Online Tribunal Judgments and The Limits of Open Justice

*Zoe Adams, Abi Adams-Prassl, Jeremias Adams-Prassl**

Forthcoming, (2021) 41 Legal Studies

Abstract

The principle of open justice is a constituent element of the rule of law: it demands publicity of legal proceedings, including the publication of judgments. Since 2017, the UK government has systematically published first instance Employment Tribunal decisions in an online repository. Whilst a veritable treasure trove for researchers and policy makers, the database also has darker potential – from automating blacklisting to creating new and systemic barriers to access to justice. Our scrutiny of existing legal safeguards, from anonymity orders to equality law and data protection, finds a number of gaps, which threaten to make the principle of open justice as embodied in the current publication regime inimical to equal access to justice.

Keywords: Employment Tribunals, Online Judgments, Open Justice, Machine Learning, Data Protection, Algorithmic Discrimination

1. INTRODUCTION

Since February 2017, Employment Tribunal decisions from England, Wales, and Scotland have been made freely available online at www.gov.uk. At the time of writing, the website lists over 70,000 decisions, uploaded as individual .pdf or text files, searchable by keywords, jurisdiction code, and date. Accessible with relative ease, the repository provides a veritable treasure trove

* Research Fellow, King's College, University of Cambridge; Senior Research Fellow and Associate Professor, Department of Economics, University of Oxford; Professor of Law, Magdalen College and Faculty of Law, University of Oxford. We acknowledge funding from the Economic and Social Research Council, Grant No ES/ S010424/1, and are grateful to Stergios Aidinlis, John Armour, Alysia Blackham, Natalie Byrom, David Erdos, Aislinn Kelly-Lyth, Ravi Naik, Ernest Ryder, Hannah Smith, Fabian Stephany, and Stefan Theil, as well as participants at a workshop in Oxford in March 2020 and the anonymous reviewers for feedback and discussion. The usual disclaimers apply.

for researchers interested in understanding the operation of employment and discrimination law in practice, as well as the justice system more broadly.

This resource is urgently needed: employment tribunals are a high priority in the Ministry of Justice's and HMCTS' digitalisation reforms, and the Covid-19 pandemic has put significant further strain on the system. Whilst championed in some quarters,¹ the reforms have come in for sustained criticism, with both the National Audit Office and the House of Commons Justice Select Committee questioning the speed, scale, and effectiveness of the proposed reforms – as well as their potential impact on the constitutional right of access to justice.² What all parties seem to agree on is the need for high-quality empirical evidence about the current operation of the system, and the impact of any changes. Existing data sources, however, make any detailed evaluation or broader understanding difficult: the regular statistical bulletins compiled by HMCTS are published with significant time-lag and consist of only highly aggregated data; official in-depth surveys were discontinued in 2013.³

The publication of employment tribunal decisions, moreover, is not just of importance to researchers and policy makers: it is a constitutional imperative. There is a statutory duty on the 'Lord Chancellor [to] maintain a register containing a copy of all judgments and written reasons issued by a Tribunal'.⁴ Public access to the details of judicial proceedings is a central component of the principle of open justice, intimately related to the common law right of access to a court.⁵ As the Supreme Court emphasised in *UNISON*, albeit in the context of tribunal fees, without 'unimpeded access' to courts, the 'law is liable to become a dead letter.'⁶ Courts and tribunals must be configured in such a way that claims can be tried and adjudicated effectively; there must be transparency, and accountability, in judicial decision-making, and

¹Richard Susskind, *Online Courts and the Future of Justice* (OUP 2019).

² National Audit Office, *Transforming Courts and Tribunals – A Progress Update*, HC 2638 Session 2017–2019 13 September 2019 <<https://www.nao.org.uk/wp-content/uploads/2019/09/Transforming-Courts-and-Tribunals.pdf>>; Justice Select Committee, *Courts and Tribunal Reforms*, Second Report of Session 2019 HC 190, 30 October 2019 <<https://publications.parliament.uk/pa/cm201919/cmselect/cmjust/190/190.pdf>>

³ See, for example, <<https://www.gov.uk/government/statistics/tribunal-statistics-quarterly-january-to-march-2020>>.

⁴ SI 2013/1237, reg 14(1); see also The Employment Tribunal Rules of Procedure 2013 (as amended), rule 67.

⁵ Ryder SE, 'Securing Open Justice' in Burkhard Hess and Ana Koprivica Harvey (eds), *Open Justice* (Nomos Verlagsgesellschaft mbH & Co KG 2019) <<https://www.nomos-library.de/index.php?doi=10.5771/9783845297620-125>> accessed 21 January 2020

⁶ *UNISON, R (on the application of) v Lord Chancellor* [2017] UKSC 51, at [68]

clarity, and predictability, in the law. To this end, the publication of court and tribunal judgments, the opening up of courts and judgments to public scrutiny, is essential.⁷

At the same time, however, it is also important to consider key countervailing considerations – in particular, parties’ right to protection of their private lives. While the importance of these principles and their inherent tradeoffs are widely accepted, the way in which they are to be realised in practice is far more contested. Take anonymization as an example. In the UK, it has long been assumed that open justice demands the publication, in all but the most exceptional cases,⁸ of the names and personal details of claimants and respondents.⁹

This sets the UK apart from its European counterparts, where the anonymisation of judgments is more often routine.¹⁰ It also has important implications for data protection:¹¹ Insofar as the processing of personal data by the judiciary in the course of its judicial functions is concerned, the public interest in open justice is deemed to outweigh, in all but the most exceptional cases, the private interest in control over one’s personal data.¹² Publication of claimants’ names and details thus constitutes an express exception to the rights provided to data subjects under the GDPR (The EU General Data Protection Regulations) and DPA (Data Protection Act 2018).¹³

⁷ These principles are also embodied in Article 6(1) European Convention of Human Rights, and have been reiterated by the Supreme Court in: *R (Guardian News and Media) v City of Westminster Magistrates’ Court* [2012] EWCA Civ 420.

⁸ The general practice in the Family Division is for judgments in ancillary relief cases, if published, to be anonymised (see *Lykiardopulo v Lykiardopulo* [2010] EWCA Civ 1315 [79]). However, ordinary civil claims brought by children are not routinely anonymised (see *JXF v York Hospitals NHS Foundation Trust* [2010] EWHC 2800 (QB)). The fact that the case concerns private information is not, of itself, a sufficient basis for making an anonymity order (see eg *Bernard Gray v UVW* [2010] EWHC 2367 (QB)).

⁹ ‘Judiciary and Data Protection: Privacy Notice’ <<https://www.judiciary.uk/about-the-judiciary/judiciary-and-data-protection-privacy-notice/>> accessed 13 January 2020. Indeed, an order for anonymity is seen as a derogation from the principle of open justice and an interference with the Article 10 rights of the public at large (See eg *Bernard Gray v UVW* [2010] EWHC 2367 (QB) [1]).

¹⁰ For a detailed comparison of the position in different countries, see: Marc van Opijnen and others, ‘On-Line Publication of Court Decisions in the EU’ the European Union 178.

¹¹ ‘Judiciary and Data Protection: Privacy Notice’ <<https://www.judiciary.uk/about-the-judiciary/judiciary-and-data-protection-privacy-notice/>> accessed 13 January 2020

¹² See, for example: “*A threatened interference with the Art 8 rights of a claimant is not, by itself, always sufficiently serious to necessitate the imposition of an injunction or anonymity order*” (*JIH v News Group (No.2)* [2010] EWHC 2979 (QB) [30]).

¹³ The exemption is set out in article 23(1)(f) GDPR and section 15(2)(b) and Schedule 2, part 2, para6 and Schedule 2, part 2, para 14(2) Data Protection Act 2018. The reason for the exception, according to the judicial Working Group, being ‘to secure the constitutional principles of judicial independence and of the rule of law’

What is at stake here is not simply how to strike an appropriate balance between the public and private interest, however.¹⁴ Rather, the issue is how to ensure that the principles of open, and equal access to justice can be reconciled in practice. Different contexts may well require distinct balances to be struck. Employment Tribunal judgments provide a particularly salient case study in this regard, given the inequality of bargaining power between employer and employee, and the social significance of work to the individual.

Indeed, whilst the publication of data extracted from tribunal judgments serves, at first glance, unequivocally to further the principle of open and equal access to justice, upon closer inspection, the ensuing datasets may also be used to undermine it – from evading employment law and exacerbating existing labour market inequalities to creating new barriers to access to justice. As automated online reputation screening tools begin to play an increasingly salient role in the recruitment process,¹⁵ for example, it is not difficult to envisage how an unscrupulous provider of such services might rely on the tribunals database to facilitate the de-facto blacklisting of applicants who have brought claims in the past.

When thinking about the sort of data to publish, and to whom that data should be available, then, we cannot think exclusively in terms of the ideal purposes to which we would like that data to be put: possible abuses of that data, and appropriate legal safeguard to avoid them, must also enter into the equation.

In this article, we take digitally available Employment Tribunal decisions as a case study in order to explore the promise and perils of publishing tribunal decisions online. Two different sorts of questions must be answered. First, what sort of knowledge about the tribunal system and/or legal framework can data extracted from tribunal judgments actually facilitate, and for the pursuit of what sort of purposes can that information be put? Second, what risks of abuse arise from the publication of the judgment data, and how might we prevent or manage those risks?

In developing our answers, discussion is structured as follows. Section 2 outlines how data on employment tribunal litigation can be extracted from the Government’s online database. Section 3 turns to the broader range of potentially beneficial uses to which the corpus of information contained in the judgments might be put: from identifying litigation patterns to predicting individual legal outcomes. In contrast, section 4 explores some of the more

¹⁴ This tension is implicit in article 6 ECHR which states: “*the press and public may be excluded from all or part of the trial in the interest of morals, public order or national security in a democratic society, where the interests of juveniles or the protection of the private life of the parties so required, or to the extent strictly necessary in the opinion of the court in special circumstances where publicity would prejudice the interests of justice*”.

¹⁵ See, for example, <https://fama.io/>.

questionable ends to which the same data could be deployed, be it blacklisting or skewed settlement negotiations. Section 5, finally, surveys a range of existing safeguards – from anonymity orders to data protection and equality law – to highlight the gaps which threaten to make the principle of open justice as embodied in the current publication regime inimical to equal access to justice. A brief conclusion charts the short- and longer-term reforms required to realign these fundamental constitutional principles.

2. ONLINE JUDGMENTS

The availability of online judgments has huge potential value for research into the operation of employment law, and the legal system more broadly. While publicly available statistics provide some insight into the overall level of cases in the justice system (albeit with a considerable time delay), they provide few details about claimants' and defendants' characteristics, the legal points argued, or damages awarded. Official statistics on claims received and disposed of by the ETs are furthermore only available with a considerable time lag and in an aggregated format. The Quarterly Tribunal Statistics *Employment Tribunal and Employment Appeals Tribunal Annual Tables*, for example, reports the average compensation awarded at tribunal for seven jurisdictions with a year's delay. No information is available on many other important indicators, from gender composition of claimants to regional breakdown.

There is a long tradition in labour law research of using social science techniques in the study of law, with published judgments recognized as a potentially rich source of data about the social world and the law's role within it.¹⁶ Judgments will usually contain information on the characteristics of claimants, respondents, the facts of the case, and an indication of the relevant legal principles applied by the judge in coming to a decision. Analysis of these texts can allow for a deeper understanding of how legal decisions are made, how policy is translated into legal reasoning, and how, and why, policy goals may or may not be met in practice.

Prior to 2017, any analysis of ET judgment data would have involved the collection of paper copies of tribunal decisions. This process was laborious, time-consuming, and resource-intensive. Even relatively small sample sizes required extensive work: research assistant visits to Bury St Edmunds, the historical repository of physical ET records,¹⁷ fortuitous access via

¹⁶ See, for example, the work of legal realist scholars: Sally Engle Merry Edited by Heinz Klug, *The New Legal Realism: 2* (Cambridge University Press 2016); Karl N Llewellyn, 'A Realistic Jurisprudence -- The Next Step' (1930) 30 Columbia Law Review 431.

¹⁷ Catherine Barnard, Amy Ludlow, Sarah Fraser Butlin, Beyond Employment Tribunals: Enforcement of Employment Rights by EU-8 Migrant Workers, *Industrial Law Journal*, Volume 47, Issue 2, July 2018, 226.

other bodies that collect particular sets of decisions,¹⁸ and/or the patient marrying up of electronic database information with physical files.¹⁹ Relatively straightforward computational techniques, akin to those explored in detail, below, go some way towards overcoming these challenges.

a. Data Access

The online database of decisions on Employment Tribunal cases in England, Wales, and Scotland is accessed via the UK government's central online portal, gov.uk. The website allows users to select judgments by country, jurisdiction code (e.g., age discrimination, equal pay act, health and safety, etc.), or date range. For each individual decision, a brief subpage giving the judgment document(s) is available. Figure 1 gives a screenshot of the website interface and the page corresponding to an individual decision (reached by clicking the name of the case).

Figure 1. Employment Tribunals Online²⁰

¹⁸ G Lockwood, P Rosenthal and A Budjnaovcanin, 'A Quantitative and Qualitative Analysis of Sexual Harassment Claims 1995-2005' (2011) 42 Industrial Relations Journal 86; Patrice Rosenthal and Alexandra Budjanovcanin, 'Sexual Harassment Judgments by British Employment Tribunals 1995-2005: Implications for Claimants and Their Advocates' (2011) 49 British Journal of Industrial Relations 236.

¹⁹ LD Irving, 'Challenging Ageism in Employment: An Analysis of the Implementation of Age Discrimination Legislation in England and Wales' (Coventry University 2012) 71
<https://curve.coventry.ac.uk/open/file/ffc88163-6994-4400-bead-121298f52bd1/1/Irving%202012.pdf> Cf also Lyndsey Bengtsson, 'Addressing Age Stereotyping against Older Workers in Employment: The CJEU and UK Approach' (2020) 62 International Journal of Law and Management 67.

²⁰ Screenshot of <https://www.gov.uk/employment-tribunal-decisions> taken on 3/3/2020.

The website makes accessing judgments radically easier and cheaper when compared with paper files. A search function is provided to filter judgments by country, jurisdiction, and decision date. Furthermore, the online database can easily be searched by the names of claimants and employers using a free text search function (see Figure 1a). It is also possible to process online judgments in bulk rather than relying on individual search requests; modern scraping and machine learning algorithms facilitate the automatic harvesting of online judgments and extraction of relevant information from within the judgments themselves with only minimal human involvement.

Both academic researchers and private companies are already availing themselves of this resource. A recent study by Blackham, for example, uses 1,208 judgments downloaded from the website to provide a qualitative and quantitative analysis of age discrimination judgments.²¹ Vizlegal, a company that provides searchable judgment information across a range of countries, have added ET and EAT judgments to their product offering.²² Solomonic, a legal intelligence firm, are using ET online judgments to ‘build a platform to assist employees and SMEs who are involved in employment tribunal claims to make more informed, data-backed decisions.’²³

In the following paragraphs, we outline in general terms how tribunal data can be bulk downloaded, processed, and analysed from the online repository. It is important to note that any such project would, in principle, itself fall within the ambit of the GDPR, as it involves the processing of personal data.²⁴ That said, as long as the processing is carried out with a view to publication of academic material, where the controller ‘reasonably believes that the publication of the material would be in the public interest’,²⁵ a series of exemptions and derogations apply.²⁶

²¹ Blackham, ‘Enforcing Rights’ (n) 6-8 provides a detailed explanation of her methodology.

²² See <https://www.vizlegal.com/>.

²³ <https://www.solomonic.co.uk/news/solomonic-to-develop-employment-dispute-resolution>

²⁴ GDPR, Art 4.

²⁵ DPA 2018, Schedule 2, part 5, para 26.

²⁶ For a detailed overview (*mutatis mutandis*), see M Mourby, H Gowans, S Aidinlis, H Smith, and J Kaye, ‘Governance of Academic Research Data under the GDPR – Lessons from the UK’ (2019) 9 *International Data Privacy Law* 192.

b. Extracting & Aggregating Information

Crawling, parsing, and scraping algorithms facilitate the aggregation of judgment information on a large scale, with low set-up costs.²⁷ In a first step, all the subpage links ('urls') for individual decisions are extracted from the main gov.uk website. This can be done using a web-crawling algorithm to identify the hyperlinks to individual decisions' subpages.

Each individual decision subpage contains a lot of relevant information on the case in a structured format that can be easily accessed by web-scraping. This information is usually referred to as metadata. For example, consider Figure 1b. Decision date (13th February 2020), country (England and Wales), jurisdiction codes (Breach of Contract, Public Interest Disclosure, Unlawful Deductions from Wages) are uniformly stored in the html structure of the subpage. Each subpage contains the link to the full decision reporting. In many cases, the hyperlink to this file contains keywords that allow further characteristics of the decision, such as name of claimant and respondent, case id or decision outcome, to be easily identified.

In a second step, the documentation underlying the individual decision can be downloaded for further analysis. The majority of these documents are pdfs, although some are in text format. Compared with individual subpages, information in these judgment files is harder to extract using automated tools, given the greater heterogeneity in judgment document structure and wording. However, natural language processing and supervised machine learning tools provide ample opportunities to harvest information on a number of dimensions: from legal representation and the value of any damages awarded to the relevant facts of the case.²⁸

3. THE USES OF JUDGMENT DATA

Assuming for the moment that online judgments are representative of all decisions in the ETs,²⁹ the data contained within them can be put to a number of uses: metadata, for example, can easily be extracted from online judgments to answer important questions for policy design and research, such as analysing the composition of ET cases. Moving beyond mere metadata analysis, more sophisticated methods have increasingly become viable through advances in machine learning, and other forms of 'artificial intelligence': when applied to judicial precedents, they may allow predictions of case outcome, potentially even elucidating judicial reasoning.

²⁷ This is usually achieved through a series of Python scripts.

²⁸ E Ash, D Chen and A Ornaghi, 'Gender Attitudes in the Judiciary: Evidence from the US Circuit Courts', (2020) CAGE Working Paper 462.

²⁹ We return to this question, below.

a. Descriptive Analysis of Case Characteristics from Metadata

A key potential benefit of the online judgment data is the ability quickly to track much more fine-grained data than is made available in publicly available statistics. To some extent, such useful information can be extracted from judgments even without the use of machine learning techniques. Take the analysis of the gender composition of claimants as an example. The disproportionate impact of the fees on low-value claims and on claims brought by women were central to the arguments presented before the Supreme Court in *UNISON*:³⁰ to what extent has this impact been reversed?

Case outcomes are not aggregated by gender in the QTS. However, gender differences can be analyzed using the title by which claimants self-identify (Ms/Mr etc); information that is readily available in the judgment hyperlink, thus not even requiring access to information contained within the judgment text.

b. Extracting information from judgment text

Many important case characteristics cannot simply be scraped directly from individual decision webpages on the gov.uk website: from the value of any monetary compensation awarded to whether parties were represented. These can only be extracted by accessing information within the judgment text itself.

Compared with individual subpages, information in these judgment files is harder to extract using automated tools, given the greater heterogeneity in judgment document structure and wording. Each judgment document is composed of a header and a flow text that explains the proceeding of the tribunal and the decision made. While the judgment header contains standardised information, such as the hearing date or name of the Employment Judge, the position and wording indicating these elements varies considerably. For example, the location of the tribunal might be indicated as “held at: Sheffield court” or “Court: Sheffield” or “Hearing at: Sheffield court”. Some judgments list representatives, while others do not.

In scraping information from the main (‘flow’) text of the decision, similar impediments arise, but on a much bigger scale. The flow text, written by individual judges, allows for much greater individual deviation of wording. There is no consistent structure to the flow text, making identification of outcomes and remedies challenging using a simple rule-based

³⁰ *UNISON, R (on the application of) v Lord Chancellor* [2017] UKSC 51. For a discussion of the impact on low-value claims, see Abi Adams and Jeremias Prassl, ‘Vexatious Claims: Challenging the Case for Employment Tribunal Fees’ (2017) 80 *Modern Law Review* 412.

approach. Take the value of any monetary compensation awarded by the tribunal as an example. This is not uniformly structured within the webpage or judgment document and so must be extracted from the text itself. Simply extracting the monetary values contained in the flow text is insufficient to identify compensation – some judgments include reference to salary levels, while others decompose any compensation into various sub-categories. This renders the identification of total monetary compensation challenging: a method that simply selects the maximum or first monetary value appearing in the flow text, for example, risks providing inaccurate and misleading results.

However, these challenges are increasingly remedied by advances in Natural Language Processing (NLP) techniques. NLP can help to overcome the difficulties associated with inconsistent formatting and heterogeneity in language, especially when combined with a supervised machine learning approach. Application of these methods typically requires a ‘training set’ of judgments that have been labelled with the features of interest (e.g. the value of any monetary compensation, whether a claimant has legal representation and so on). This allows for representation of the judgment text in a numerical matrix format on which an algorithm can then be trained to predict correctly key features from the judgment text. Once sufficient accuracy is achieved within the training set, the algorithm can be applied to extract the relevant features from all available judgments.

It is worth noting that the complex nature of tribunal judgments can still cause problems, even for sophisticated NLP-aided analysis. For example, a substantial number of cases have two or more claimants or one claimant bringing forward multiple claims. Hence, while one of these claims might result in a dismissal, others might succeed. While humans find it relatively easy to identify which parts of the judgment text correspond to the various complaints, this is more challenging for automated tools. However, these difficulties are not insurmountable and, as described in Section 2, have not stopped companies and researchers from drawing on this new resource. Further, a more standardised approach to judgment formatting would alleviate the remaining technical issues by improving the identification of relevant information on outcomes, compensation, and representation in a common structured format.

c. Predictive Analysis of Case Outcomes

In principle at least, the use of judgment data is by no means limited to the creation of descriptive statistics. Machine learning algorithms are increasingly applied to judicial precedents to *predict* the outcome of a case, or to identify which relevant legal provisions are triggered by a given set of facts, rather than merely describing the characteristics of parties. The current research frontier involves the use of ‘deep learning’ techniques to improve the

accuracy of predictive algorithms. A range of studies have applied such techniques to European Court of Human Rights, achieving an accuracy of approximately 80% for predicting the outcome of a case.³¹ It has proven harder to train models to predict the relevant legal issues from a set of facts; attempts to predict which (if any) ECHR article has been violated from the facts of European Court of Human Rights cases have only achieved accuracy of approximately 60%, though accuracy is evolving rapidly.³²

Even more ambitiously, a recent strand of the literature has begun to investigate the extent to which quantitative analyses of judgment text can elucidate legal reasoning. As Armour notes ‘[outcome prediction] may be useful for citizens (and practitioners) who wish to be able to assess their likely prospects of success in legal proceedings. However, methods that simply “black box” the prediction process are incapable of attracting normative significance, as the authoritative force of law comes from its character as reason-based norms’.³³

d. Representativeness

Before turning to a detailed treatment of potential legal and societal drawbacks in the application of machine learning algorithms to judicial precedents, it is crucial to stress an urgent need for greater transparency over which judgments are posted online. If the set of judgments posted online is not representative of all decisions, any conclusions about the composition of cases and any machine learning models trained on this body of precedent will be biased.

The process by which judgments are posted online is unclear. It is therefore important to assess whether all judgements are reported online, and if there are systematic patterns in the characteristics of decisions that are missing. To this end, we compared the number of judgements with the number of decided cases as reported in the Quarterly Tribunal Statistics. In the financial year 2017/18, the judgments of 73% of disposed cases appear online; the QTS report that 18,360 cases were disposed of compared to 13,483 online judgments. In 2018/19, 18,287 decisions appeared online compared to 25,877 in QTS (a coverage of 71%).

Figure 2 shows the quarterly number of judgments appearing online and the number of cases disposed of as recorded in the QTS: it is evident that online judgments are an incomplete

³¹ See discussion in John Armour, ‘AI and Judicial Precedents: A Review of the Literature’, February 2020 (on file with the authors).

³² Ilias Chalkidis, Manos Fergadiotis, Prodromos Malakasiotis, Nikolaos Aletras, & Ion Androutsopoulos, ‘Extreme ‘Multi-Label Legal Text Classification: A case study in EU Legislation’ (2019) arXiv preprint arXiv:1905.10892.

³³ Armour (n 30).

record of tribunal decisions. The impact of *UNISON* on case load, for example, appears much more quickly in the QTS. Closer inspection suggests that this is partly because of differences in the recording of ACAS conciliated settlements. In the QTS, ACAS conciliated settlements are fully included in the official statistics; in 2017/18, they made up 25% of disposed complaints. However, ACAS mediated decisions are very rare within the corpus of online judgments.

Figure 2. Comparison Between Online Judgments & Disposed Cases

Given how disposal statistics are presented in the QTS, it is not possible accurately to compare the scope and representativeness of online judgments to all disposals other than ACAS mediated settlements. Future research would greatly benefit from having outcome and jurisdictional information separately recorded for single and multiple claimant cases in the QTS to facilitate the representativeness of online judgment information for all cases decided.

Finally, it is important to stress that even if online judgments are representative of *decided cases* in the ETs, they are highly unlikely to be representative of all employment disputes arising.³⁴ This should be of concern to those hoping to advise potential claimants of their likelihood of success in court.

e. Summary

The online availability of employment tribunal judgments is already facilitating new empirical research and commercial products related to labour litigation. Small changes to the formatting

³⁴ Linda Mulcahy, 'The Collective Interest in Private Dispute Resolution' (2013) 33 *Oxford Journal of Legal Studies* 59.

of judgments, and greater transparency over the process by which judgments are posted online, would reduce barriers to the use of this resource even further. While this is to be welcomed in principle given the absence of timely official statistics and high-quality survey evidence, we must also recognise the potentially serious drawbacks which can ensue. It is to this question that discussion now turns.

4. DRAWBACKS

The problems with publicly available employment tribunal judgment data are not limited to technical and statistical issues. Whilst insights about the tribunals and their users at the system level are of considerable value to a range of stakeholders, the availability of fine-grained, personalized data can quickly become much more problematic from the perspective of the individual employee. As a result, considerations cannot be limited to selecting the data variables that will be most useful for shedding light on research and policy concerns. When thinking about what data to publish, we also need to bear in mind the potentially abusive purposes to which that information might be put, and against which adequate legal regulation has to provide protection. In this section, we identify potential abuses, ranging from blacklisting and skewed settlement negotiations to evasion and algorithmically-driven discrimination in hiring, promotion, and retention in order then to be able to highlight the challenges for existing legal regimes in employment law and beyond in developing appropriate safeguards.

a. Blacklisting

In the narrowest sense of the term, blacklisting refers to the practice of collating information on trade union members, and in particular activists, with a view to avoiding their hiring on future jobs.³⁵ This may be done by large employers individually, or by third party services collating industry-wide information, such as for example in the construction sector, on employers' behalf. Whilst historically particularly prevalent in industries such as construction and infrastructure, and with a focus on collective activities, blacklisting today can encompass a much broader set of informal practices designed to punish workers perceived as 'trouble-makers'. Indeed, there is growing evidence on black-listing practices in the context of data-analytics.³⁶ At present, the names of claimants and respondents are published in tribunal

³⁵ D. Smith and P. Chamberlain, *Blacklisted; The Secret War Between Big Business and Union Activists* (Oxford: New Internationalist, 2015).

³⁶ For examples of novel black-listing practices in the context of data-analytics, see: Newman, Nathan. 2017. 'Reengineering Workplace Bargaining: How Big Data Drives Lower Wages and How Reframing Labor Law Can Restore Information Equality in the Workplace'. U. Cin. L. Rev. 85: 693. For a discussion of blacklisting in the specific context of 'big data', see: Hu, Margaret. 2016. 'Big Data

judgments, and the .gov.uk website search tools allow anyone easily to find judgments associated with particular individuals or firms. It is therefore a *real* risk that employers or related third-party organisations might use this easy source of information to effectively blacklist employees based on their litigation history.

Not only does this practice allow employers to refuse to hire applicants on grounds that are irrelevant to their capacity to do the job, thereby introducing new forms of discriminatory (and arguably irrational) hiring, it would also potentially discourage employees from bringing employment claims in the first place, out of a fear that to do so might negatively impact their future job prospects. As Barnard, Ludlow, and Fraser Butlin have demonstrated, there are clear ‘fears about the consequences of raising a dispute [including in] particular concerns about the possibility of reprisal by way of being “blacklisted”’.³⁷ Such concerns are further exacerbated if the feared consequences of reprisal are extended beyond a claimant’s current workplace or local community to potentially any future employer, whether domestically or internationally: the publication of tribunal judgments might thus actually undermine the broader goal of facilitating access to justice. The concern is not new – over a century ago, the House of Lords noted ‘that the principle of open justice may be restricted where publication would “reasonably deter a party from seeking redress, or interfere with the effective trial of the cause”’.³⁸

b. Settlement negotiations

The availability of case-level information about the composition of claims and awards will be invaluable for facilitating detailed quantitative studies of the tribunal system, and, potentially, for developing policies capable of effectively targeting inefficiencies. In principle such information could be very useful for employees when deciding whether both to launch a claim and in any subsequent settlement negotiations.

Blacklisting’. *Florida Law Review* 67: 77. For a discussion of blacklisting practices in the UK more specifically, see: Chamberlain, Phil, and Dave Smith. 2016. *Blacklisted: The Secret War between Big Business and Union Activists*. New Internationalist; and on present-day concerns about blacklisting practices in the UK and elsewhere, see: Kessler, Sarah. 2020. ‘Companies Are Using Employee Survey Data to Predict — and Squash — Union Organizing’. Medium. 30 July 2020.

<https://onezero.medium.com/companies-are-using-employee-survey-data-to-predict-and-squash-union-organizing-a7e28a8c2158>; ‘New Report Recommends Public Inquiry into Blacklisting Scandal & Criminal Sanctions for Blacklisters’. 2017. IER. 14 December 2017. <https://www.ier.org.uk/press-releases/new-report-recommends-public-inquiry-blacklisting-scandal-criminal-sanctions-blacklisters/>.

³⁷ Catherine Barnard, Amy Ludlow, Sarah Fraser Butlin, *Beyond Employment Tribunals: Enforcement of Employment Rights by EU-8 Migrant Workers*, *Industrial Law Journal*, Volume 47, Issue 2, July 2018, 226, at 242.

³⁸ Emily Gordon Walker and Gus Baker, ‘Litigants anonymous’ (n) 9, citing *Scott v Scott* [1913] AC 417, 446.

However, the availability and statistical analyses of such data might also provide employers with the sort of information required to gain an upper hand in settlement negotiations.³⁹ Settlement negotiations in the employment context are potentially highly problematic: in a context of inequality of bargaining power, and with the looming stress of tribunal proceedings, settlement offers potentially enable employer to place undue pressure on vulnerable employees to give up their rights.⁴⁰ Given the superior resources at their disposal, it is likely that employers will in a better position to access and understand quantitative information on the chance of tribunal success and the likely level of financial remedies than are employees. This risk might be mitigated if a public dialogue about the costs and benefits of employment litigation could be stimulated, and data about these matters were made accessible, and digestible, to the wider public. However, we are currently far from this scenario.

c. Evasion

A related point arises in the context of ex-ante avoidance of employment law. As we have already seen, detailed information about tribunal decision-making processes could provide private sector start-ups with the data inputs they need to develop predictive tools, such as tools for predicting employment status.⁴¹ Nonetheless, there is a real possibility that this data, and such predictive tools themselves, might be used by employers to creatively draft complex contractual provisions with a view to better obscuring the employment status of their workers. While the ability to better predict the outcomes of legal claims is essential for ensuring clarity and certainty in the law, and thus, is an essential component of open justice; in the employment context in particular, there is a real risk that complete publicity might actually *frustrate* this goal, facilitating evasion if called upon by “armies of lawyers” contriving documents in their clients’ interests which simply misrepresent the true rights and obligations on both sides’, as the Employment Tribunal so starkly put it in *Aslam v Uber*.⁴²

d. Hiring Algorithms

³⁹ The QTS currently provides only annual statistics on the average level of compensation awarded for seven types of complaints.

⁴⁰ This is particularly so in the context of pre-termination negotiations which, since 2013, cannot be introduced as evidence into an unfair dismissal hearing, except if there is evidence of improper behaviour. This is, however, not likely to include attempts by employers to impose an unfair bargain on workers where no form of discrimination, aggression, or victimisation is employed.

⁴¹ Samuel Dahan, ‘Determining Worker Type from Legal Text Data Using Machine Learning. Pervasive Intelligence and Computing’ (2020) IEEE PICom.

⁴² *Aslam v Uber BV ET* [96], citing the observations of Elias J in *Kalwak*.

A final danger arises in the context of statistical and machine learning algorithms increasingly being used by employers to make hiring decisions.⁴³ These algorithms work by drawing on data about a given population with a view to identifying correlations between certain characteristics (and combinations of characteristics), and particular forms of behaviour. They thus generate *predictions* about how likely an individual is, given their characteristics, to engage in a specific behaviour in the future.⁴⁴ If details about claimant characteristics are made available to the public, then, there is scope for employers to rely on such techniques to predict how likely an individual is to take them to court (to challenge unlawful decisions, and/or engage in litigation in the future) with a view to refusing them employment on that basis.

Such practices would not only introduce new barriers to access to justice; they are also likely to exacerbate and/or reproduce existing inequalities. The reason for this is two-fold. First, algorithms embed, in code form, pre-existing human biases. Even where algorithms operate on the basis of clear decision-rules, because these rules are formed and conceptualised by humans, they can still operate in a biased way.⁴⁵ Second, where a particular group is under or over-represented in the data sample (the input data), as will inevitably be the case with data about tribunal claimants, the algorithm will tend to generate predictions that are biased against a particular group (or, very often, more than one).⁴⁶ This is exactly what happened with the now-infamous hiring algorithm used by Amazon: drawing on information about past

⁴³ J Adams-Prassl, ‘What if your Boss was an Algorithm? Economic Incentives, Legal Challenges, and the Rise of Artificial Intelligence at Work’ (2020) 41 *Comparative Labor Law and Policy Journal*

⁴⁴ Similar practices are already widespread, with some firms using them to identify candidates most likely to stay long-term in a job; to engage in fraudulent activities, notwithstanding no past record of misbehaviour; and/or to identify, in advance, any workers likely to quite their job. See: ‘JPMorgan Chase Develops “Early Warning System”’ <<https://www.bankinfosecurity.com/jpmorgan-chase-develops-early-warning-system-a-12855>> accessed 15 January 2020; ‘JPMorgan Algorithm Knows You’re a Rogue Employee Before You Do’ *Bloomberg.com* (8 April 2015) <<https://www.bloomberg.com/news/articles/2015-04-08/jpmorgan-algorithm-knows-you-re-a-rogue-employee-before-you-do>> accessed 15 January 2020. Jennifer Liu, ‘This Algorithm Can Predict When Workers Are about to Quit—Here’s How’ (*CNBC*, 10 September 2019) <<https://www.cnbc.com/2019/09/10/this-algorithm-can-predict-when-workers-are-about-to-quitheres-how.html>> accessed 15 January 2020; Eric Rosenbaum, ‘IBM Artificial Intelligence Can Predict with 95% Accuracy Which Workers Are about to Quit Their Jobs’ (*CNBC*, 3 April 2019) <<https://www.cnbc.com/2019/04/03/ibm-ai-can-predict-with-95-percent-accuracy-which-employees-will-quit.html>> accessed 15 January 2020.

⁴⁵ ‘Discrimination, Artificial Intelligence, and Algorithmic Decision-Making’ 51; Ulrich Leicht-Deobald and others, ‘The Challenges of Algorithm-Based HR Decision-Making for Personal Integrity’ (2019) 160 *Journal of Business Ethics* 377.

⁴⁶ Sandra Wachter and Brent Mittelstadt, ‘A Right to Reasonable Inferences: Re-Thinking Data Protection Law in the Age of Big Data and AI’ [2019] *Colum. Bus. L. Rev.* 494; Kristian Lum and James Johndrow, ‘A Statistical Framework for Fair Predictive Algorithms’ [2016] arXiv:1610.08077 [cs, stat] <<http://arxiv.org/abs/1610.08077>> accessed 22 January 2020; Betsy Anne Williams, Catherine F Brooks and Yotam Shmargad, ‘How Algorithms Discriminate Based on Data They Lack: Challenges, Solutions, and Policy Implications’ (2018) 8 *Journal of Information Policy* 78.

employees in which women were under-represented, the algorithm proved to discriminate against women in predicting an employee's suitability for the job.⁴⁷ While touted for their apparent objectivity, and freedom from human bias, then, these algorithms can actually perpetuate historical discrimination. Given the lack of transparency inherent in much of the technology deployed at the moment, furthermore, it may be extremely difficult to detect, yet alone prove, that an algorithm is discriminating in any one case.

To be clear, what is at stake here is not *only* the possibility of facilitating evasive practices that have long troubled the labour market; such is always a risk where individuals are aware of the law and have resources to deploy when it comes to creatively evading it. The risk is also that the publication of tribunal data might actively facilitate the abuse of individuals' personal data for the purposes of making decisions about them that may well have a significant and material impact on their every-day lives, and over which they have little or no control. Not only would such practices be inconsistent with the broader values, and goals, of labour law, they also fundamentally conflict with the very principles, of open and equal access to justice, on which the publication of that data is based. Before we can resolve the question of what data to publish, then, we need to think carefully about what safeguards might exist, and/or be necessary, to adequately regulate its use.

5. EXISTING SAFEGUARDS

In this section we consider the various safeguards that might be required to protect against some of the abuses of tribunal data identified thus far. For reasons of space, we focus attention on the safeguards that might be required to prevent the abuse by employers of the *personal data* of tribunal claimants, whether for the purposes of blacklisting, or as an input into hiring algorithms and machine learning systems designed better to predict employee behaviour.

In this vein, we identify four existing sources of protection, and expose their respective limits and potential: the Blacklisting Regulations 2010, anonymity orders, data protection rules, and equality law. The conclusion is relatively sobering: whilst the former two regimes could be expanded, in principle at least, to address some of the most grievous instances of tribunal judgment misuse, ultimately, the problem lies not just in the non-anonymised nature of the data that is publicly available, or even, the processing of that data by employers; the real problem lies in a lack of control, by individuals, over the inferences drawn from data, and over the decision-making processes that rely upon them. As such both data protection and

⁴⁷ See: 'Amazon Scraps Secret AI Recruiting Tool That Showed Bias against Women' *Reuters* (10 October 2018) <<https://www.reuters.com/article/us-amazon-com-jobs-automation-insight-idUSKCN1MK08G>> accessed 20 January 2020.

equality law, as presently conceived, are inadequate. The only way to adequately safeguard against abuse in this context is to regulate, and open to scrutiny, the process by which the machine-learning systems using various forms of input data are designed, and the process by which the recruitment decisions that make use of those systems are made.

a. Blacklisting Protection

There is an explicit statutory instrument to outlaw blacklisting in the narrow sense of that term: the Blacklisting Regulations 2010,⁴⁸ which stipulate that ‘no person shall compile, use, sell or supply a [blacklist]’, ‘compiled with a view to being used by employers or employment agencies for the purposes of discrimination in relation to recruitment or in relation to the treatment of workers.’⁴⁹ Upon closer inspection, however, of relatively little aid here: first and foremost, because its protective scope is limited to trade union membership and activities,⁵⁰ as well as lacking meaningful remedial teeth.⁵¹ Even in particularly egregious cases, technical questions such as employer status in multilateral contractual setups can quickly defeat the Regulations.⁵²

b. Anonymity Orders

Rule 50 of the 2013 Employment Tribunals Rules of Procedure provides for a limited exception to the publication in full of employment tribunal judgments:⁵³ a Tribunal may order, at any stage of proceedings, ‘that the identities of specified parties ... referred to in the proceedings should not be disclosed to the public, by the use of anonymisation or otherwise, whether in the course of any hearing or in its listing or in any documents entered on the Register or otherwise forming part of the public record’.⁵⁴ Such orders, however, are not to be granted lightly: the Rules of Procedure explicitly demand that ‘[i]n considering whether to make an

⁴⁸ The Employment Relations Act 1999 (Blacklists) Regulations 2010: SI 2010 No 493.

⁴⁹ Ibid, Reg 3(1), 3(2)(a).

⁵⁰ Ibid, Reg 3(2)(a)).

⁵¹ For a full discussion, see Charles Barrow, The Employment Relations Act 1999 (Blacklists) Regulations 2010: SI 2010 No 493, *Industrial Law Journal*, Volume 39, Issue 3, September 2010, Pages 300. See also Zahn, Rebecca (2014) Recent developments in blacklisting. *Greens Employment Law Bulletin*, 124, pp. 4, who concludes that ‘[t]he Regulations seem to have had little impact on restricting the practice of blacklisting.’

⁵² *Smith v Carillion (JM) Ltd* [2015] EWCA Civ 209, [2015] IRLR 467; see further A. Bogg, ‘Common Law and Statute in the Law of Employment’ (2016) 69 CLP 67, 109.

⁵³ Employment Tribunals (Constitution and Rules of Procedure) Regulations 2013, SI 2013/1237. The Rules are set out in Schedule 1.

⁵⁴ *ibid* Rule 50(3)(b).

order under this rule, the Tribunal shall give full weight to the principle of open justice and to the Convention right to freedom of expression.’⁵⁵

This emphasis on open justice, even at the expense of parties’ right to private and family life under Article 8 ECHR, has been consistently reflected in Tribunals’ approach to Anonymity Orders.⁵⁶ Despite early predictions that the ‘wider [online] availability of tribunal judgments should herald a greater willingness on the part of tribunals to make anonymisation orders’,⁵⁷ little seems to have changed in the balancing exercise between the right to a fair trial and freedom of expression in Articles 6 and 10 ECHR, respectively, on the one hand, and Article 8 on the other.⁵⁸

Whilst at first glance, the principles of access to justice and open justice seem to support unequivocally the wide dissemination of tribunal judgments, upon further inspection, the case is less clear-cut, not least when considering the serious implications of ‘automated blacklisting’. Potential anonymity, furthermore, is not just attractive for claimants – decisions might not infrequently criticise the behaviour of Respondent employees and managers, which may in turn impact the latter’s future job prospects:⁵⁹ spare a moment’s thought for Ms Bertram, Uber’s London General Manager, whose ‘grimly loyal evidence’ reminded the London Central Employment Tribunal ‘of Queen Gertrud’s most celebrated line: The lady doth protest too much, methinks.’⁶⁰

As Judge Eady QC (as she then was) explained in *Ameyaw*, however, ‘it is likely to be a rare case where other rights (including those derived from Article 8 ECHR) are so strong as to grant an indefinite restriction on publicity’.⁶¹ Indeed, in undertaking that balancing exercise Employment Tribunals were to remember that ‘the burden of establishing any derogation from the fundamental principle of open justice or full reporting lies on the person seeking that derogation’, and second, that ‘it must be established by clear and cogent evidence that harm will be done by reporting to the privacy rights of the person seeking the restriction on full reporting so as to make it necessary to derogate from the principle of open justice’.⁶²

⁵⁵ *ibid* Rule 50(2).

⁵⁶ See e.g. *BBC v Roden* [2015] ICR 985.

⁵⁷ Emily Gordon Walker and Gus Baker, ‘Litigants anonymous: the tribunal database and anonymity’ 2007 (24) ELA Briefing 8.

⁵⁸ *Re S (A Child) (Identification: Restrictions on Publication)* [2005] 1 AC 593, at [17].

⁵⁹ Stammeringlaw ‘Online Tribunal Decision Database, and Anonymity Orders’ (Online Blog, 15 September 2020) <https://www.stammeringlaw.org.uk/employment/resolving-employment-disputes/online-tribunal-decisiondatabase-anonymity-orders/>

⁶⁰ *Aslam v Uber* (ET) [87]

⁶¹ *Ameyaw v Pricewaterhousecoopers Services Ltd* [2019] UKEAT 0244_18_0401 (4 January 2019) [45].

⁶² *Ibid* [48]

In practice, this means that Anonymity Orders will only be granted rarely – even though, as the Employment Tribunal recognised in *Kirkham v UKRI*, there ‘will be many employees who will be cautious, if not reluctant, to issue a claim in the Employment Tribunal against a former employer in case they are subsequently identified as a “trouble-maker” and thus less suitable for employment.’⁶³ That said, the grant of such Orders is not unheard of. In *X v Y*, the Employment Appeals Tribunal granted an application, lodged after the ET’s judgment, to anonymise the parties’ details to protect the claimant’s personal details. At the same time, however, xx J emphasised that there was no ‘definite rule or principle of law that every time a judgment refers to a party’s transsexual status, or refers to sensitive mental health issues, the Tribunal should anonymise the names of the parties’, and refused to delete any substantive elements of the Tribunal’s judgment.⁶⁴

In any event, even where Courts agree to party anonymity, and subject only to national security exemptions, there has been strong resistance to arguments that a judgment not be published at all, or that particular details such as references to a claimant’s disabilities and their consequences, should be fully redacted. Whilst Bean LJ decided that it was ‘unnecessary to go so far as to say that there will never be a case (other than one concerning national security) in which an ET judgment can be kept secret by not being entered on the Register’, in *L v Q Ltd* he found ‘it hard to imagine the circumstances in which it would be right for an ET acting under Rule 50 to withhold publication of a judgment altogether.’⁶⁵

c. Data Protection Law

The GDPR (General Data Protection Regulations), as implemented by the Data Protection Act 2018, introduces a number of rights for individuals with regard to the processing of their personal data.⁶⁶ In this context, ‘personal data’ includes all information relating to a person from which they can be identified,⁶⁷ and processing includes a wide variety of activities, including the collection, recording, organisation, structuring, storage, adaptation or alteration, retrieval, consultation, use, disclosure by transmission, dissemination or otherwise making available, alignment or combination, restriction, erasure or destruction, of personal data.⁶⁸

⁶³ *Kirkham v United Kingdom Research and Innovation* [2019] UKET 2501482/2018 [44].

⁶⁴ *X v Y* [2019] UKEAT/0302/18/RN [43], [47].

⁶⁵ *L v Q Ltd* [2019] EWCA Civ 1417 (09 August 2019) [26].

⁶⁶ 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 (General Data Protection Regulation), OJ L 119/1, 4.5.2016 (‘GDPR’).

⁶⁷ Article 4(1) GDPR; Section 3(2) DPA 2018

⁶⁸ Article 4(2) GDPR Section 3(4) DPA 2018

By virtue of this regime, all data controllers must identify a lawful basis for their processing of personal data. In practice, express consent will often be relied upon to justify the processing, but in employment, the inequality of bargaining power between the parties makes this ground difficult to establish. As such, in the employment context, employers will likely be required to point to a legitimate interest that the processing serves, and for which the processing is necessary and proportionate.⁶⁹ The GDPR also provides certain rights to data-subjects vis-a-vis the processing of their personal data, which, in practical terms, means that all employers must inform job applicants of the categories of personal data that will be obtained, including data obtained from public sources, how that data will be used, and for what purposes. Individuals must furthermore be able to gain access to that data, and to make amendments where necessary.

Under the DPA, moreover, employers are required to have clear policies in place relating to the processing, and retention, of personal data, and a specific policy for the processing of ‘special category’ data, data which can only be processed on certain conditions (with the explicit consent of the data-subject, and/or for employment and social security purposes where authorised by law (equal opportunities monitoring etc)).⁷⁰ Notably, special category data includes data that would often be revealed in the course of employment claims, such as information about a person’s religion, political opinion, race, trade union membership, and/or sexual orientation.⁷¹

Blacklisting

The above provisions go some way towards protecting the interests of employees and prospective employees with regard to the processing of their personal data by employers. However, when it comes to blacklisting in the broader sense identified earlier, *viz* searching the names of job applicants against those listed as claimants with a view to potentially refusing them employment and the use of hiring algorithms to predict an individual’s likely litigation behaviour, its protective scope is limited. Any use of tribunal data about a specific claimant (searching a claimant’s name) constitutes ‘processing’ of *personal data* within the scope of the GDPR and DPA. Nonetheless, it might not be difficult for an employer to establish that the processing of such data is lawful, for such processing serves the employer’s legitimate interest in being able to assess a candidate’s ‘suitability’ for the job. It is irrelevant for Data Protection purposes that an employer is taking a candidate’s litigation history into account when assessing

⁶⁹ GDPR, Art 5(1)(b).

⁷⁰ Section 56(2) DPA 2018.

⁷¹ Article 9(1) and Recital 51 GDPR; Sections 10 and 11 DPA 2018.

‘suitability’. Data Protection law is concerned only with the lawfulness of the *processing*, and thus, the consultation of the database. Provided that the job applicant is informed that information from certain public data bases will be consulted as part of the recruitment process, and provided the employer has a clear policy as to the retention of any such data, such ‘processing’ will be lawful.⁷²

Hiring Algorithms and Behavioural Predictions

Where data about tribunal claimants is used as an input into a hiring algorithm, the GDPR and DPA will often not apply for the reason alluded to above: the data will often be anonymised, and so, not necessarily constitute *personal* data within the scope of the data protection regime. Employers may, for example, deploy a machine learning system programmed to draw inferences from statistics about tribunal claimants, with a view to constructing behavioural profiles. These profiles then allow employers to pin-point those applicants deemed likely to pose a high ‘litigation’ risk, and, potentially, to refuse such applicants employment on the basis of this prediction. The problem, however, is that such predictions can be made *even on the basis of anonymised data*.⁷³ Provided that the data is anonymised prior to its being used by the employer (or rather, the algorithm), then, it will fall outside the scope of the DPA and GDPR.

If, however, the input data was not anonymised (or had been anonymized, and thereby processed, by the employer), the GDPR and DPA *would* apply to its processing by the algorithm. The legislation would nonetheless *not* regulate the processing of the transformed data; that is, the *inferences drawn from that data*.⁷⁴ Thus, the reliance, by an employer, on a prediction generated by an algorithm using tribunal data as an input, such as a prediction that Asian Women are more likely to bring a claim in an employment tribunal, would not be ‘processing’ within the scope of GDPR or DPA, and so would not be something about which an individual applicant would have to be informed (nor something the accuracy of which he/she would have the right to rectify/challenge), nor would those *inferences* be ‘personal data’ to which the individual could request access (and, potentially, correct).⁷⁵

⁷² Wachter S and Mittelstadt B, ‘A Right to Reasonable Inferences: Re-Thinking Data Protection Law in the Age of Big Data and AI’ [2019] Colum. Bus. L. Rev. 494

⁷³ Williams BA, Brooks CF and Shmargad Y, ‘How Algorithms Discriminate Based on Data They Lack: Challenges, Solutions, and Policy Implications’ (2018) 8 Journal of Information Policy 78

⁷⁴ For an in-depth discussion of this issue, see: Wachter, Sandra and Mittelstadt, Brent, A Right to Reasonable Inferences: Re-Thinking Data Protection Law in the Age of Big Data and AI (October 5, 2018). Columbia Business Law Review, 2019(2). Available at SSRN: <https://ssrn.com/abstract=3248829>

⁷⁵ Even where inferences have been recognised as personal data, the courts have made it clear that if the inference or opinion came about in a situation where the individual is seeking to be assessed, as in a job application, to

True, if the decision to refuse employment on the basis of such a prediction was automated, the DPA *would* apply, because it establishes specific rules relating to automated decision-making processes. However, in such a context, the Act merely requires that an employer identify a legitimate interest for using an automated decision-making process (streamlining the recruitment process), that it inform the applicant that such a process will be used, and that it have a policy in place governing its use.⁷⁶ It does not, then, provide individuals with a right to challenge the content of automated decisions, nor the decision rules on which they are based- to challenge, in other words, the inferences drawn from tribunal data, and the particular use to which those inferences are put.

The problem here is that Data Protection Law has been conceived not as a mechanism for ensuring fairness and transparency in the decision-making processes that rely on, and analyse, personal data, but as a tool for individuals (data subjects) to assess whether the ‘input’ data was legally obtained, and whether the purpose for processing it was lawful. The CJEU has thus reiterated time and again, that Data Protection Law is not concerned with the *accuracy* (or fairness) of decisions and decision-making processes *involving* personal data, nor with making *those* processes fully transparent.⁷⁷ It is simply concerned with the legitimacy of the input stage of personal data processing.

Even if an employer can explain which data and variables have been used to make a decision, (such as tribunal decisions and details about litigation behaviour) the decision itself turns on inferences drawn from these sources, namely, that the individual is not a suitable employee. This, however, is an assumption or prediction about future behaviour that Data Protection Law does not allow us to verify, or refute.⁷⁸ Because many of the predictions

request access to, and to rectify, those inferences, would be anathema to the very purposes of the processing, and should not be allowed under the GDPR. The right to rectification is also limited by the requirement that it not affect the fundamental rights and freedoms of others, including IP and trade secrets (and, according to the recitals, the code underpinning the algorithm). For the most recent discussion of inferences, see: Case C-434/16 *Peter Nowak v Data Protection Commissioner*, 2017 ECR, I 994, particularly the Opinion of Advocate General Jott, 9-13. Joined cases C-141/12 *Y.S v Minister voor Immigratie, Integratie en Asiel* C-372/12 *Minister voor Immigratie, Integratie en Asiel v M. and S.*

⁷⁶ In this respect, the DPA is not as strict as the GDPR which requires that an individual be able to object to being subject to automated profiling and request that a decision be made by a human (Article 22).

⁷⁷ See, for example, Case C—553/07, *College van burgemeester en wethouders van Rotterdam v M.E.E Rijkeboer*, 2009, E.C.R. I-293, 48-52 and Case C-435/16 *Peter Nowak v Data Protection Commissioner*, 2017 E.C.R. (I-994, 54-55. See also: C-141/12 *Y.S v Minister voor Immigratie, Integratie en Asiel* C-372/12 *Minister voor Immigratie, Integratie en Asiel v M. and S.*

⁷⁸ Wachter S and Mittelstadt B, ‘A Right to Reasonable Inferences: Re-Thinking Data Protection Law in the Age of Big Data and AI’ [2019] *Colum. Bus. L. Rev.* 494

generated by hiring algorithms are simply probabilistic assumptions, moreover, even if a candidate was able to access the assumptions on the basis of which they had been refused employment, they would have a particularly difficult job in proving that they were *false*. If an algorithm were to establish that Asian Women were 80% more likely to engage in employment litigation than applicants without the characteristic of being female and Asian, for example, the real objection is not that this is empirically false (it would be difficult to say either way), but that a discriminatory decision is being made and a prediction that may well have no application to the individual job applicant is being used as a justification for refusing them employment.

It should be noted here that the above issues could not be solved by excluding from tribunal judgments certain sensitive categories of data (social category data). Not only would this be difficult to do in practice, given that such data is sometimes material to the particular legal issue in dispute (in the context of discrimination), it would also be largely ineffective. As Williams et al have shown, machine-learning techniques can learn *correlates*, or *proxies* for personal characteristics, generating behaviour predictions on the basis of such *proxies*, without the need for the characteristics of individuals to be explicitly stated.⁷⁹ As such, any included variables that are correlated with the protected variables will still contain information about the relevant characteristic.

d. Equality Law

A final alternative in regulating the use of tribunal data for the purpose of screening applicants would be to rely on the prohibitions on direct and indirect discrimination in the Equality Act 2010. By virtue of these prohibitions, employers cannot refuse to hire a potential employee because that person has, or is assumed to have, a particular protected characteristic, nor can they adopt job criteria that have the effect *in practice* of putting persons that share a particular protected characteristic at a disadvantage compared with those who do not (unless that criterion can be shown to be a proportionate means of achieving a legitimate aim).⁸⁰

Direct discrimination, in relation to an employer's use of hiring algorithms, might occur where an algorithm predicts, on the basis of that (and potentially additional) data that Asian people are 90% more likely to engage in litigation than other groups. If an Asian applicant were refused a job on the basis of this assessment, we might argue that the person was refused

⁷⁹ Betsy Anne Williams, Catherine F Brooks and Yotam Shmargad, 'How Algorithms Discriminate Based on Data They Lack: Challenges, Solutions, and Policy Implications' (2018) 8 Journal of Information Policy 78.

⁸⁰ Sections 13 and 19 Equality Act 2010.

the job *on grounds of race*. That is, but for the protected characteristic of race, the applicant would have been offered the job (or would not have been excluded from the running).

Things are more difficult, however, where the decision to exclude the applicant is made by the algorithm (via an automated decision).⁸¹ First, while employers can be questioned about the reasons for a particular decision, the decision-rules of algorithms may be tightly protected by IP. It would thus be difficult to prove that the *reason* for the decision was race. This is particularly so given that, in practice, hiring algorithms, and other forms of ‘profiling,’ rarely pin-point *a single* characteristic. That is, they rarely identify a correlation between just *one* protected characteristic and the particular behaviour in which the employer is interested (such as a tendency to make use of the employment tribunal system or seek the enforcement of employment rights). Drawing on a variety of data sources and data points, a more likely conclusion would be (for example) that Afro-Caribbean Women who did not go to University, and who live in Northern England, demonstrate an 80% higher risk of engaging in litigation than persons *without this combination* of characteristics. In this situation, the law relating to direct discrimination would be of little assistance. UK law does not allow for multiple factor discrimination; if an Afro-Caribbean Woman were refused a job based on the algorithm’s prediction, she could not prove that but for being Afro-Caribbean, or but for being female, she would have got the job. As such, a direct discrimination claim would fail.⁸²

This observation highlights a further limitation of the Equality Act: it only applies to discrimination on the basis of *listed* characteristics. In the above example, then, discrimination based on where one lives, or one’s educational background, would not be covered, even if one *could* show that *but for* that characteristics, the applicant would have got the job. Where, moreover, certain characteristics are used as *proxies* for protected characteristics (such as a claimant’s postcode being a proxy for race), it would be extremely difficult to make out the basic requirements for a claim: the refusal of employment is not based on a prohibited characteristic but something which, on its face, has nothing to do with those characteristics at all.

The prohibitions on indirect discrimination are slightly more promising. Indirect discrimination in this context might arise where an employer adopts a policy of refusing jobs to any person who, according to the hiring algorithm, poses a litigation risk of 80% or above. If this apparently neutral policy has the effect, in practice, of excluding persons who share a

⁸¹ Though this is itself problematic: GDPR, Art 22.

⁸² There already exists provision in the Equality Act for ‘dual factor’ discrimination and this would go some way towards addressing some of the issues identified: Section 14 Equality Act 2010.

particular protected characteristic from the chance to get the job, a claim in indirect discrimination would lie, unless that policy could be justified.

If, however, an automatic filtering system were to be implemented whereby persons posing (according to the algorithm's prediction) a litigation risk of 80% or above were *automatically* excluded, things would be more difficult: there is no legal *person* responsible for the disadvantage. Moreover, while an alternative route might be to challenge the decision *by the employer* to use an automatic filtering system, or to adopt a criterion of litigiousness, on the basis that this is itself a policy that disadvantages a particular group; given that algorithms are continuously processing data and updating their predictions accordingly, it would be difficult to show that *a particular protected group was systematically disadvantaged* at a given time.

* * *

In conclusion, then, it appears that neither blacklisting protection and anonymity orders, nor data protection and equality law suffice to address the issues we identified in the previous section. The former are limited in their scope and practical use, whereas the latter appear ill-equipped to deal with the complexity of discriminatory hiring practices, based on a *finite* list of protected characteristics, a narrow range of prohibited conduct, and the inability to challenge the various decision-making (automated and human) processes that data processing facilitates.

6. CONCLUSION

The above discussion reveals both the real potential of the systematic publishing of employment tribunal decisions online, and the not insignificant risks associated with such publication given its potential uses and abuses. In concluding, we suggest some steps that might be taken with a view to minimising these risks without, however, undermining the potential benefits available.

The first, and simplest, suggestion, is the systematic anonymization of the published data. This would be the most effective remedy against blacklisting, for example, for it would prevent data about specific individuals being used as a basis for employment and hiring decisions, while also avoiding a number of the challenges associated with Data Protection.⁸³ Such an

⁸³ Unfortunately, the questions of anonymity and online reporting of judgments were not included in the terms of reference for the Law Commission's recent report on Employment Law Hearing Structures Law Commission Report No 390, 'Employment Law Hearing Structures: Report' (2020) HC308.

approach would not, moreover, significantly restrict the data to be published, nor its potential benefits.

Beneficial uses of employment tribunal data very rarely, if ever, require the names of the parties. Rather, three different types of information are required: information about the management and composition of claims and awards (the category of claim, the outcome of the various complaints the type of remedy and its magnitude, and/or success rates); information about the characteristics of claimants and respondents (including size of business and sector, or race, sexuality, gender, respectively), and finally, information about the tribunal decision-making process itself (factors taken into account and/or given most weight in deciding different legal questions etc). All of these could be readily published without the need to include the names and/or addresses of claimants, and so, would not be hampered by a general policy of anonymity.

It is also worth noting that publication of *anonymized* data would not be contrary to the principle of open justice; indeed, it would allow for a fairer balance between open justice and the importance of protection of personal data and privacy. Of course, there is a public interest in exposing to scrutiny the way in which different legal decisions are made, and legal principles applied; in allowing the public to gain insight into how different categories of claimant are treated, and different categories of cases decided, and in opening up to the public the reasons given for those decisions. But none of these benefits *necessarily* require that the name and personal details of individual claimants and defendants be published online, or, at least, not as a matter of course. A number of jurisdictions committed to open-justice already adopt a much more fine-grained approach to anonymization, drawing a distinction between legal, and natural, persons, for example; between persons professionally involved in a case (judges, barristers, clerks etc) and those not so-involved and/or between cases that do and do not involve public figures.⁸⁴ In so doing, they allow for names to be published where a clear, and specific, public interest can be established, without, however, insisting on publication as a matter of course. It should also be noted that just because a judgment is anonymised for the purposes of online publication, does not mean that the names of defendants and claimants might not be available in different formats, and for different purposes, and/or at specific request. It would be possible, and arguably necessary, to regulate access to this personalised data, in a way that better balances the requirements of data protection and open justice, and to do this entirely independently from the online repository of anonymised judgments.

Despite these advantages of anonymization, it is important at the same time to remember that anonymization alone will not be sufficient to address the full range of concerns identified. Nor

⁸⁴ Marc van Opijnen and others, 'On-Line Publication of Court Decisions in the EU' the European Union 178, 24.

is it enough simply to restrict the sort of data that is made available in tribunal judgments, by excluding, for example, social category data from the public domain. Indeed, to exclude such data not only risks simply exacerbating the risk that AI systems will learn inadequate proxies for such data, introducing new forms of discrimination, but there is the added risk that removing such categories of data would actually take the published data outside of the scope of many of the provisions of the GDPR and DPA entirely.

In order to ensure that the risks associated with the publication of even anonymised data do not outweigh the potential benefits, then, a number of other, more far reaching, steps should be considered in the long run.

First, it will be imperative to add a claimant's litigation history to the list of prohibited grounds on which individuals cannot be refused employment.⁸⁵ In this way, it would be possible to regulate *directly* the employer's decision-making process, and its attempt to *exploit* data about an individual's litigation history in the context of recruitment.

Second, we need to open up the list of protected characteristics in the Equality Act to include 'analogous grounds', thereby allowing flexibility in the event that new characteristics emerge as potential grounds for discriminatory decision-making as decision-making processes become more complex and elaborate with the result of improved access to fine-grained data.

Third, and in more general terms, it will be necessary to make it explicit that the 'persons' to whom anti-discrimination provisions apply include persons legally responsible for the decisions of algorithms – the employer, or user, of that algorithm – so as to ensure that the provisions of the Equality Act can more effectively be used to counter-act possible side-effects in terms of potential discrimination.

Finally, as part of a more ambitious approach to the regulation of artificial intelligence more generally, much more comprehensive steps will need to be taken to regulate the use of judgment data. This might mean, effectively, opening up the 'black box' of hiring algorithms, subjecting particular decision-rules as well as their role in the exercise of managerial prerogatives, to scrutiny. This would require not just transparency in the processing of data, but transparency in the construction of the underlying systems and the decision-making processes to which they are applied.

⁸⁵ At present such grounds are limited to the list of protected characteristics in the Equality Act, and trade union membership.

The publication of an online repository of Employment Tribunal decisions marks a watershed in open justice. Present discussion has primarily focused on the benefits to academic research, though the potential upsides are by no means limited thus. At the same time, however, it is important to remember the potential downsides outlined. Concrete steps need to be taken to strike an appropriate balance between these competing considerations, and ensure that this resource is not used in a way that is positively harmful to the very persons whose interests the Employment Tribunal system exists to uphold.