Guest Editorial: Regulating algorithmic management

**Jeremias Adams-Prassl, Halefom Abraha,
Aislinn Kelly-Lyth, Sangh Rakshita, and
Michael 'Six' Silberman***

Abstract

This special issue of the European Labour Law Journal, edited by Jeremias Adams-Prassl, Halefom Abraha, Aislinn Kelly-Lyth, Sangh Rakshita and Michael 'Six' Silberman, explores the regulation of Algorithmic Management in the European Union and beyond. In our guest editorial, we set out the background to the project, introduce the reader to the key themes and highlights of the papers to follow, and acknowledge the support that the project has enjoyed.

Keywords

algorithmic management, AI at work, discrimination, data protection, impact assessments, artificial intelligence, data trusts, collective rights

Introduction

Digitalisation is revolutionising the world of work. Fears of widespread technological unemployment driven by the rise of artificial intelligence continue to prove unfounded—but this should not detract us from the fundamental impact the deployment of emerging technologies has on the organisation of work. The rapid pace of technological innovation has set the stage for the rise of algorithmic management (ARM): the potential automation of the full range of traditional employer

Corresponding author:
Jeremias Adams-Prassl
jeremias.adams-prassl@law.ox.ac.uk

functions, from hiring workers and managing the day-to-day operation of the enterprise through to the termination of the employment relationship.[1]

The origins of ARM are closely linked to the advent of the gig economy.[2] But they are no longer limited thus: boosted by the Covid-19 pandemic, ARM systems are quickly becoming omnipresent in the labour market, not least through integration into existing systems, from word processing to enterprise management at large.[3] Vendors promise solutions covering the full range of employer functions, in workplaces across the socio-economic spectrum.

The promise—and perils—of ARM have been extensively documented in the literature.[4] Policymakers and regulators are starting to take note of the need to tackle the problems associated with the deployment of ARM systems, from algorithmic opacity to mental and physical harm.[5] They face a range of complex questions: To what extent can existing norms evolve to address these problems? Which harms are genuinely novel? And, more fundamentally, how can we foster genuine innovation whilst also protecting fundamental rights at work?[6]

The race to regulate AI: Global comparative perspectives

These questions sit at the core of iManage, a five-year interdisciplinary research project funded by the European Research Council to explore the promise—and perils—of algorithmic management systems. As AI systems are quickly becoming ubiquitous, so too are the regulatory challenges: across the world, new models have begun to emerge at national, regional, and international levels. There is general consensus that existing instruments are insufficient, potentially even inadequate, to deal with the myriad challenges this development poses—but also considerable divergence on what new AI regulatory frameworks should look like.

In the course of a high-level conference convened at Oxford University in June 2022, we explored AI regulation through a comparative lens and examined the extent to which existing and proposed approaches across the world adequately address risks whilst maximising the benefits of AI systems. Discussions ranged from data protection and omnibus regimes such as the European Union's proposed AI Act to transnational approaches and sectoral regulation, focused on employment as a central case study.

The conference was followed by a one-day workshop bringing together academics, technical experts, and policymakers to discuss different regulatory options for AI at work, based around a draft blueprint for regulating algorithmic management, developed by the iManage project

---

1. Jeremias Adams-Prassl, 'What if Your Boss was an Algorithm? Economic Incentives, Legal Challenges, and the Rise of Artificial Intelligence at Work' (2019) 41 Comparative Labor Law and Policy Journal 123.
2. Jeremias Prassl, *Humans as a Service: the Promise and Perils of Work in the Gig Economy* (Oxford University Press 2018); Antonio Aloisi and Valerio De Stefano, *Your Boss Is an Algorithm: Artificial Intelligence, Platform Work and Labour* (Hart 2022).
3. Wolfie Christl, 'Digitale Überwachung und Kontrolle am Arbeitsplatz' (Cracked Labs 2021).
4. For early comprehensive accounts, see the special issue of the Comparative Labor Law & Policy Journal (2019) 41.
5. For recent initiatives in the EU and North America, see for example the provisions on algorithmic management (Chapter III) of the proposed 'Platform Work Directive' (Proposal for a Directive of the European Parliament and of the Council on improving working conditions in platform work, COM (2021) 762 final (9 December 2021) (Platform Work Directive)); Canada's Directive on Automated Decision-Making 2019; California Bill AB-1651 Workplace Technology Accountability Act 2022; Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People October 2022.
6. European Commission, 'Fostering a European Approach to Artificial Intelligence' (Communication) COM (21) 205 final, 8.

team. We were joined in discussion by the authors of the contributions collected in this Special Issue, ranging from discrimination and data protection law to broader technical and comparative questions.

Regulating AI at work

Writing in March 1999, Spiros Simitis set out to offer a series of 'Prolegomena to an EU Regulation on the Protection of Employees' Personal Data', 're-iterat[ing] the need for a Regulation on the protection of employees' data.' His aim was both to outline 'the nature, provisions and scope which such a regulation should entail', and to 'reflect [on] both the reality of the modern employment relationship, and a new normative vision of the workplace which aims to inject such relationships with a measure of communicative participation.'[7]

More than two decades on, his aspirations are reflected in the opening piece of this collection, a comprehensive blueprint which sets out regulatory options in response to the rise of algorithmic management. A brief technical introduction sets out the underlying normative case for regulating algorithmic management by identifying two regulatory gaps: the exacerbation of privacy harms and information asymmetries, and the loss of human (especially, but not only, managerial) agency. There follow eight concrete policy measures designed to address these gaps, each with a detailed rationale explaining the regulatory choices involved. The first set of options focuses on protection against privacy harms and overcoming information asymmetries, with options ranging from explicit prohibitions ('redlines') and purpose limitations to information and data access rights.

Discussion then turns to options designed to re-establish agency for management as well as workers and their representatives. Instead of focusing exclusively on bans on fully automated decision-making (the oft-discussed requirement for a 'human in the loop'), the blueprint proposes a series of interventions across the life cycle of algorithmic management, from design and deployment to operations and review. This includes a role for 'humans before the loop', establishing requirements for ARM design and deployment; 'humans after the loop', *viz*, rights for workers to challenge, request explanations for, and request human review of decisions that affect them; and 'humans above the loop', monitoring the broader implications of ARM through dedicated impact assessments.

Following on from the blueprint, a first set of articles explores different ways in which existing norms in closely related fields apply to the deployment of algorithmic management systems. *Aislinn Kelly-Lyth*'s examination of European *Anti-Discrimination Law* finds the acquis to be 'remarkably robust', with most of the gaps identified in the literature hitherto predating the digitalisation of the workplace. With careful judicial development, existing norms will be able to tackle many supposedly novel legal questions: algorithmic 'accuracy' or a counterfactual biased human-decision maker, for example, are much less likely to be accepted as justifications for indirect discrimination than broadly assumed. The prohibition of direct discrimination similarly holds significant promise in tackling key mechanisms of algorithmic bias, despite potential difficulties in liability attribution against a backdrop of complex AI supply chains, as does a broad reading of employers' duty of reasonable accommodation in the disability context.

---

7. Spiros Simitis, 'Reconsidering the Premises of Labour Law: Prolegomena to an EU Regulation on the Protection of Employees' Personal Data' (1999) 5 European Law Journal 45, 45.

Algorithmic discrimination, Kelly-Lyth argues, thus confronts equality law with a paradox: whilst automated decision-making processes are frequently opaque, the digitalisation of previously human decision-making processes can render discriminatory criteria more traceable and outcomes more quantifiable. Our primary focus should therefore not be on wholesale reinvention, but rather on a renewed effort at developing and enforcing established norms, not least by overcoming the significant information asymmetries inherent in the deployment of algorithmic management systems: '[g]etting cases of algorithmic discrimination before courts is critical'.

The results of an in-depth review of European *Data Protection Law* are somewhat more mixed: *Halefom Abraha* suggests that existing provisions of the General Data Protection Regulation (GDPR) offer some protection against the harms arising from algorithmic management. This includes the right to be informed and a right of data access under Article 15, recently deployed against major warehouse operators in several jurisdictions. Article 22 provides further protection in the context of fully automated algorithmic management decisions. Significant gaps remain, however. First, due to a lack of specificity: both the scope of the Article 15 right and its carve-outs and limitations are unclear, whereas Article 22 'remains the most complex and controversial provision, in both theory and practice' in the GDPR. Second, due to context-specific adequacy of the GDPR as a general data protection instrument in the employment field: issues specific to the world of work do not neatly fit within the omnibus remit of data protection law, whether it is consent as a legal ground for processing or the individualistic nature of data protection law.

These gaps matter: while algorithmic management is regulated in a host of legal domains, Abraha suggests, data protection law has traditionally been the area of law that is most engaged, owing to the vast troves of personal data processing underpinning the majority of these tools. The solution thus lies in strengthening, complementing, and particularising existing protections, through both legislative and non-legislative measures, both domestically and at the Union level.

Discussion then turns to occupational safety and health (OSH): *Aude Cefaliello, Phoebe V Moore, and Robert Donoghue* explore the applicability—and shortcomings—of the existing acquis to algorithmic management, with a particular emphasis on its *Psychosocial Risks and Safety by Design*. The gaps in existing OSH regulations are significant: many of the risks faced by workers resulting from ARM fall in the domain of the psychosocial, which existing measures fail to address adequately. The harms identified include threats to the psychological contract and trust; bias, discrimination, and unfairness; deskilling; loss of worker autonomy and privacy; function creep; and increased opportunities for disciplining workers and for work intensification and acceleration.

Against this backdrop, the authors turn to scrutinise the prevailing 'one-off', 'safe-by-design' approach to OSH risk assessments—and identify four main deficiencies in existing regulation: its narrowness in scope, a lack of explicit reference to ARM harms, a lack of recognition of the organisational realities of the employment context, and, crucially, an inability to account for the dynamic emergence of previously unanticipated OSH risks from the ongoing deployment of algorithmic management technologies in specific workplaces. There is, they conclude, a clear need for standalone legislation to address the OSH harms of algorithmic management: 'safety' means different things in the design, implementation, and usage phases. This difference can only be fully addressed by a shift towards a new approach to risk assessment and mitigation in which formal channels and structures are created through which workers can surface, and work together with management to address, risks as they become apparent. They describe this approach as 'design for responsibility'—and here, crucially, what must be 'designed for responsibility' is not only the technology, but also the organisational systems by which its impacts are apprehended and mitigated.

In regulating algorithms at work, the collective dimension is at least as important as the focus on individual rights. In their contribution exploring the *Collective Regulation of Algorithmic Management*, *Zoe Adams and Johanna Wenckebach* make a strong case for co-determination at work, with a particular emphasis on the governance of algorithmic management. Technology's exacerbation of the inequalities of bargaining power inherent in employment relationships are key to the underlying normative framework 'focused [not] on regulating technologies; [but] instead on regulating power relations': this understanding of workers' 'systematic disadvantage' explains why consultation and negotiation alone are not enough: joint regulation is required.

It is against this yardstick that both existing domestic (UK and German) and Union law, as well as the proposed *Blueprint for Regulating Algorithmic Management* are evaluated, and key obstacles identified. Whilst welcoming key policy options in the blueprint, Adams and Wenckebach identify two problems that persist notwithstanding the proposed improvements to the status quo: the absence of any explicit rights of co-determination; and a failure to specify how the collective provisions will operate in workplaces in which structures of representation are wholly absent, or inadequate. They provide specific proposals for the further protection of collective governance, including mandatory Data Protection and Technology Committees, compulsory sector-level agreements concerning algorithmic red lines for the sector, and a strengthening of the autonomous collective self-regulation of work; they conclude with an emphasis on the need to ensure that working conditions generally allow collective regulation to flourish, whilst also acknowledging the inherent limitations of 'top-down mechanisms' in achieving this goal.

This concluding point is taken up in the contribution by *Aislinn Kelly-Lyth and Anna Thomas*, who explore the role of *Algorithmic Impact Assessments* in 'complement[ing], as well as inform[-ing], an overarching "top-down" framework for the governance of algorithmic management systems': identifying risk mitigations on a case-by-case basis permits context-specific responses, ensures that appropriate measures can be built into system design and deployment *ex ante*, and strikes the appropriate balance between broad-brush generalised requirements and free-for-all self-regulation. In order to ensure that the duty to conduct an impact assessment leads to effective worker protection, the authors propose a number of detailed criteria, guided by the Good Work Charter developed by the Institute for the Future of Work. These criteria cover the timing and actors involved (including distinct obligations on different parties at relevant points in the AI life cycle); impact assessments' substantive scopes (*viz*, the full range of impacts to be covered at both the individual and collective level); and appropriate procedures (from stakeholder involvement to publication of outcomes at defined intervals).

When analysed through this lens, existing obligations in the GDPR and domestic data protection law to carry out Data Protection Impact Assessments (DPIAs) in high-risk processing contexts fall short, whether it is in terms of scope or consultation process. In concluding, Kelly-Lyth and Thomas therefore set out concrete proposals to overcome these weaknesses, whilst also avoiding unnecessary duplication and a proliferation of impact assessments.

Reflecting on a similar theme, *Dan Calacci and Jake Stein* then explore the potential of *Collective Data Governance for Workers*: analysing workers' current data access and use practices, they suggest that existing data protection law fails workers, as the primary focus of current regulatory frameworks of data collection and use in the workplace should be on improving working conditions, rather than merely ensuring data protection: a focus on privacy and the individual data subject alone ignores how data are deployed fundamentally to reshape the workplace. The complexities of worker data autonomy are developed in a hypothetical case study of call centre workers to demonstrate how the deployment of a multitude of data-intensive systems at work, and the

increasing complexity of data processing supply chains through the growing importance of third-party software providers, is liable to render a data-protection focused approach to worker protection intractable and thus ineffective.

Received data subject rights and data protection should therefore not be the basis for data rights in the workplace: what is required instead is support for nascent approaches grounded in labour law. The authors conclude that only a combination of access rights, liability mechanisms, and worker co-determination could help support workers in algorithmic and data-driven workplaces. Algorithmic management should be met with a combination of strategies, including the use of collective data subject access requests; the enactment of new, sector-specific employment regulation; and the construction of 'data intermediaries'.

Whilst the main focus of this special issue has been on the regulation of algorithmic management at the European level, regulators around the world have begun to grapple with the underlying phenomenon. *Antonio Aloisi and Valerio De Stefano* assess the ensuing *Transatlantic Race to Govern AI-driven Decision-making* from a comparative perspective, with a particular emphasis on the *Tension between Risk Mitigation and Labour Rights Enforcement*. In comparing approaches from the EU (including the GDPR and recent proposals for an AI Act and a Platform Work Directive), the United States, and Canada, a common emphasis on risk assessment and risk management is identified. This, the authors argue, is wholly unsuitable to the enforcement and protection of fundamental rights at work: non-negotiable workplace rules cannot give way to 'cosmetic audits, vague checklists and courtesy toolkits'.

On the basis of a detailed scrutiny of different regulatory approaches, the shortcomings of a 'scalable, decentralised' governance model, outsourced in part to key actors themselves, are identified, before highlighting the potentially stark consequences of the current approach: a 'frenzied race to regulate AI', with all the ensuing complications and contradictions which abound, particularly in competing EU proposals, 'could slow technological innovation, undermine national frameworks regulating the introduction of workplace technologies, and erode fundamental rights.'

Taking up the challenges of technical and regulatory innovation, *Michael Veale, Michael 'Six' Silberman, and Reuben Binns* turn to the European Union's proposed Platform Work Directive (PWD), setting out concrete proposals designed further to *Fortify the Proposal's Algorithmic Management Provisions*. They identify a series of challenges beyond platform workers' misclassification as self-employed contractors, ranging from unaccountable instances of non-payment and account suspension to ineffective and unfair automated decision-making and uncommunicative customers and platform administrators.

Potential responses include improvements to *ex ante* and *ex post* algorithmic transparency (notably as regards the meaning of algorithmic transparency); identifying and strengthening the standard against which human reviewers of algorithmic decisions review (in particular as regards applicable substantive standard(s) of review, including not least a restatement of the reasons for a decision in a simple, human comprehensible manner, and supported by appropriate evidence); anticipating challenges of representation and organising in complex platform contexts (far from traditional paradigms of workplace organisation, from an absence of traditional representation structures to the complexities of consulting across language and jurisdictional borders); creating realistic ambitions for digital worker communication channels (whether in terms of the challenges inherent in the creation of civic technologies, or the proposed confidentiality requirements); and accountably monitoring and evaluating impacts on workers while limiting data collection. On balance, however, the authors conclude that none of the requirements of the PWD as proposed present insurmountable technical challenges, nor do they run counter to the interests of users—

the aim is to ensure that EU (and Member State) regulators can better ensure that the PWD successfully protects platform workers in practice.