Italian Labour Law e-Journal
Issue 2, Vol. 17 (2024)
ISSN 1561-8048
https://doi.org/...

Differential Privacy and Collective Bargaining over Workplace Data

Sandy J.J. Gould*

Abstract

The datafication of workplaces permits employers an increased information advantage over workers during bargaining. Instrumenting tools (and workers themselves), aggregating data and applying sophisticated analytic techniques can give employers greater insight into what is happening in workplaces. Data privacy laws like GDPR provide protections for individual workers but can make it more difficult for worker representatives to access workplace data for collective bargaining because employers can reasonably argue that releasing such data would put them in breach of their legal responsibilities in relation to data privacy. Aggregate summaries of datasets provided by employers would comply with data privacy laws, but are susceptible to manipulation. I argue that using differential privacy, a technique for processing data that makes it harder to determine who contributed data to a dataset, would remove an obstacle to employers sharing workplace data with worker representatives.

Keywords: Datafication, workplace data, collective bargaining, data privacy regulations, differential privacy.

1. Introduction

Using technology to monitor the productivity of individual workers is regulated to a significant degree in several jurisdictions.[1] Where it isn't, collective bargaining over monitoring has been proposed as a way for workers to exercise some control over what gets collected and how it gets used.[2] Colclough proposes a four-stage collective bargaining model

* Reader (Human-Centred Computing) at the School of Computer Science and Informatics, Cardiff University.

[1] See the case of Spain: Rodríguez Fernández M.L., *Collective bargaining and AI in Spain*, in Castillo A.P.D. (ed), *Artificial intelligence, labour and society*.

[2] Cecchinato M.E., Gould S.J.J., Pitts F.H., *Self-Tracking & Sousveillance at Work: Insights from Human-Computer Interaction & Social Science*, in Moore P.V., Woodcock J. (eds), *Augmented Exploitation Artificial Intelligence, Automation and Work*, Pluto Press, 2021; De Stefano V., Taes S., *Algorithmic management and collective bargaining*, in *Transfer: European Review of Labour and Research*, 29, 1, 2023, 21–36.

covering data collection, data analyses, data storage and data-offboarding.[3] Trade unions and worker representatives, Colclough suggests, need to be alert to individuals' data rights being compromised at any of these stages; getting overly focused on any particular stage will leave workers exposed. Aleks et al. enumerate some of the challenges facing trade unions in engaging in bargaining on this topic, noting that often wages and job security have been at the top of workers' agendas (especially in the UK context, where both have been eroded over the last decade), and data gets forgotten.[4] They suggest several practical strategies ranging from job mappings in relation to datafication to bargaining for formal protocols for evaluating the functioning of a particular algorithm. Unions have themselves investigated digital workplace surveillance; the UK TUC conducted a large-scale survey of workers,[5] and it has continued to document the rising use of these technologies.[6]

The extent to which various protocols and strategies are effective for bargaining over digital monitoring in particular contexts is not something that we know much about. Rather than speculating on strategies for bargaining, in this paper, I speculate on how worker representatives might go about removing barriers to collective bargaining over workplace monitoring. One perennial challenge for worker representatives is getting hold of relevant workplace data to form a basis for bargaining.[7] Many organisations are collecting more data about staff than they should – specifically violating Article 5(1)(c) of GDPR, which requires the principle of data minimization to be upheld by employers.[8] Even if superficially anonymized, amassing data without good reason has the potential to dilute protections for individuals and their data.

Perhaps differential privacy could help? This is a statistical technique that adjusts datasets so that analyses remain consistent, even as individuals are made very difficult to identify.[9] It is advantageous over simple aggregations of data (e.g., averages) because it permits stakeholders with different goals and priorities to perform their own analyses without having to accept theory-laden aggregations from the party with control over the raw data. It also means that no parties can discuss 'anonymous' data while being able to know implicitly who individual contributors might be.

[3] Colclough C.J., *Righting the Wrong: Putting Workers' Data Rights Firmly on the Table*, in Graham M., Ferrari F. (eds), *Digital Work in the Planetary Market*, The MIT Press, 2022, 291–302.
[4] Aleks R., Maffie M.D., Saksida T., *The role of collective bargaining in a digitized workplace*, in *Reimagining the governance of work and employment*, 85, 2020.
[5] *I'll be watching you | TUC*, 2018, https://www.tuc.org.uk/research-analysis/reports/ill-be-watching-you, accessed 18 October 2024.
[6] *What is monitoring at work? | TUC*, 2022, https://www.tuc.org.uk/guidance/what-monitoring-work, accessed 18 October 2024; Allen J., *Wales TUC: Worker surveillance on the rise in Wales*, 2021, https://www.tuc.org.uk/news/wales-tuc-worker-surveillance-rise-wales, accessed 18 October 2024.
[7] Sherman H.L.Jr., *Employer's Obligation to Produce Data for Collective Bargaining*, in *Minnesota Law Review*, 35, 1, 1950, 24–46.
[8] *HDPA (Greece) - 31/2023*, https://gdprhub.eu/index.php?title=HDPA_(Greece)_-_31/2023, accessed 18 October 2024; *CNIL (France) - SAN-2020-003*, https://gdprhub.eu/index.php?title=CNIL_(France)_-_SAN-2020-003, accessed 18 October 2024.
[9] Dwork C., *Differential Privacy*, in Bugliesi M. and others (eds), *Automata, Languages and Programming*, Springer Berlin Heidelberg, 2006, 1–12; Dwork C. and others, *Calibrating Noise to Sensitivity in Private Data Analysis*, in Halevi S., Rabin T. (eds), *Theory of Cryptography*, Springer Berlin Heidelberg, 2006, 265–284.

https://doi.org/10.6092/issn.1561-8048/14085

This paper makes the case for differential privacy as a tool for collective bargaining over workplace data. I will describe the problems associated with collective bargaining over data collection, introduce differential privacy, and explain the ways in which it could help to solve some contemporary challenges for worker representatives bargaining over workplace data collection.

2. Challenges of collective bargaining over data

Access to data about employers and workplaces has always been important for collective bargaining. Information about an employer's financial situation, pricing strategies and macro measures of worker productivity would be put into the broader economic context to help worker representatives form a bargaining strategy.[10] The structure of organizations and the nature of their governance means there is normally an imbalance between the information available to employers and the information available to workers and their representatives that favours employers. The datafication of workplaces using digital technologies has tilted the balance further in employers' favour: they have an even more fine-grained view of their operations and access to new analytic tools to make sense of patterns.[11]

Let's consider an example to illustrate how the growing imbalance in information between employers and worker representatives might influence collective bargaining. Imagine a set of production lines for bottles of beer. At a quarterly joint meeting, the employer brings production averages for each line and each shift. A union with sufficient density and wherewithal brings its own estimates, because workers can themselves observe how each line and shift is running. It'd be impractical to count every bottle coming off every line, but it would be conceptually possible with a sufficiently dedicated workforce. The nature of the work and the nature of the way that work is measured in such a context limits an employer's capacity to obfuscate.[12] Contrast this with a software engineering company where engineers' machines are instrumented to record twenty different things about workplace behaviour, and these measures are fed into a machine classifier that ranks the output of the various teams. The quarterly join meeting comes around; the employer gets to choose how to draw on their panoply of data, data which represents things that no human could ever observe. How many key presses did a given engineer produce last week? What was their affect, as measured by facial expressions, and how did that correlate with the number of emails that they sent? Technology lets the employer measure a huge amount about their workforce. Can worker representatives even know what is produced, given that the outputs of the business are not something tangible? How will they be able to challenge the employer's version of events when there is no way for them to arrive at their own shadow representation of what they think is happening?

[10] Sherman H.L.Jr., nt. (7).
[11] Aloisi A., *Regulating Algorithmic Management at Work in the European Union: Data Protection, Non-discrimination and Collective Rights*, in *International Journal of Comparative Labour Law and Industrial Relations*, 40, 1, 2024.
[12] The idea of a production line whose output could be obfuscated is an interesting one. How could the number of items produced in a factory be successfully hidden from workers? A kind of production line panopticon where no set of workers can know how much is being produced.

https://doi.org/10.6092/issn.1561-8048/14085

Employers often have a vested interest in withholding data, both to enhance their bargaining positions with their own staff,[13] and to limit vectors for sensitive data reaching competitors, media or regulators. There is an argument that data privacy laws make it easier for employers to refuse to share data because of the risk of exposing individual staff in any data they release. While the employer may have had a legitimate purpose for collecting data under Article 5(1) of GDPR, they would need a lawful basis for sharing personal data with, say, a trade union (And individual workplace telemetry would be personal data). Such sharing would require, at least, explicit consent from workers. What incentive would an employer have as a data controller to gain such consent? It might very well suit an employer to keep a trade union from having access to data collected through workplace monitoring.

Richards describes the challenge presented by the co-option of privacy in this way, noting that GDPR provides "opportunities for mischief as privacy rules are used pretextually to serve other ends."[14] Calacci and Stein suggest that the focus of data privacy legislation on individual data subjects is 'myopic'.[15] They relate the law surrounding privacy at work to the broader collective importance of data in workplaces. Their argument is that there is a balance to be struck between these competing imperatives. By focusing on individual data privacy and not on collective rights to data, individuals are still disadvantaged by proxy. Employers are able to develop an understanding of the cohort-level behaviours of workers, but workers cannot. In this sense, workers' individual privacy rights are leveraged to disenfranchise workers as a collective.

The aim of protecting individuals through data privacy laws is understandable. Even in instances where data sharing has been secured with the goal of helping disadvantaged groups, it has had the opposite effect.[16] Consider an example scenario where an employer measures five things about one hundred workers' productivity. These measures go into a database and are used for performance evaluations. Simply stripping names of this database does not automatically render the data anonymous (i.e., not personal) under GDPR. This is because Article 4(5) specifically covers instances of pseudonymization: "[this] means the processing of personal data in such a manner that the personal data can no longer be attributed to a specific data subject without the use of additional information, provided that such additional information is kept separately and is subject to technical and organisational measures to ensure that the personal data are not attributed to an identified or identifiable natural person".

Finck and Pallas provide a comprehensive analysis of the relationship between GDPR and the potential for individuals to be identified in datasets that appear to have been

[13] Chung R. and others, *Do Managers Withhold Good News from Labor Unions?*, in *Management Science*, 62, 1, 2016, 46–68.
[14] Richards N., *The GDPR as Privacy Pretext and the Problem of Co-Opting Privacy The Internet and the Law: Legal Challenges in the New Digital Age: Essays*, in *Hastings Law Journal*, 73, 5, 2022, 1511.
[15] Calacci D., Stein J., *From access to understanding: Collective data governance for workers*, in *European Labour Law Journal*, 14, 2, 2023, 253–282.
[16] Hancock J. and others, *The tensions of data sharing for human rights: A modern slavery case study*, in *The 2024 ACM Conference on Fairness, Accountability, and Transparency*, 2024, 974–987.

https://doi.org/10.6092/issn.1561-8048/14085

anonymized.[17] They describe the often-contradictory judgments that surround the question of anonymity – some courts have determined that it must be *impossible* to identify individuals from datasets. Others have taken a less monolithic approach to the question (recognising that 'impossible' is, a security researchers would make clear, a word that should be used with special care). Finck and Pallas relate computing techniques for anonymizing data under GDPR. By considering the information content of data that has undergone processing from a technical perspective, it becomes clear that "[to] achieve absolute protection, any processing of personal data would have to be outlawed."[18] They emphasise the importance of a risk-based approach to anonymizing data rather than an absolute, categorial approach that renders data as either wholly anonymous or completely identifiable.

An employer might argue that this privacy conundrum could be solved if they were to provide aggregate summaries of the data they hold. On the surface, this seems like an effective argument. If a group of worker representatives are collectively bargaining, why is it necessary for them to identify individuals in a sample? Why not just use summaries of collective behaviour? This does little to address the informational imbalance between employers and worker representatives, however, and may even increase it. Aggregations necessarily mean a reduced information content. But the actor that chooses which aggregations to perform and on which data to perform it is critical to the story that aggregations tell. Aggregations can be manipulated.

Anscombe's quartet[19] illustrates the limitations of summary statistics. Figure 1 shows four separate datasets that Anscombe developed. The mean x and y values of these datasets are identical in all four facets (9 and 7.5, respectively). Other properties (e.g., variance) are also identical in all four facets. Are these datasets representing the same things? No: once the raw data is plotted, it becomes clear that we are dealing with very different phenomena.

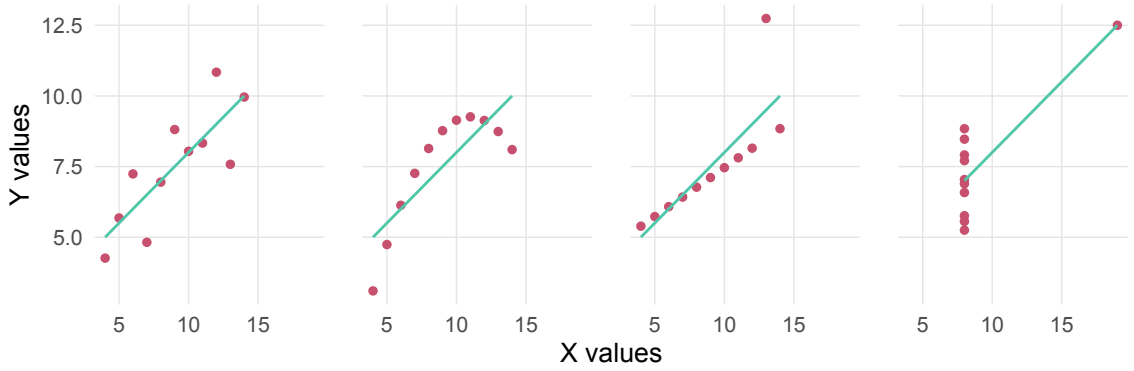

Figure 1: A plot of Anscombe's quartet. Each of the four datasets has the same properties in aggregate, but it is obvious that they are representing radically different things.

The datasets developed by Anscombe were synthetic – made-up to make a point. For Anscombe, the point was the importance of plotting data to understand the effects of outliers. Here, the point is slightly different; it's that in collective bargaining, which is often

[17] Finck M., Pallas F., *They who must not be identified—distinguishing personal from non-personal data under the GDPR*, in *International Data Privacy Law*, 10, 1, 2020, 11–36.
[18] ibid 34.
[19] Anscombe F.J., *Graphs in Statistical Analysis*, in *The American Statistician*, 27, 1, 1973, 17–21.

https://doi.org/10.6092/issn.1561-8048/14085

adversarial in nature, worker representatives having access to aggregate summaries of a dataset may not yield a comprehensive picture of what is happening in the workplace.[20] The two-dimensional nature of Anscombe's quartet really exacerbates the problem with summary statistics, but the problem is common in all datasets.

Let's imagine a call centre operation with four sites: Aberystwyth, Brynmawr, Caerphilly and Dolgellau. Some of the staff in Brynmawr have complained to their trade union representatives that there is a technical issue with how calls are being routed, which means some operators' wait times are being exaggerated while others are being underestimated: this is not fair. The management respond by reporting that their analysis shows that there is no problem, and that all four sites have identical call wait data, with a mean wait of 5-minutes and a standard deviation of 2-minutes. It seems that perhaps the operators' anecdotal reports are just that until the trade union receives an anonymous email with Figure 2 attached.

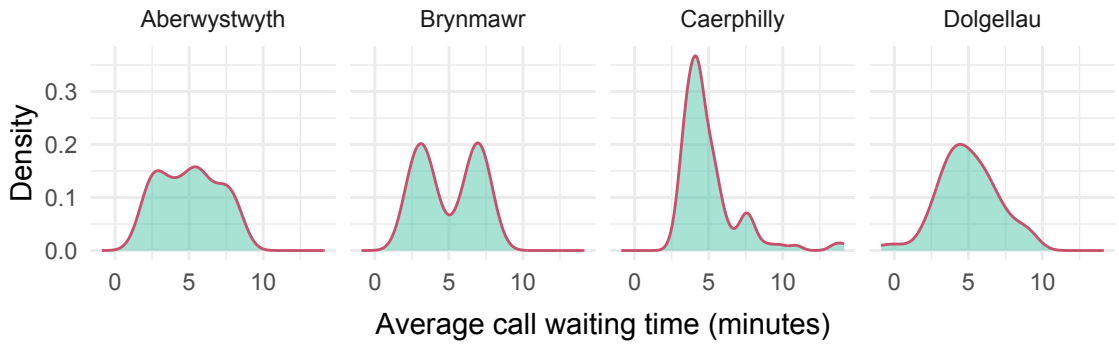

Figure 2: Summary data is not sufficient to understand what is happening in a dataset. Each of these fictional call centre sites have the same aggregate call waiting time: 5-minutes (SD=2-minutes). Plotting the underlying data tells a very different story, though: there are radically different profiles of call waits across the four sites.

It is clear from Figure 2 that although the summary data suggest four similar sites, digging into the data shows them to be very different. We can imagine this scenario playing out with groups and individuals. Each Key Performance Indicator (KPI) provides a new opportunity to tell a story, so by being selective with how each is reported, the overall narrative can be controlled. It is possible to construct the summaries that are provided to highlight certain aspects of the dataset while minimizing others.[21] This puts worker representatives at an essential disadvantage, and does not represent a 'compromise' position because it may be

[20] And this is just an expression of the status quo in all aspects of labour. As Adams and Wenckebach note, workers "suffer a systematic disadvantage when it comes to in influencing the various decision-making processes which ultimately impact on the organisation of work" (Adams Z., Wenckebach J., *Collective regulation of algorithmic management*, in *European Labour Law Journal*, 14, 2, 2023, 211–229.) This is a further, specific manifestation of a structure of the labour market, and Adams and Wenckebach explore the specific challenges of collective bargaining over –in their case– algorithmic management.

[21] And unfortunately 'plot the data' does not provide a simple solution to this problem. Consider a different version of Figure 2 where the continuous probability density plot is replaced with a histogram in which the data are 'binned' into ranges (say 0-2, 2-4, 4-6,6-8,8-10). Suddenly the details shown by a continuous plot are subject to another form of summarisation, and therefore potential obfuscation. Plots are just another form of summarization and can equally be bent to a particular narrative.

https://doi.org/10.6092/issn.1561-8048/14085

more misleading than having no data at all. A happy balance between maintaining individual privacy rights and workers' capacity to bargain collectively is difficult to find.

The problem is clear, then: in workplaces that have undergone datafication, data defines what work takes place, how it takes place, who does it, and how they have performed. To understand what is happening in a workplace, it is necessary to have some insight into that data (Because it underpins the decision-making of the employer). Employers have an interest in selectively sharing aggregate descriptions of data with worker representatives, and the individualizing nature of data privacy legislation gives them reasonable cover to do so. Where aggregated summaries are provided by an employer, there is potential for them to be misleading.

There is a way, I think, to remove some of the employer 'cover' provided by data privacy laws: differential privacy. This way of handling data would mean employers could share more than just summaries without falling foul of data privacy legislation. It does not *compel* them to share, but it would remove the cover provided by the law.

3. Differential privacy for collective bargaining

One way of avoiding leakage of data about individual workers, either directly through personally identifiable records or indirectly through pseudonymous ones, is to make use of differential privacy. This is a statistical technique for altering datasets so that the patterns within a dataset are maintained *in aggregate* while limiting the capacity of someone with access to the dataset to identify a given individual in it.

Consider again the example scenario of productivity data introduced earlier. As I noted, simply dropping names off the dataset does not make it sufficiently anonymous that an employer could share it with, say, a trade union. Through inference, the names could be 'reattached' to the 'anonymized' dataset. This is what the provisions of GDPR are designed to guard against. But what if the dataset could be 'scrambled', such that it is no longer possible to make inferences about individuals in the dataset while aggregate summaries were still accurate? This would permit worker representatives to conduct their own analyses, build their own narratives about what the data shows, and bring them to collective bargaining. This is, in essence, what differential privacy offers.[22]

Dwork et al. did much of the work to extend older ideas about data leakage from databases in codifying differential privacy.[23] Dwork notes the impossibility of preventing *any* inferences from being drawn from data, especially in the presence of auxiliary information (i.e., "information available to the adversary other than from access to the statistical

[22] In this paper I consider bargaining with a single employer. As Adams and Wenckebach note, in the UK context, "multi-employer bargaining is now largely absent" (p.215) and this is the context that I am familiar with. Conceptually these tools could be applied, to a degree, to multi-employer bargaining scenarios. Without data sharing, each employer would need to produce its own output from a differential privacy system. The commensurability of the output would be contingent on the value of ϵ that was chosen – different values of ϵ provide different levels of guarantee of anonymity (i.e., noise, ultimately) and so influence the degree to which the outputs from different employers could be compared.

[23] Dwork C., nt. (9); Dwork C. and others, nt. (9).

https://doi.org/10.6092/issn.1561-8048/14085

database"[24]). In the context of employment, it's likely that employers and worker representatives will have access to such auxiliary data. Instead, Dwork proposes that the risk to an individual being exposed by a dataset "should not substantially increase" because of their presence (or absence) in a dataset.[25] The technique protects individuals from deanonymization without meaningfully affecting the overall results obtained when producing summaries like means or medians. They cover multidimensional data and so could be used to maintain anonymity across, for example, a set of KPIs. The level of protection provided by these techniques is defined by ϵ, with lower values of ϵ representing a lower risk of a given individual's contributions to a dataset being identifiable.

The capacity of differential privacy to help prevent anonymization from being "undone" makes it useful for compliance with data privacy legislation. Huang et al.[26] note that the EU has explored differential privacy as a mechanism for providing an objective criterion of privacy preservation in datasets. They discuss the technique in relation to data privacy legislation, showing how it allows parties (employers, workers, worker representatives) to agree an acceptable 'privacy budget' in advance of any analysis that meets regulatory standards and then make data available accordingly, potentially choosing 'riskier' but more accurate datasets or less risky and less accurate datasets, depending on the nature of the data contained within them. Determining which value of ϵ is acceptable can be difficult, as Huang and Zheng point out: getting participants in datasets to vote on an acceptable value will itself provide auxiliary information that can be used to help identify those same individuals in datasets. Collective bargaining provides a good way of determining an acceptable value for ϵ, with workers having indirect input into an agreed value without having to put themselves at risk of leaking auxiliary information.

The critical thing about differential privacy is that it lets a given actor produce their *own* aggregate analysis of a dataset without being able to make inferences about who contributed to the dataset (i.e., which workers' data might be included). For workplaces that have undergone datafication, it is essential that worker representatives are able to conduct their own independent analysis of workplace data in order to effectively represent workers. Differential privacy techniques permit this to happen without risking the unnecessary leakage of information about individual workers (and preventing this concern from being used to obstruct the ability of worker representatives to understand the operation of the workplace). The challenge is making sure that stakeholders in a given dataset have a sufficiently well-developed understanding of how such techniques work and what their limits are, especially where those stakeholders are heterogeneous in terms of their skills and knowledge.[27]

[24] Dwork C., nt. (9), 2.
[25] ibid.
[26] Huang T., Zheng S., *Using Differential Privacy to Define Personal, Anonymous, and Pseudonymous Data*, in *IEEE Access*, 11, 2023, 109225–109236.
[27] Klymenko O. and others, *Understanding the Implementation of Technical Measures in the Process of Data Privacy Compliance: A Qualitative Study*, in *Proceedings of the 16th ACM / IEEE International Symposium on Empirical Software Engineering and Measurement*, 2022, 261–271.

https://doi.org/10.6092/issn.1561-8048/14085

4. A general case for differential privacy and workplace data

This paper focuses on the use of differential privacy in a particular context: the sharing of workplace data from employers to employee representatives to facilitate collective bargaining. This is a 'weak' case for the use of this technology, focused on a specific interaction between employers and worker representatives.

There is another 'strong' case for the use of differential privacy techniques with workplace data. In this formulation, *all* workplace data would be funnelled through a differential privacy system.[28] Everyone, whether employer or employee representative would view data through differential privacy. Only in specific circumstances would individually identifiable data be available – and the specification of these circumstances could be subject to collective bargaining. Organizations that routinely handle datasets that could identify individuals, like national statistics bodies[29], are often organized around workflows with differential privacy as an essential component. The focus in this organisation is often on protecting data about third parties (e.g., citizens, customers), but the same principles could be applied to data about staff that are collected internally.

A better industrial relations context for workers would be one where employers did not enjoy advantages through privileged access to workplace information and greater capital to apply to making sense of it. This would require a significant – perhaps implausible – reconfiguration of the political economy of work. However, Dencik et al.[30] have addressed defeatist positions on datafication of workplaces. They argue that the kind of technical intervention I have suggested in this paper is fundamentally limited. They suggest that "the focus on access to and ownership of data is based partly on an assumption that the trend towards datafying workplaces cannot be reversed or substantially resisted".[31] The interviewees in their study suggested that issues around datafication reflect a fundamental weakness of trade unions in the United Kingdom, and that with stronger unions, issues around data might be treated like other contested issues in the workplace, rather than needing bespoke legislation or in situ technical expertise. There seems scope for new longitudinal empirical work in this area to assess the limits of particular interventions (technical, legal, organisational) on behalf of workers.

[28] A differential privacy approach relies on there being a store of raw data to sample from, and there are other good reasons to have raw data stored *somewhere* including for compliance reasons. There are few day-to-day scenarios where direct usage of raw data would be necessary, though. Organisations' data architectures often include abstraction layers; reporting tools used by non-technical staff draw on raw data sources, reformatting to make them more digistible. The same principle should apply to all data, but the abstraction layer should be differential privacy. If your role means that you need worker data, then your view of 'the data' is the output of the differential privacy output. All requests that could be made for data are funneled through this layer. Access to underlying raw data would only be possible in an enumerated set of circumstances. The power of the differential privacy approach is that conclusions drawn from data that have come through the differential privacy layer are commensurable with those that could be drawn from the raw data. It is functionally invisible to consumers of the data.

[29] B. Hawes M., *Implementing Differential Privacy: Seven Lessons From the 2020 United States Census*, in *Harvard Data Science Review*, 2020.

[30] Dencik L., Brand J., Murphy S., *What do data rights do for workers? A critical analysis of trade union engagement with the datafied workplace*, in *Transfer: European Review of Labour and Research*, 2024, 10242589241267006.

[31] ibid 10.

https://doi.org/10.6092/issn.1561-8048/14085

5. Conclusion

Information asymmetries and asymmetries in sense-making resources put worker representatives in a challenging position in collective bargaining, especially where there has been significant datafication of the workplace, that can make performance measurement less transparent. Employers have myriad reasons to be reticent to share data with worker employees, but data privacy laws often provide a plausible excuse for employers to refuse to share data.

I have proposed differential privacy as a way to nullify the 'cover' of data privacy laws being an obstacle to data sharing as part of collective bargaining. This is a statistical technique that allows data to be made sufficiently anonymous that it can be shared without removing the capacity for recipients of data (like worker representatives) to perform their own analyses of the data and develop their own stories about what it shows. It's not a panacea – it is not a legal mechanism for requiring sharing of the kind that some authors have proposed. It does, however, make it an employer's decision not to share, rather than (by proxy) a data regulator's. And collective bargaining exists to put pressure on employers' decisions.

Acknowledgements

I am grateful to an anonymous EPSRC grant reviewer whose comments ultimately produced the idea for this paper.

Bibliography

Adams Z., Wenckebach J., *Collective regulation of algorithmic management*, in *European Labour Law Journal*, 14, 2, 2023, 211–229.

Aleks R., Maffie M.D., Saksida T., *The role of collective bargaining in a digitized workplace*, in *Reimagining the governance of work and employment*, 85, 2020.

Allen J., *Wales TUC: Worker surveillance on the rise in Wales*, 2021, https://www.tuc.org.uk/news/wales-tuc-worker-surveillance-rise-wales, accessed 18 October 2024.

Aloisi A., *Regulating Algorithmic Management at Work in the European Union: Data Protection, Non-discrimination and Collective Rights*, in *International Journal of Comparative Labour Law and Industrial Relations*, 40, 1, 2024.

Anscombe F.J., *Graphs in Statistical Analysis*, in *The American Statistician*, 27, 1, 1973, 17–21.

B. Hawes M., *Implementing Differential Privacy: Seven Lessons From the 2020 United States Census*, in *Harvard Data Science Review*, 2020.

Calacci D., Stein J., *From access to understanding: Collective data governance for workers*, in *European Labour Law Journal*, 14, 2, 2023, 253–282.

Cecchinato M.E., Gould S.J.J., Pitts F.H., *Self-Tracking & Sousveillance at Work: Insights from Human-Computer Interaction & Social Science*, in Moore P.V., Woodcock J. (eds), *Augmented Exploitation Artificial Intelligence, Automation and Work*, Pluto Press, 2021.

Chung R. and others, *Do Managers Withhold Good News from Labor Unions?*, in *Management Science*, 62, 1, 2016, 46–68.

Colclough C.J., *Righting the Wrong: Putting Workers' Data Rights Firmly on the Table*, in Graham M., Ferrari F. (eds), *Digital Work in the Planetary Market*, The MIT Press, 2022, 291–302.

De Stefano V., Taes S., *Algorithmic management and collective bargaining*, in *Transfer: European Review of Labour and Research*, 29, 1, 2023, 21–36.
Dencik L., Brand J., Murphy S., *What do data rights do for workers? A critical analysis of trade union engagement with the datafied workplace*, in *Transfer: European Review of Labour and Research*, 2024, 10242589241267006.
Dwork C. and others, *Calibrating Noise to Sensitivity in Private Data Analysis*, in Halevi S., Rabin T. (eds), *Theory of Cryptography*, Springer Berlin Heidelberg, 2006, 265–284.
——, *Differential Privacy*, in Bugliesi M. and others (eds), *Automata, Languages and Programming*, Springer Berlin Heidelberg, 2006, 1–12.
Finck M., Pallas F., *They who must not be identified—distinguishing personal from non-personal data under the GDPR*, in *International Data Privacy Law*, 10, 1, 2020, 11–36.
Hancock J. and others, *The tensions of data sharing for human rights: A modern slavery case study*, in *The 2024 ACM Conference on Fairness, Accountability, and Transparency*, 2024, 974–987.
Huang T., Zheng S., *Using Differential Privacy to Define Personal, Anonymous, and Pseudonymous Data*, in *IEEE Access*, 11, 2023, 109225–109236.
Klymenko O. and others, *Understanding the Implementation of Technical Measures in the Process of Data Privacy Compliance: A Qualitative Study*, in *Proceedings of the 16th ACM / IEEE International Symposium on Empirical Software Engineering and Measurement*, 2022, 261–271.
Richards N., *The GDPR as Privacy Pretext and the Problem of Co-Opting Privacy The Internet and the Law: Legal Challenges in the New Digital Age: Essays*, in *Hastings Law Journal*, 73, 5, 2022, 1511–1538.
Rodríguez Fernández M.L., *Collective bargaining and AI in Spain*, in Castillo A.P.D. (ed), *Artificial intelligence, labour and society*.
Sherman H.L.Jr., *Employer's Obligation to Produce Data for Collective Bargaining*, in *Minnesota Law Review*, 35, 1, 1950, 24–46.
I'll be watching you | TUC, 2018, https://www.tuc.org.uk/research-analysis/reports/ill-be-watching-you, accessed 18 October 2024.
What is monitoring at work? | TUC, 2022, https://www.tuc.org.uk/guidance/what-monitoring-work, accessed 18 October 2024.
CNIL (France) - SAN-2020-003, https://gdprhub.eu/index.php?title=CNIL_(France)_-_SAN-2020-003, accessed 18 October 2024.
HDPA (Greece) - 31/2023, https://gdprhub.eu/index.php?title=HDPA_(Greece)_-_31/2023, accessed 18 October 2024.

https://doi.org/10.6092/issn.1561-8048/14085