁰⁰

4C. Defining the procedure

The fourth and final question is procedural: impact assessments can quickly become a box-ticking exercise if inadequately defined. Different approaches are clearly required for assessments carried out by the developer and the deployer, but some common themes can be identified in the literature—particularly participatory identification of risks and responses (which must be meaningful and imply some level of transparency) and ongoing review of implementation and outcomes (at least periodically).¹⁰¹ Existing research also includes a variety of model approaches.

The IFOW’s proposal, for example, describes an assessment structured around four distinct stages: identification of individuals and communities who might be impacted; an *ex ante* risk and impact analysis; the taking of appropriate response action; and continuous evaluation to ensure assessment and appropriate action is ongoing.¹⁰² The content of the assessment, like the steps taken in response to it, should be reasonable and proportionate in the circumstances, respecting and balancing the rights and interests of employee and employer, and considering the size, resources, and capabilities of the business alongside the severity and proximity of the harm or other adverse impacts.

⁹⁷ Adams-Prassl and others (n 7).

⁹⁸ Nadezhda Purtova, ‘Default Entitlements in Personal Data in the Proposed Regulation: Informational Self-Determination off the Table ... and Back on Again?’ (2014) 30 Computer Law & Security Review 6. See also CAHAI (n 70).

⁹⁹ For a proposal on incorporating the principles into statute, see HL Deb 10 March 2021, vol 810, col 693GC et seq. The Charter provisions have already been the subject of wide-ranging consultation and debate: see Future of Work Commission, ‘Report of the Future of Work Commission’ (2017) 7 <<https://perma.cc/EKZ6-3EL5>>.

¹⁰⁰ Anna Thomas and others, ‘Artificial Intelligence in Hiring: Assessing Impacts on Equality’ (Institute for the Future of Work 2020). The Platform Work Directive art 6(1)(b) provides a list of this type in relation to ‘working conditions’, the impact on which must be monitored per art 7(1).

¹⁰¹ See, for example, McGregor, Murray and Ng (n 61) 320–324; Therivel (n 67) 37–38, 40. There is strong worker support for legally mandated consultation: Trades Union Congress, ‘Technology Managing People: The Worker Experience’ (2020) 41.

¹⁰² Sheir and others (n 24). Anna Thomas drafted the model in Annex I with input by Helen Mountfield, Sa’ad Hossain, Abby Gilbert, Stephanie Sheir, David Leslie, and Reuben Binns.

The framework for ethical assurance proposed by Burr and Leslie similarly points to four main procedural ‘buckets’, which broadly equate with the four stages of the APPG and IFOW models and should be seen across the life cycle of the algorithmic system.¹⁰³

1. Stakeholder identification and analysis, which can serve to guide the appropriate level and scope of engagement activities.
2. Evaluation of risks and harms on this group, combining impact assessments required by law (such as data protection impact assessments, discussed below) with best practice (such as equality, human rights, and bias assessments) and exploring impacts against a normative framework (in this case, the Charter). The evaluation should allow for reflection, challenge, and the identification of new and unforeseen risks.
3. Effective co-determination of reasonable and proportionate adjustments or other steps in response to the assessment (such as revisiting the variable selection or training).
4. Finally, the establishment of a process to enable contextually informed, ongoing monitoring of both known and unforeseen impacts.

An empirical analysis of the approaches taken by data protection authorities led Mantelero and Esposito to develop a model methodology for design-stage impact assessments,¹⁰⁴ which assesses impacts by reference to likelihood and severity and builds in stakeholder engagement at the scoping stage. Other proposals have been designed for use in the public sector or for public-private intersection.¹⁰⁵

In all cases, stakeholder consultation is key. In the work context, it is critical that such consultation has a collective dimension. There are two reasons for this. First, many algorithmic management tools are inherently relational:¹⁰⁶ standards are set by reference to average or best performers. Algorithmic systems work by making predictions about groups based on common features, and workers are evaluated against their colleagues. Focusing on the individual might mean missing collective harms. Secondly, collective voice serves to temper the inequality of bargaining power inherent in the employment relationship: the collective is stronger than the individual.¹⁰⁷ This is particularly important in the context of growing information asymmetries at work and concentration of information assets more widely.

Newman, for example, highlights that worker analytics tools can tend to result in ‘least attached workers’ receiving preferential treatment. The tools enable employers to better identify the workers who are most likely to exit the firm, so that retention efforts can be better targeted—to the detriment of ‘older and less marketable’ workers.¹⁰⁸ Newman suggests that a ‘collective-action approach’ is necessary to challenge such problems.

Where design and development take place outside of the firm, consultation could take place via worker representative bodies, such as the European Trade Union Confederation (in the EU), or the Trades Union Congress (in the UK). The feasibility of such consultation is well-

¹⁰³ Burr and Leslie (n 76).

¹⁰⁴ Mantelero and Esposito (n 41) 17–21.

¹⁰⁵ See, for example, Dillon Reisman and others, ‘Algorithmic Impact Assessments: A Practical Framework for Public Agency Accountability’ (AI Now 2018); Lara Groves, ‘Algorithmic Impact Assessment: A Case Study in Healthcare’ (Ada Lovelace Institute 2022).

¹⁰⁶ See similarly, in relation to data governance, Salomé Viljoen, ‘A Relational Theory of Data Governance’ (2021) 131 *The Yale Law Journal* 370.

¹⁰⁷ See, generally, Richard B Freeman and James L Medoff, *What Do Unions Do?* (Basic Books 1984).

¹⁰⁸ Nathan Newman, ‘Reengineering Workplace Bargaining: How Big Data Drives Lower Wages and How Reframing Labor Law Can Restore Information Equality in the Workplace’ 85 *University of Cincinnati Law Review* 68, 703–704, 716–721.

established: the ETUC has been involved in the work of standardisation organisations at the Union level since 2017, and their early involvement in standard-setting activities has led to positive results.¹⁰⁹

Meanwhile, procedural requirements for firm-level consultation processes should be drawn from existing approaches in industrial relations. In Europe, these include collective bargaining, works councils, and joint consultative committees.¹¹⁰ In all instances, the worker voice must have some force behind it. Such force might stem from the possibility of collective action (in the form of strikes, for example) or from co-determination rights (as where works councils can block certain corporate decisions). By contrast, light-touch duties to share information and consider stakeholder views may be ineffective.¹¹¹ Mandated algorithmic impact assessments should therefore articulate the information to be documented and shared, and by whom; and should require the establishment of a dedicated process for stakeholder participation and the period for wider consultation.¹¹²

Transparency is a crucial component for effective consultation, starting with the existence, purpose, and remit of the algorithmic system and its anticipated outputs, as well as the ARIA itself: blinkered input will be of little use. Transparency also has independent value, making it more likely that workers are able to understand algorithmic decisions made about them, so long as they have the capacity, time, and space to interrogate and challenge these decisions and the socio-technical context in which they are made.¹¹³

Such transparency is urgently required. In one recent study, 50 percent of UK-based employees reported that algorithmic management technologies might be in use in their workplaces *without their knowledge*, and only 21 percent felt that they would be able to effectively challenge algorithmically-made decisions.¹¹⁴ In another survey, 52 percent of workers were ‘not at all confident’ that they knew ‘why and for what purposes’ their employers used data collected about them, and 67 percent were ‘not at all confident’ that they understood how their data was used to assess their performance.¹¹⁵ Current requirements and forums for information-sharing, consultation, and participation have been found to lack bite in the context of the use of algorithmic systems at work, and the ARIA provides one example of the new processes which we anticipate will be developed over time in response to this shortfall.¹¹⁶

Crucially, transparency cannot be mandated only *during* the assessment itself: access to information about the options for adjustment, the balancing exercise involved as trade-offs are made, and the outcomes of the assessment and likely effectiveness of mitigations, are

¹⁰⁹ Directorate-General for Internal Market, Industry, Entrepreneurship and SME and European Commission, ‘Study on the Implementation of the Regulation (EU) No. 1025/2012 (Article 24)’ (Publications Office of the EU 2021) 86–88 <<https://data.europa.eu/doi/10.2873/504681>> accessed 27 April 2022.

¹¹⁰ For discussion, see Adrian Wilkinson and others (eds), *Handbook of Research on Employee Voice* (Edward Elgar 2020) chs 14, 15, and 16.

¹¹¹ See, for example, the limited impact of the European Information and Consultation Directive: Tony Dobbins and others, ‘Employment Regulation, Game Theory and Weak Employee Voice in Liberal Economies’ (2017) 156 *International Labour Review* 395.

¹¹² The IFOW model suggests the provision of a statement for the period of consultation and process for stakeholder participation.

¹¹³ Trades Union Congress (n 101) 37.

¹¹⁴ *ibid* 38.

¹¹⁵ Thomas and others (n 17) 24–26.

¹¹⁶ See IFOW proposal as adopted by Community: Community and IFOW, ‘Technology Agreements: A Partnership Approach to Use of Technology at Work’ (2021) <<https://perma.cc/5TA8-2RZB>>.

equally important.¹¹⁷ Transparency is also required about the nature and process of participation in the substantive assessment.¹¹⁸ The evidence is that effective impact assessments are iterative processes rather than one-time obligations, and oversight has been identified as a key element in this framework.¹¹⁹ Publication of impact assessments and follow-up materials, in redacted form where absolutely necessary, should therefore be a non-negotiable procedural obligation.¹²⁰ Where such transparency is not realised, the consequences can be severe, as shown below in relation to the GDPR-mandated data protection impact assessments (DPIAs).

By answering the four questions posed at the beginning of this section, a concrete proposal for an ARIA has emerged:

1. *Temporal points:* Unless the tool is developed in-house for a designated purpose, a minimum of three impact assessments should be carried out and then updated on an ongoing basis: at the design stage, the development stage, and the deployment stage. For larger employers, assessments at all key seven stages of human decision-making and continual monitoring is best practice.
2. *Duty bearers:* These assessments should be carried out by the algorithm designer, developer, and employer respectively.
3. *Contents:* The assessments should consider risks to, and impacts on, workers' rights, freedoms, and other interests at an individual and group level. The Good Work Charter should be used as a framework for this exercise, alongside concrete and non-exhaustive particularised examples for each of the ten Charter principles. Legal, ethical, and social impacts identified should be recorded in the ARIA. The assessments should also document the proportionate technical and non-technical steps taken as a result of the evaluation to address or mitigate any harms identified, and should enable consideration of potentially 'good' impacts on work conditions and quality. For consistency against a baseline, this exercise can also be undertaken by reference to the Good Work Charter.
4. *Process:* The procedure should require mandatory recording of documentation, involvement of stakeholder workers and representatives, publication of outcomes, and published reviews at defined intervals. Procedures for consultation and participation should build on existing work on public sector algorithm impact assessments (in the case of vendors), or traditional labour approaches (in the case of employers).¹²¹

5. A need for new law

In the UK and EU, the existing DPIA obligation under the GDPR might appear to go a long way towards realising this proposal. Indeed, Kaminski and Malgieri analyse the DPIA *as an*

¹¹⁷ See, for example, Public Authority Algorithm HL Bill clause 4(4), which would require impact assessments carried out by public bodies to be 'published in accessible 20 format within 30 days of the results being known'; and Sheir and others (n 24) 23, which sets out clear disclosure requirements.

¹¹⁸ Gilbert and others, 'Worker involvement in auditing' (IFOW, forthcoming 2022).

¹¹⁹ Therivel (n 67) 36; Kaminski and Malgieri (n 38) 133.

¹²⁰ The EU's proposed AI Act would enhance transparency by requiring certain AI systems to be registered in a public database, although the scope of the publishable information remains limited: see AI Act art 60. In the US, the Algorithmic Accountability Act of 2022 s 6 would require the Federal Trade Commission to publish a report providing information on the use of tools for automating certain high-risk decisions, drawn from information mandatorily provided by regulated entities.

¹²¹ IFOW's guidance on Good Work Audits (forthcoming) with the support of the Information Commissioner's Office.

‘algorithmic impact assessment’.¹²² The relevant provision, Article 35(1) of the GDPR, provides (in part) that:¹²³

Where a type of processing in particular using new technologies, and taking into account the nature, scope, context and purposes of the processing, is likely to result in a high risk to the rights and freedoms of natural persons, the controller shall, prior to the processing, carry out an assessment of the impact of the envisaged processing operations on the protection of personal data.

The obligation is therefore triggered by a risk to rights and freedoms (not just to data protection rights), and Article 35(7) clarifies that the DPIA must include ‘an assessment of the risks to the rights and freedoms of data subjects’ as well as the ‘measures envisaged to address’ these risks. The term ‘rights and freedoms of natural persons’ can be understood to encompass the rights guaranteed by the Charter of Fundamental Rights.¹²⁴ Recital 75 further explains:¹²⁵

The risk to the rights and freedoms of natural persons . . . may result . . . in particular: where the processing may give rise to discrimination . . . or any other significant economic or social disadvantage; . . . where personal aspects are evaluated, in particular analysing or predicting aspects concerning performance at work . . .

The DPIA obligation will therefore apply in many of the circumstances with which this contribution is concerned, and should essentially cover the same or similar impacts as those identified at section 4. Indeed, the majority of national data protection agencies in EU Member States have expressly specified that even employee monitoring requires a DPIA,¹²⁶ and EU-level official guidance also gives systematic monitoring of employees’ activities as example of a situation requiring a DPIA.¹²⁷

The material scope of the DPIA is not, however, equivalent to that of the proposed ARIA. Although data protection can operate as a gateway to access other rights and freedoms relevant to algorithmic management, the data protection regime rests on assumed human ability to control and manage information (human ‘sovereignty’ over data¹²⁸), an idea that has been challenged by the latest wave of algorithmic management tools. Further, there will be some situations in which tools used to inform the managerial prerogative do not require a DPIA simply because they do not process personal data. For example, where non-personal (anonymised) data on supply chain efficiency leads to an entire team being relocated,¹²⁹ no

¹²² Kaminski and Malgieri (n 38). Emre Kazim and Adriano Koshiyama, ‘The Interrelation between Data and AI Ethics in the Context of Impact Assessments’ (2021) 1 AI and Ethics 219, identify that DPIAs are a ‘natural place to start’ for AI impact assessments. We note, at the time of writing, that the Data Protection and Digital Information Bill proposes to change this DPIA obligation.

¹²³ Emphasis supplied.

¹²⁴ This can be inferred from the references to the Charter at GDPR recitals 1 and 4; and see, for example, Case C-645/19 *Facebook Ireland and others v Gegevensbeschermingsautoriteit*, ECLI:EU:C:2021:483, para 46, noting that the GDPR ‘respects all fundamental rights and observes the freedoms and principles recognised in the Charter’, and proceeding to interpret its provisions on that basis.

¹²⁵ Emphasis supplied.

¹²⁶ Fieldfisher, ‘Summary of DPIA “blacklists”’ (January 2020) <<https://perma.cc/U9H5-N4RW>>, summarising Member States’ implementation of GDPR art 35(4).

¹²⁷ Article 29 Working Party, ‘Guidelines on Data Protection Impact Assessment (DPIA) and Determining Whether Processing Is “Likely to Result in a High Risk” for the Purposes of Regulation 2016/679’ (adopted on 4 April 2017, as last revised and adopted on 4 October 2017, 17/EN WP 248 rev.01) 11.

¹²⁸ Patrik Hummel and others, ‘Data Sovereignty: A Review’ (2021) 8 Big Data & Society 2053951720982012.

¹²⁹ The GDPR applies only to information relating to an identified or identifiable natural person (‘personal data’), and not ‘anonymous information’: GDPR art 4(1) and recital 26.

DPIA obligation arises—even though the team move is an exercise of the managerial prerogative with impacts on workers.

Returning to the four key questions, some further gaps can be identified. First, section 4 argued that to be effective, impact assessments must be carried out during at least three stages. The DPIA obligation will generally only arise at the second stage, because it applies to the ‘data controller’, defined as the party who determines the ‘purposes and means’ of data processing.¹³⁰ In almost all cases, the employer will decide why and how to use algorithmic management tools. If the developer processes personal data for the employer, it may be obliged to ‘assist’ the employer with the DPIA,¹³¹ but such assistance will only be relevant after the design phase, thus losing the value which comes from the embedding harm mitigation at the design stage.¹³²

The first and second elements of the proposal—at *what stage* and *by whom*—are therefore only partially addressed by the DPIA obligation. Moreover, the proposed AI Act would not fill in this gap: the proposal would mandate ‘conformity assessments’, but these would be by reference to technical standards, not human rights (or other legal, ethical, or social impacts), and would not meet the consultation or transparency standards identified at section 4.¹³³

The transparency and consultation points both relate to the fourth element of the algorithmic impact assessment proposal: procedure. Even though the harms assessed in the DPIA are broadly similar to those that would be assessed in the ARIA, there is significant procedural divergence. Most strikingly, the DPIA obligation contains no transparency requirements whatsoever. Official EU-level guidance suggests that DPIA publication could help foster trust, adding that publication could just consist of ‘a statement that a DPIA has been carried out’¹³⁴—but there is no obligation for data controllers to do even this.

Kaminski and Malgieri suggest that the absence of any DPIA transparency obligation is its ‘biggest shortcoming’,¹³⁵ and argue that the absence of DPIA transparency has undermined the effectiveness of related obligations. For example, the GDPR requires data controllers to consult relevant national supervisory authorities about proposed processing where the DPIA reveals unmitigable high risks.¹³⁶ While this obligation theoretically provides an opportunity to access expert advice on risk mitigation and acceptability, Kaminski and Malgieri suggest that in reality it means that companies can ‘decide . . . whether [they] should be subject to regulatory oversight’.¹³⁷ The authors’ concerns are valid: as of December 2021, the UK data protection authority had been approached only twice by data controllers proposing high-risk processing in an employment context.¹³⁸ Research reveals that high-risk data processing was indeed being carried out by various UK-based employers during the same period,¹³⁹ but since DPIAs are very rarely published, there is little scope for assessing the mismatch.

¹³⁰ GDPR art 4(7).

¹³¹ GDPR art 28(3)(f). Occasionally the vendor may have sufficient input into the ‘purposes and means’ to be a ‘joint processor’, in which case the obligation to carry out the DPIA is a joint one.

¹³² Edwards (n 61) 20.

¹³³ AI Act, at chs 4 and 5; Kelly-Lyth (n 24) 6–8.

¹³⁴ Article 29 Working Party (n 127) 18.

¹³⁵ Kaminski and Malgieri (n 38) 133.

¹³⁶ GDPR art 36.

¹³⁷ Kaminski and Malgieri (n 38) 130.

¹³⁸ Freedom of information request submitted by the author to the Information Commissioner’s Office, response received 8 December 2021.

¹³⁹ Gilbert and Thomas (n 1).

A closely related shortcoming of the DPIA is the inadequacy of its consultation element. The GDPR provides that '[w]here appropriate, the controller shall seek the views of data subjects or their representatives' as part of the DPIA.¹⁴⁰ This obligation to 'seek views' does not afford the workers' voice any force: there is no requirement for the employer to *address* the views it receives.¹⁴¹ Once again, the shortcoming is compounded by the absence of transparency. The employer need not even inform the workers or their representatives about the outcome of the assessment, so oversight is severely curtailed. We note that this requirement does not form part of the UK's proposed Data Protection and Digital Information Bill.

The DPIA obligation exists alongside other laws, and in some jurisdictions, existing labour consultation obligations may fill in the gaps. In Germany, for example, the use of 'technical devices designed to monitor the behaviour or performance of the employees' is subject to co-determination.¹⁴² At the other end of the European spectrum, even employers in the UK are required to disclose some information to recognised unions for the purposes of collective bargaining.¹⁴³ This obligation only applies in a narrow set of circumstances, however, and is of limited value given the country's low union density.¹⁴⁴ The challenge is that there is significant heterogeneity in industrial relations across Europe, and the inadequacy of the DPIA is not always mitigated by other laws.

In summary, although the DPIA is similar to the proposed ARIA in terms of aim and scope, the two are not functionally equivalent. On the other hand, a proliferation of impact assessments is clearly undesirable in practical terms. How, then, to marry the two?

Kazim and Koshiyama suggest three options: (i) an AI impact assessment which 'sits on top of' the DPIA; (ii) a DPIA which is 'adapted and modified' for AI; and (iii) an AI impact assessment which is independent of the DPIA.¹⁴⁵ Despite recognising that a DPIA 'may be required when using any AI system',¹⁴⁶ the authors suggest that options (i) and (ii) pose a fundamental challenge: while data protection is 'mainly an expression of privacy', 'AI impact' is 'an expression of the value of fairness', and these can come into 'direct conflict' with one another.¹⁴⁷ They therefore argue that integration is possible only if a 'fundamental value judgment' can be made about the prioritisation of privacy versus fairness. If such *ex ante* prioritisation is not possible, for example because value trading is 'context specific', then 'data [processing] and AI [systems] . . . will each need [their] own impact assessment[s] and how the two relate will have to be worked out through another mechanism'.¹⁴⁸

This argument does not hold, for two reasons. First, it relies on a false equivalence being drawn between the right to privacy and the right to data protection.¹⁴⁹ Although courts

¹⁴⁰ GDPR art 35(9).

¹⁴¹ Nicholas Martin and others, *The Data Protection Impact Assessment According to Article 35 GDPR* (Fraunhofer Verlag 2020) 15.

¹⁴² Works Constitution Act (Section 87 (1) No 6 BetrVG), although there are indications that co-determination rights are being underused in this area and also require strengthening: Deutscher Gewerkschaftsbund, 'DGB Concept Paper: Artificial Intelligence (AI) for Good Work' (2020) 10–11 <<https://perma.cc/SL6L-XC5R>>.

¹⁴³ Trade Union and Labour Relations (Consolidation) Act 1992 s 182(2). The disclosure duty only applies for the purposes of collective bargaining in relation to matters specified at TULRCA s 178(2).

¹⁴⁴ Department for Business, Energy & Industrial Strategy, 'Trade Union Membership, UK 1995-2020: Statistical Bulletin' (2021) <<https://perma.cc/R93T-EDX5>>.

¹⁴⁵ Kazim and Koshiyama (n 122).

¹⁴⁶ *ibid* 222.

¹⁴⁷ *ibid* 219–220.

¹⁴⁸ *ibid* 223.

¹⁴⁹ See, for example, *ibid* 221, suggesting that recital 1 of the GDPR, which states that '[t]he protection of natural persons in relation to the processing of personal data is a fundamental right', 'appears to equivocate respect for human dignity with respect for privacy'.

have tended to conflate these rights in the past,¹⁵⁰ the right to protection of personal data is protected independently from the right to privacy under the Charter of Fundamental Rights¹⁵¹ and is inspired by different goals, including the reduction of power and information asymmetries between data subjects and data controllers.¹⁵²

Secondly, even if the right to data protection *were* a subset of the right to privacy, Kazim and Koshiyama's argument seems to suggest that the DPIA process requires privacy to be promoted above all else. This is not so: as above, the DPIA should consider the full set of 'rights and freedoms of data subjects'.¹⁵³ To the extent that contextual value trade-offs need to be made, the DPIA is precisely the forum for such deliberation. Indeed, at a more general level, it makes more sense to consider value trade-offs within a single process than designing separate processes which are 'fundamentally in tension', leaving the conflicting outcomes to be 'worked out through another mechanism'.¹⁵⁴

In practical terms, a single impact assessment should be carried out for each (set of) algorithmic management tool(s). Duplication could be prevented by making two simultaneous legislative changes: introducing the ARIA obligation outlined herein; and adjusting the DPIA obligation such that it is automatically satisfied if the impacts of the high-risk data processing have been considered as part of the ARIA.¹⁵⁵

Framing the obligation in this way makes the DPIA subservient to the ARIA in the labour context. This is sensible when one considers that the ARIA is filling in the gaps of the DPIA: reversing the approach would mean starting from a weaker position. Rolling up assessments should be no anathema to DPIA advocates: Article 35 itself specifies that '[a] single assessment may address a set of similar processing operations that present similar high risks'.¹⁵⁶ The proposal would therefore enable the added value of the ARIA to be realised while minimising any additional burden.

6. Conclusion

This contribution has proposed the introduction of a framework for undertaking algorithmic impact assessments for tools which exercise, or inform the exercise of, the managerial prerogative at work. To be effective, these new impact assessments must be imposed at least at the design, development, and deployment stages; must consider impacts on workers' rights and freedoms and impacts on work conditions and quality; and must meet certain procedural and substantive minima, including recording and disclosure of relevant documents, effective consultation and genuine transparency. While DPIAs imposed by the GDPR go some way towards realising this proposal, marked gaps remain, both in terms of the scope of application and the procedural standards. The new ARIA should therefore incorporate but build on the DPIA obligation and model, combining it with other dimensions for assessment and expanding its remit and bite to consider all the primary dimensions of good quality work that shape the employee experience from legal, ethical, and social perspectives. We propose the Good Work

¹⁵⁰ Orla Lynskey, 'Deconstructing Data Protection: The 'Added-Value' of a Right to Data Protection in the EU Legal Order' (2014) 63 International & Comparative Law Quarterly 569, 574–581.

¹⁵¹ CFR arts 7 and 8.

¹⁵² Lynskey (n 150) 588–597. See also Mantelero (n 19) 761, suggesting that although 'part of the legal doctrine has pointed out the role of human dignity as a foundational ground of data protection in Europe', there are indicators that 'a broader range of values' in fact 'underpin data protection'.

¹⁵³ See text to n 124.

¹⁵⁴ Kazim and Koshiyama (n 122) 223–224.

¹⁵⁵ The assessment would also have to include the components specified at GDPR art 35(7).

¹⁵⁶ GDPR art 35(1).

Charter as a practical framework for undertaking such a hybrid model, and for legislating for it.

While this chapter has considered the granular details of impact assessments in the labour context, it is possible that eventual legislative action on algorithmic impact assessments will be cross-sectoral, and that sectoral detail will be left to secondary legislation or regulatory instruments.¹⁵⁷ It is therefore useful to identify those points which apply to algorithmic impact assessment proposals more broadly.

The first is that despite the tendency towards omnibus AI regulation, sectoral specificity enables a shift away from ‘risk’ in the abstract and the consequent scope for creative interpretation to curtail scope of application. Regulated technologies, impacts, and procedural steps should be identified as concretely as possible, and this is more feasible at the sectoral level, suggesting that guidance on methods and measurements for the ARIA should follow.

Secondly, experience with the DPIA has demonstrated that transparency and stakeholder power are crucial for technological impact assessments to be effective. An absence of transparency and meaningful stakeholder engagement seriously undermines oversight mechanisms, and ‘consultation’ can be spurious when it takes place without the information, capabilities, or space for active participation in a context of significant power imbalance.

Thirdly, and perhaps more critically, impact assessments are both part of a broader governance regime and will inform its development over time, as ARIAs generate new evidence and understanding about impacts on people, and detailed guidance and case law is built up.¹⁵⁸ ARIAs enable co-regulation where red lines are too blunt, but they do not displace the need for red lines; they provide space for stakeholder voice, but they do not address broader power structures; and they create *ex ante* consideration of impacts, but do not secure access to justice for those who nonetheless experience harm themselves. While this contribution highlights the significant potential value of impact assessments in the employment context, they cannot tackle the novel concerns posed by algorithmic management in isolation: that will require a rethinking of the entire regulatory toolbox.

¹⁵⁷ See, for example, Public Authority Algorithm HL Bill cl 4(5), which would require the Secretary of State to ‘prescribe the form of an Algorithmic Impact Assessment’; and Algorithmic Accountability Act of 2022, which would require the Federal Trade Commission to detail the requirements of impact assessments of automated decision systems and augmented critical decision processes.

¹⁵⁸ Selbst (n 48).