

# Not Even Nice Work If You Can Get It; A Longitudinal Study of Uber’s Algorithmic Pay and Pricing

Reuben Binns\*  
firstname.lastname@cs.ox.ac.uk  
Department of Computer Science,  
University of Oxford  
Oxford, United Kingdom

Jake Stein\*  
firstname.lastname@cs.ox.ac.uk  
Department of Computer Science,  
University of Oxford  
Oxford, United Kingdom

Siddhartha Datta  
firstname.lastname@cs.ox.ac.uk  
Department of Computer Science,  
University of Oxford  
Oxford, United Kingdom

Max Van Kleek  
firstname.prefix.lastname@cs.ox.ac.uk  
Department of Computer Science,  
University of Oxford  
Oxford, United Kingdom

Nigel Shadbolt  
firstname.lastname@cs.ox.ac.uk  
Department of Computer Science,  
University of Oxford  
Oxford, United Kingdom

## ABSTRACT

Ride-sharing platforms like Uber market themselves as enabling ‘flexibility’ for their workforce, meaning that drivers are expected to anticipate when and where the algorithm will allocate them jobs, and how well remunerated those jobs will be. In this work we describe our process of participatory action research with drivers and trade union organisers, culminating in a participatory audit of Uber’s algorithmic pay and work allocation, before and after the introduction of dynamic pricing. Through longitudinal analysis of 1.5 million trips from 258 drivers in the UK, we find that after dynamic pricing, pay has decreased, Uber’s cut has increased, job allocation and pay is less predictable, inequality between drivers is increased, and drivers spend more time waiting for jobs. In addition to these findings, we provide methodological and theoretical contributions to algorithm auditing, gig work, and the emerging practice of worker data science.

### ACM Reference Format:

Reuben Binns, Jake Stein, Siddhartha Datta, Max Van Kleek, and Nigel Shadbolt. 2025. Not Even Nice Work If You Can Get It; A Longitudinal Study of Uber’s Algorithmic Pay and Pricing. In *The 2025 ACM Conference on Fairness, Accountability, and Transparency (FAccT ’25)*, June 23–26, 2025, Athens, Greece. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3715275.3732099>

## 1 INTRODUCTION

Uber is an on-demand ride-hailing service which connects passengers with drivers, and sets both passenger prices and driver pay. Until a 2021 Supreme Court case, Uber drivers in the UK were unlawfully misclassified as independent contractors rather than workers, and for the purposes of minimum wage law, were only considered ‘working’ while carrying passengers from pickup to dropoff. The Court upheld the Employment Tribunal’s view that working

time also includes whenever a driver is logged into the app and ‘available to accept a trip request’ [1]. While Uber has implemented some changes in partial compliance with the Court judgement, they still only pay drivers for time on trip and en route to pickup. This means Uber can increase availability of rides without increasing their costs, as they can grow the pool of drivers waiting around for trips without pay. A key conflict therefore remains between Uber’s restricted definition of working time, and a more expansive definition which includes all the additional time in between jobs.

When Uber first launched in the UK, prices were determined by the distance and duration of the trip, with Uber taking a fixed percentage of the passenger’s fare, first 20% and then later raised to 25%. In Q1 of 2023, Uber introduced ‘dynamic’ pricing in London, which means fares are no longer a simple function of time and distance. Instead the price the passenger pays and the fee the driver receives vary independently of each other, and are calculated dynamically based on location of pick up and drop off, time of day / week / year, probability of driver / passenger cancellation, and other factors undisclosed by Uber. This means that Uber’s take rate (the percentage they keep) is no longer stable but varies trip-by-trip. Dynamic pricing has been opposed by drivers on the grounds that it creates opacity and uncertainty around how pay is calculated, what cut Uber may take, and how much a driver can expect to make in a given hour or location [30].

The amount that a passenger pays for a trip is not available to drivers in the app, and Uber bars drivers from asking their passengers for this information directly, so it is not possible for drivers to work out Uber’s cut on any given journey. In December 2024, Uber began disclosing a weekly average take rate to drivers in the app, but this still does not include a trip-by-trip breakdown. However, some media reports suggest Uber’s cut can be as high as 47% [31]. An Uber spokesman responded that such take rates are ‘not the case when averaged out over the week’, and the take rate ‘remained stable for many years’. However, without both trip-level and aggregate analysis over time and across multiple drivers, the impacts of dynamic pricing are hard to assess.

In this work, we present an audit of Uber’s pay and work allocation algorithms, to understand: how much Uber drivers are being paid for their work (contrasting the Employment Tribunal’s definition of working time with Uber’s own); how dynamic pricing has

\*Both authors contributed equally to the paper



impacted the cut that Uber takes from the rider fare, and how this distributed; and related metrics like time spent waiting for jobs, and changes in the predictability of work opportunities. We partnered with Worker Info Exchange (WIE), a non-profit organisation representing gig workers in the UK and Europe. WIE helps gig workers access and gain insight from data collected from them at work, and has empowered hundreds of workers to use their rights under data protection law to access a copy of all their data from Uber, including detailed trip data. 258 drivers volunteered their data covering over 1.5 million trips between 2016-2024. We engaged with WIE and drivers in participatory action research over a period of two years, mainly on-site in London, UK. Our study design, methods, research questions, and findings have emerged through close collaboration with WIE and individual gig workers.

The primary aims of the audit, developed in collaboration with WIE and drivers, were to measure the following:

- **RQ1:Pay:** How much are drivers paid per hour worked (given contrasting definitions of working time)?
- **RQ2:Take Rate:** How are fares split between Uber vs drivers?
- **RQ3:Utilisation:** How much working time is spent on trips vs on standby?
- **RQ4:Predictability:** Has the predictability of work opportunities changed?

Our longitudinal data allows us to assess how these key aspects of work and pay are changing over time. In the case of take rates, we can see how dynamic pricing has impacted what percentage of a passenger's fare goes to the driver and how much is kept by Uber. We find that in the year following dynamic pricing, Uber took a larger median average cut, with take rates as high as 50% on some trips. In real terms, average pay per hour has been roughly stagnant since 2016, and is lower in the year following the introduction of dynamic pricing, even on Uber's own (limited) definition of working time. While a small number of drivers are earning more per hour worked, the majority are earning less. Post-dynamic pricing, Uber's passengers now pay higher prices, but drivers are not better off. Utilisation has decreased, meaning drivers spend more time on unpaid 'standby' while waiting to be allocated a job. Finally, work is less predictable after the introduction of dynamic pricing, meaning how much a driver can expect to earn from a trip of a certain duration / distance, at a given time, can no longer be predicted based on similar trips from before dynamic pricing.

In addition to reporting findings from our algorithm audit, we provide methodological and theoretical contributions to existing literature on algorithm auditing, gig work, and the emerging practice of 'participatory worker data science'[9, 16]. Methodologically, we demonstrate an alternative approach to algorithm audits which uses large-scale data subject access requests (DSARs). Unlike other approaches to external audits which often rely on black box observations or crowdsourced screenshots, this data derives directly from the algorithm deployer's own back-end, revealing schemas, attributes, identifiers, profiles and scores applied to data subjects. To our knowledge, this is the first contribution to the algorithm auditing field of a large-scale audit based on DSARs.

## 2 BACKGROUND

### 2.1 Algorithmic Management and Worker Resistance

Algorithmic, data driven management of human workers on gig platforms like Uber has been studied extensively in a variety of research domains. Lee et al. studied drivers' use of online forums to make sense of work allocation algorithms [26]. Compared to human managers, workers perceive algorithmic management decisions to be less fair and trustworthy, [25]. Further work in this vein has used participatory design to co-design interventions with workers, including: to understand their needs for well-being [49]; to elicit stakeholder preferences for addressing worker well-being [19]; and to explore 'data probes' — interactive visualizations of workers' data - as training tools for policymakers [50]. Ma et al. study how gig work platform designers sense of morality regarding the plight of workers affects their own practice [27]. Through a study of mobility platforms in Jakarta, Qadri and D'Ignazio chart how workers develop their own 'ways of seeing' urban mobility markets, which can both contest and complement the algorithmic visions of the platform [34].

Recent work has explored various tools which could be deployed by and for workers themselves. Calacci et al., designed and deployed Shipt Calculator, a tool developed in collaboration with non-profit worker groups to track and share aggregate data about worker pay, increasing wage transparency [9] (this provides direct inspiration for parts of our audit below). Do et al. explore a variety of 'sousveillance' tools to allow workers to monitor their algorithmic managers [13]. Similarly, Gallagher et al. develop 'participatory worker data science' [16], through the 'Workers' Observatory', a collaboration between workers and researchers. In previous work, also in collaboration with WIE, we explored possible data governance and infrastructure to enable bottom-up data institutions for gig workers [41].

Some prior studies involve auditing of Uber's pricing algorithm. These include a 2015 study of Uber's surge pricing algorithm in major US cities [11]; a study of passenger pricing which found disparate impacts in pricing between neighbourhoods [32]; and Brodeur and Nield's 2018 analysis of weather shocks on mobility gig work platforms, which demonstrates how dynamic pricing can reduce driver income on Uber (among other platforms) [7]. Most recently, the FairFare project crowdsources and analyzes workers' data to estimate Uber's take rate in the US[10].

Despite confirmation of their legal status as workers in the UK following the 2021 Supreme Court judgement, Uber continues to inculcate the ideal of drivers as entrepreneurs. Prior work has argued that the ideological construction of the driver as an atomised individual, an 'entrepreneurial self' [18, 33], is an important aspect of Uber's legitimization in the eyes of policymakers, regulators, and drivers themselves. Gregory & Sadowski argue that gig workers are encouraged to develop their own human capital in perverse ways to make themselves 'more productive, more desirable workers'[17].

While worker resistance and solidarity does persist[44], such efforts are often made harder under algorithmic management, which replaces human supervisors 'with a set of highly depersonalised and disembodied commands' [46]. Rosenblat and Stark argue that

Uber relies on ‘information and power asymmetries’ and ‘rhetorical invocations of digital technology and algorithms’ to control workers and structure asymmetric corporate relationships in their favour [37]. Many workers buy into the idea of themselves as entrepreneurs, which aims merely to improve individual conditions without challenging the platforms’ business model [3]. Nevertheless, surveys of platform workers in the UK suggest support for more labour rights, representation, and voice is high [28]. Finally, not all gig work platforms treat their workers the same. Kusk et al conducted an ethnographic study of workers on the Wolt food delivery app [24], finding that its practices of algorithmic management are perceived as less ‘harsh’ and ‘despotic’ than those of Uber. Alternative ownership structures, with better working conditions, are also imaginable [22, 38].

## 2.2 Algorithm Auditing with Data Subject Access Requests

The nascent sub-field of algorithm auditing has developed rapidly over the last decade. Audits can be performed internally by the organisation deploying the algorithm, or externally by a third party, either consensually (i.e. upon invitation by the deployer) or non-consensually (e.g. by researchers and activists without permission from the deployer) [35]; our audit is external and non-consensual. While many organisations are voluntarily engaging in audits [48], external audits remain an important tool to uncover algorithmic harms and seek redress for affected communities where internal audits and harm mitigation are not forthcoming. External audits have covered topics including: racial discrimination in online ad delivery [42], job ads [21], personalisation in recommender systems such as Youtube [20], and gender and racial bias in facial recognition technology [8]. However, such external audits can be limited by their reliance on system behaviour that is publicly observable to researchers from the outside. They do not have access to the internal, back-end systems of the algorithm deployer. As such, researchers are often in the dark about exactly what data is collected, and how users are scored or classified on the basis of that data. Privacy concerns mean individual-level data is also not easily accessible, and even where scrapable, its unconsented use is at odds with prevailing research ethics norms. Such issues can significantly affect the kinds of questions that can be asked and the ecological validity of external audits.

However, a complementary data source for algorithm audits is also available: legally mandated data access rights such as Data Subject Access Requests (DSARs). Research participants in the UK and EU can access their data from organisations using a DSAR under Article 15 of the General Data Protection Regulation, as well as a subset of portable machine readable data under Article 20. DSARs have already been used by researchers and data rights activists to investigate organisations’ data practices [2], including those of complex tech infrastructures and data pipelines such as Apple’s Siri voice assistant [45].

So far, such research has mainly sought understanding with relatively small samples. But with larger samples, it may be possible to investigate systemic effects. Recent proposals include open-source data donation frameworks (e.g. OSD2F) which combine multiple research participants DSARs while preserving their privacy [6]. A

similar approach has been pursued for data obtained under Article 20 [52], and for data download functions (e.g. to study ad targeting on X/Twitter [47]). Data rights activists have also employed DSARs for studying algorithms; notable examples include Open SCHUFA<sup>1</sup>, an attempt to reveal the biases of a ubiquitous German credit scoring algorithm, and Hestia Labs, which provides DSAR tools for collectives, including Uber drivers.<sup>2</sup> Our approach is inspired by these efforts.

## 3 METHODOLOGY, DATA SOURCE, & PARTICIPATORY AUDIT DESIGN

In this section, we describe our methodology, data source, and approach to participatory audit design.

### 3.1 Methodology

Our approach to this work was informed by previously identified challenges and limitations associated with algorithm auditing and the use of data rights in worker justice. First, we aimed for our research and auditing efforts to be aligned to and embedded within existing struggles for worker justice (following [12]). We drew from Katell, Young et al.’s framework for developing ‘situated interventions for algorithmic equity’ [23], which emphasises first working with organisers to document the degree of current efforts, resources available and setting milestones or conditions for success. We also followed a participatory action research (PAR) methodology, meaning that we iteratively identified problems, co-designed solutions, and reflected on the efficacy of both the solution produced and the design process itself. This was made possible by our partnership with WIE, which involved two researchers regularly spending days embedded in their workplace, enabling us to not only work on the specifics of the audit, but understand the issues and needs of organisers and drivers on a day-to-day level, hear about their latest campaigns, and attend trade union meetings and protests. All of these activities helped us to understand ongoing issues and actions, and to gain feedback on our ongoing auditing work.

As previous work has argued, individual-level data received in subject access requests may not enable the study of collective impacts, and its collection and schematisation will likely reflect the interests of Uber, rather than its workers [9]. We sought to overcome these shortcomings of individual-level data access, by pursuing the ‘collective leveraging of data rights’ [40], while also working closely with organisers and workers to critically interrogate Uber’s schemas and metrics, and attempt to re-construct alternative metrics from it which are more conducive to their interests.

As Birhane et al caution [5], AI audits often fall short due to the following limitations: focusing only on evaluation, without considering accountability; failing to engage with excluded stakeholders or foster collective action; or failing to use the full range of audit methodologies, focusing narrowly on popular metrics (e.g. ‘group fairness’). We aimed to overcome such limitations through our approach. First, the audit design is already informed by WIE’s day-to-day work, which focuses on worker organising, policy advocacy, strategic litigation, and other work aimed at holding platforms to account. WIE’s close links with affected stakeholders - in this case,

<sup>1</sup>openschufa.de

<sup>2</sup>hestialabs.org

Uber drivers themselves - allowed us to directly involve them in the design of the audit and formulation of research questions. Driver participation was essential, not only in making the data collection possible but also in motivating and shaping the analysis. Finally, we took an open-ended approach to audit metrics, creating a bespoke set of analyses to highlight insights of concern to the stakeholders, rather than attempting to apply existing formal metrics (e.g. from the algorithmic fairness literature).

### 3.2 Data Source

This section briefly describes the primary data source - a collection of 258 DSAR responses - including their provenance and governance. Individual drivers sign an affidavit with WIE to allow them to make and receive DSARs on their behalf, for the purposes of worker advocacy and research. We created a data transfer agreement, undertook a data protection impact assessment, and applied and received ethical approval from our institution's Research ethics committee, under reference 'CS C1A 23 016'.

When WIE receives data from gig work platforms, it takes the shape of many disorganized files that are not readily understandable without substantial interpretation, preprocessing, and cleaning. In the case of Uber driver DSARs, we typically found between 40-45 different files, of which around 35 were CSVs containing a wide variety of data from different parts of Uber's backend systems. We worked on-site initially, to understand the varying data schemas and identify fields that would be likely required for further analysis. We created an automated tool to allow staff who lacked requisite technical skills to process the data into more interpretable forms themselves. This involved cleaning, reformatting, converting and joining data from multiple tables (see 8 in Appendix). Given the sensitivity of the data, this tool ran locally on a dedicated laptop located in WIE's offices accessible only to a senior member of staff. For the audit itself, where we as researchers required access to larger-scale aggregate data on our own infrastructure, we created an aggregation pipeline which pseudonymised or stripped entirely a range of sensitive fields. Direct identifiers such as name, email, etc, were removed, and the platform-assigned driver ID was hashed, enabling the driver themselves or WIE (but not us as researchers) to re-identify the driver if required.

### 3.3 Participatory Worker Data Science

In order to facilitate worker involvement in the audit design, we began by creating a set of simple data visualisations to stimulate discussion between ourselves, drivers and organisers. We used the DSAR Guidance document provided by Uber to identify key tables and fields, cross-referencing against other sources including weekly driver earnings reports, and sense-checking our interpretation of the fields with organisers and drivers. This side-by-side exploration allowed us to better understand the data schemas, as well as map the sensitivity and contestation around particular data types. We then identified some key metrics of interest to see if they could be calculated from the DSAR data. These included monthly net income, pay per hour (distinguishing between Uber's definition of working time, and the contrary one of the Employment Tribunal); utilisation rates (how much time is spent en route to pickup, on

trip, or on standby waiting for a dispatch); and the 'take rate' (how the passenger's fare was split between Uber and the driver).

This process culminated in a session with 15 drivers based in the UK, who are members of WIE and had expressed interest in learning from their work data. The session was co-led by one researcher and one organiser, who briefly explained the data, outlined potential directions for the audit, and solicited feedback and discussion. Following [16], the purpose of this session was to put into practice 'participatory worker data science', by helping us to identify competing interpretations of the data and to explore alternative metrics. This session helped us to refine the key areas of focus for the audit, which are presented below. We include quotes from drivers to illustrate the motivation behind each focus area:

- **Pay Trends** Participants showed particular interest in graphs showing trends in monthly net pay. Reacting to stagnant or downward trends, participants said these could help them decide whether to keep working for Uber, or give them of forewarning of when they might be forced to find alternative work.
- **Working time definition** Despite some initial confusion, most participants understood and valued a graph displaying pay per hour in two ways: first, using Uber's definition of working time, and second, using the alternative expanded working time definition from the Employment Tribunal. As one driver observed of Uber's definition, *'it's only from the time you accept the job until you drop off, that's for them what's considered working time period'*, whereas the line on the graph representing the Tribunal's working time definition showed the 'real' amount of work, and was referred to as *'the real line'*. We therefore sought to measure the difference between pay per hour on both definitions in our analysis.
- **Acceptance rates** Participants raised several points in relation to the acceptance rate data (i.e. what proportion of jobs offered were accepted by the driver). One hypothesis was that there would soon be a division between more experienced, discerning drivers, for whom acceptance rates may be lower, and newer drivers, who would accept poorly paid jobs. They argued that the pay is now worse because of a recruitment drive creating excess supply. This highlighted the collective impacts of different acceptance/rejection strategies. The extent to which acceptance was really voluntary was brought into question given suspicions about Uber's dispatch algorithm treating them differently based on past acceptance rates: *'It's indirectly forcing the driver to accept; the natural instinct is 'oh I'm not getting any dispatches, is it because I just rejected two just now? Is the system trying to ignore me? Oh I'd better accept whether it suits me or not'*. This motivated us to study any relationships between pay per hour and acceptance rates.
- **Commission / Take Rate** Many drivers lamented the loss of key information that has previously (before dynamic pricing) been provided in their weekly earnings report: *'One thing they were showing before was the price the client was paying, and now they don't. That's why now I question how much they charge the client and how much they give to us, because we cannot see'*. While Uber claimed to only take 25%, many

were sceptical. One stated: *‘they are actually charging more than 25% ... They want to hide how much commission they are taking from you’*. Some described circumventing this by asking customers how much they had paid: *‘The lady paid £64 but I was getting exactly 50% which was £32’*. Such practices were seen as socially awkward (and have since been banned by Uber), but the only way they could see to uncover potentially unfair pricing for the driver and the customer. As one participant put it: *‘That’s when you discover they [Uber] are robbing us and the customer’*.

- Utilisation.** Some drivers were concerned about the seemingly increasing amount of standby time (i.e. time on the app, ready and willing to work, but waiting for a dispatch). As one driver pointed out, *‘it doesn’t cost Uber anything to have every one of us on the road at the same time, regardless how much work there is’*. Many agreed that the recent influx of new drivers - *‘10,000 new drivers on the road’*, as one highlighted - was leading to more competition for jobs and therefore more standby time. One participant argued that more conscious utilisation of the workforce would benefit drivers, suggesting that Uber *‘should introduce certain hours, they could advise drivers that you could do this, or this’*, lowering standby time and helping drivers plan their working hours around demand. Another driver agreed, adding that *‘they invest a lot in AI, but they don’t invest in our tools; we are the ones who need to have that sort of information’*. This motivated the inclusion of utilisation rate analysis in our audit.
- Dynamic pricing** The topic of dynamic pricing itself was brought up at multiple points during the session, which motivated our choice to analyse impacts before and after its introduction. The workshop took place a few months after the introduction of dynamic pricing. Drivers had many concerns about how this would change their jobs, and hypotheses about its impacts. Drivers complained about their lack of a say about such changes: *‘we all agreed [to dynamic pricing], to many things, because we do not have any other option’*.

Based on these identified focus areas, we continued working with WIE to narrow down the design of our audit. Key metrics included pay, take rate, acceptance rates, and utilisation. In particular, for all of these factors, we were interested in any changes before and after the introduction of dynamic pricing. This was first announced in London (where the majority of drivers in our dataset work) in February 2023[39]. While there has been no official communication from Uber about the exact timing of rollout in different regions, based on driver testimony and information from WIE, rollout was completed across the UK by summer 2023.

For the audit, we restricted our selection to DSAR responses received between October 2023 and December 2024. This includes over 1.5 million trips in total, and at least 100k trips for each year between 2016-2024.<sup>3</sup> We collected driver demographics including gender and age. Participants were split 96% men, 4% women. In terms of age, most were below 40 (*20-29: 23%, 30-39: 35%, 40-49:*




Figure 1: Pay per hour according to employment tribunal definition of working time (blue) vs Uber’s definition (orange)




Figure 2: Average hours/day on standby (red), en route (green), on trip (blue)

*28%, 50+: 14%*). We compare these demographics to alternative sources to assess how representative our sample was to the whole population. While few sources show demographics of Uber drivers specifically, the regulator Transport For London [43] report demographics for the private vehicle hire sector as a whole (which includes Uber and other private hire vehicles, but excludes traditional black cab taxis). TfL figures show a slightly higher proportion of men (97%), and skew older (*20-29: 6%, 30-39: 26%, 40-49: 36%, 50+: 32%*). A 2019 study, co-authored by an Uber employee with access to randomly sampled anonymised data from Uber’s administrative records, reports 99% drivers are men, and 52% aged under 40 (compared to 58% of our sample). Our sample appears to skew slightly younger than these sources, and is similar to both in terms of gender split.

## 4 AUDIT FINDINGS

### 4.1 Pay

The DSAR responses contain a primary `Payments.csv` file including a wide variety of incoming and outgoing payments associated with trips. To calculate a per hour pay figure for a given week, we aggregated all incoming and outgoing payments and divided that by hours worked in the week. This is imperfect, because some kinds

<sup>3</sup>Code and data for this analysis are available at [github.com/OxfordHCC/FACCT\\_25\\_Not\\_Even](https://github.com/OxfordHCC/FACCT_25_Not_Even)

of payments (and driver charges) may cover periods longer than a week, or occur very irregularly. However, it is assumed these will even out when aggregated over time and across drivers.

We calculate hours worked using two definitions. The first is that favored by the Employment Tribunal: any time spent on the app ready and available to work, whether that is waiting for a request ('standby'), en route to pick up, or on trip. The second is Uber's preferred definition: only time spent en route or on trip. We found a large difference between per hour pay depending on the definition of working time (see figure 1). On Uber's definition, the average is £29.46 across 2020-2024; whereas on the Employment Tribunal's definition, pay per hour was only £15.98. Both averages do not account for the costs incurred by drivers like vehicle maintenance, insurance, and fuel.

After adjusting for inflation, average pay per hour has also declined according to both measures of working time. On the Employment Tribunal definition, average pay per hour the year prior to dynamic pricing was £22.20 (adjusted for inflation), vs an average of £19.06 the year post-dynamic pricing. On Uber's definition, average pay per hour the year prior to dynamic pricing was £37.01 (adjusted for inflation), vs an average of £35.91 the year post-dynamic pricing.<sup>4</sup> We also note that average pay per hour on Uber's working time definition is less turbulent around year 2020, where demand may have been lower due to the COVID pandemic public health measures which restricted social mixing. At such times, drivers bear the cost of low demand, as they drive around waiting for trips. While this shows up in the low pay per hour according to the Employment Tribunal definition, it is less visible under Uber's definition, where working time only accounts for time en route to pickup or on trip.

Note that these figures are averaged across all drivers and trips, and therefore may not reflect the experience of individual drivers. As shown below in section 4.4, pay per hour has declined more significantly for drivers who have been driving consistently the years before and after the introduction of dynamic pricing; the averages reported here for the year after dynamic pricing are higher due to the higher rates experienced by drivers who joined after dynamic pricing.

## 4.2 Utilisation Rates

Figure 2 shows the average hours per day spent on trip, en route, and on standby, aggregated across all drivers in a given month. While on trip and en route time has remained stable, standby time has risen significantly - increasing by over 1 hour a week since 2022, and remains high. In most months since 2023, drivers spend more time waiting to be allocated their next job than they do on journeys with passengers.

## 4.3 Commission / Take Rates / Surplus

Until the introduction of dynamic pricing, Uber had publicly advertised their commission rates at 25% (having risen from 20%)[31]. However, since dynamic pricing, the customer fares were removed from drivers' weekly earnings report, meaning drivers could no longer see how much of a cut Uber takes from any given trip (the

<sup>4</sup>Inflation figures come from the UK Office for National Statistics document 'RPI All Items: Percentage change over 12 months'





Figure 3: Frequency of driver/Uber splits, post dynamic pricing

'take rate'). As explained above, drivers are banned from asking passengers about their fares, making it impossible for them to work out the take rate. However, customer fares were included in DSAR responses, albeit in an incomplete form. A unique trip identifier is not provided in the data, but by joining multiple tables and inferring which payment related to which trip based on timestamps, we were able to calculate the take rate by dividing the 'original fare' by the driver payment associated with the trip.

However, this integration was not possible for one year's worth of data due to an apparent change in the backend system which generates the DSAR data. Prior to February 2022, DSARs contained an 'original fare' field in `Trips`.CSV, which corresponded to the price paid by the customer (we independently verified this with individual drivers); as well as a 'commission' charge, which appeared at the same time as a corresponding trip, in `Payments`.CSV, which always added up to the correct commission rate (20% prior to 2021, and 25% after). However, in February 2022, the 'original fare' appeared to no longer represent the customer's fare, and the commission charges disappeared from `Payments`.CSV. This coincides with the point at which Uber introduced measures to (partially) comply with the previous year's UK Supreme Court ruling. Explaining these compliance measures, Uber stated: *'the rider [i.e. the passenger] will now pay Uber directly instead of paying the driver - a service fee is no longer suitable as a mechanism to provide payment to Uber for using our app'*<sup>5</sup>. This means that take rate data is missing from our dataset from 2022-02. However, since the introduction of dynamic pricing in 2023-02, the original fare field appears to have been restored to its previous source, once again reflecting the fare paid by the customer (as confirmed by drivers and customers), enabling us to continue our take rate analysis. As discussed above, in January 2025 Uber began to reveal weekly average take rates in the app, but still do not disclose trip-level take rates.

Figure 3 shows the distribution of take rates (in bins of 10%) for all trips since the introduction of dynamic pricing. In contrast to the standard 75/25% split prior to dynamic pricing, driver take rates are now frequently as low as 50-60%, and in some cases even below 50%. Conversely, in other cases, the driver's take rate can exceed

<sup>5</sup><https://www.uber.com/en-GB/blog/driver-terms-faq/>



**Figure 4: Surplus generated by drivers for Uber per hour on trip**


100%, i.e. the passenger’s fare does not even cover the driver’s pay, meaning Uber makes a loss on these journeys.

Mean average take rates have remained at 75%, which supports Uber’s claim in [31] that their 25% cut remains stable. However, the distribution is skewed right, with the median average take rate post-dynamic pricing at only 71%. Only 46% of drivers have managed to maintain an average take rate of 75% or above.

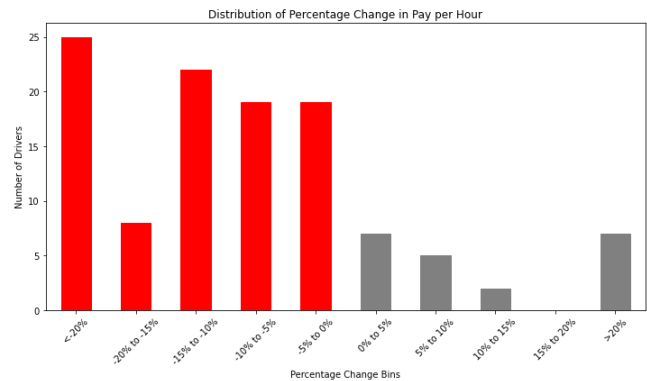
Note that Uber calculates its take rate differently in the weekly averages provided to drivers in the app, because it subtracts passenger promotions and third party fees from its take rate. We did not include these in our take rate analysis for several reasons. First, passenger promotions are discounts that Uber has decided to give to passengers; as such, the full price should be considered as Uber’s ‘real’ valuation of the trip, not the discounted amount paid by the passenger. Second, third party fees are not driver pay, but reimbursements of operating costs (such as London’s congestion zone charge) that drivers have had to pay out of their own pocket (which also aren’t included in our analysis). Finally, we were also unable to link these costs to particular trips because Uber does not provide a trip ID, and unlike with direct trip payments, we could not infer links by correlation of timestamps.

In addition to measuring Uber’s take rate, we also calculated the average surplus, defined as the amount of revenue generated for Uber by each driver for each hour they are on a trip, minus the driver’s labour cost (see Figure 4). This is not equivalent to surplus profit, as it does not account for Uber’s non-driver costs (such as employees, infrastructure, overheads, etc). However, assuming such costs are relatively low per trip, and have remained stable, this provides a reasonable proxy of the change in the surplus value generated by a driver while en route / on trip. We find that Uber’s average surplus per worker hour (minus non-driver costs) has increased by 38%, from an average of £8.47 in the period 2020/01-2022/02, to an average of £11.70 in 2023/03-2024/10.

Given that mean average take rates have remained stable, and median take rates have only decreased by 4 percentage points, this increase in surplus cannot be explained by increasing average take rates. However, looking at the distribution of take rates across



**Figure 5: Average per minute fare for varying driver/uber splits. As Uber’s cut increases, fares increase, but drivers earn less per minute (on trip) in absolute terms. Red area shows Uber’s cut as negative, i.e. when driver’s pay is more than 100% of the fare**




**Figure 6: Distribution of percentage change in pay per hour after dynamic pricing, for higher earners (grey) and lower earners (red)**

different value trips provides another explanation (see Figure 5). We find that Uber’s take rates are not evenly distributed across customer fares; rather, the higher the fare charged to the customer, the higher Uber’s take rate, and the less drivers earn per minute in absolute terms. Drivers are therefore worse off when they select higher-value fares.

#### 4.4 Inequality Between Drivers

How are changes in pay post-dynamic pricing distributed between drivers? We selected the period covering one year prior to and one year after the introduction of dynamic pricing (2022/02-2024/02), and selected data from only those drivers who had monthly data covering this period (114 drivers).

We found that average per hour pay for this subset of 114 drivers had decreased from £18.52 to £17.07 (a larger drop than that observed across all drivers in the dataset above). We found that 93 drivers were worse off in terms of their average per hour pay after dynamic pricing, and 21 were better off (Figure 6, left). Figure



**Figure 7: Pay per hour for higher (blue)/lower (red) earners after dynamic pricing.**

7 (right) shows rolling monthly per hour pay averages for those drivers whose average per hour pay was lower (red), and higher (blue). This shows that these two groups had relatively similar pay per hour pre-dynamic pricing, but after its introduction in 2023/02, they consistently diverge (with some re-convergence towards the end of 2024). This mirrors a similar pattern found in [9], where the introduction of a new algorithm on the Shipt gig work platform lead to inconsistent effects on pay among gig workers.

Based on discussion of this data with drivers and WIE, we undertook some further exploratory analysis to understand if there are any factors that might explain the differences between these groups (those who were earning less, vs same or more). We found no clear differences in hours worked, or number of years experience (see Appendix, Figure 9); however, we saw some differences in driver take rates and acceptance rates. For those who earned the same or more per hour, there is a higher concentration of trips with a driver take rate around 76%, while for those paid less, the distribution is wider (Figure 10). In terms of accept rates, those paid the same or more skew further right (Figure 11). While neither set of differences are large, when combined they may at least partially explain the differences in pay per hour.

#### 4.5 Predictability of Pay

A driver’s livelihood depends on their ability to guess what kinds of trips they will get at particular times and places, and how much those trips will pay. However, in our participatory worker data science sessions, drivers frequently complained about the unpredictability of pay post-dynamic pricing. We sought to measure this unpredictability by training and testing models to predict the pay associated with a trip, based on over 60 variables associated with each trip. These included duration, distance, time of day, day of week, airport trips, etc. The purpose of this exercise was to test whether, even with a large amount of pooled data, the distribution of pay under dynamic pricing would become unpredictable based on pre-dynamic pay patterns. Our model performance represents an upper bound on the mental models and tacit knowledge built up by drivers over the years, to guide their strategies about when and where to work. A reduction in performance of this model would

imply a reduction in the utility of such knowledge under the new dynamic pricing regime.

First, we trained linear regression models on subsets of subsequent years, and then tested them on the following year. We evaluate the regression models on a test set with R-squared (out of 1.0). In Table 1, we used the previous year ( $Y - n$ ) to predict the current year  $Y$  (e.g. train on 2021, test on 2022). This would evaluate how consistent the variables from  $n$  years ago would be on a given year. We see a drop in performance in 2020 (after the Covid-19 pandemic), and an even more significant drop post-2023 (after the introduction of dynamic pricing). Performance in previous years improved or stayed steady year on year, up until these years where it dropped off. For example, if we used any data prior to 2022 to predict pricing on 2023-4, the r2 score remains consistently poor/negative. In Table 2 (see Appendix), we used all data up until the given year 1, ...,  $Y - n$  (e.g. train on data from 2014-2021, test on 2022). We find a similar pattern as above, indicating even consistent variables over multiple years do not predict well post-dynamic-pricing.

We note that models that are both trained and tested on post-dynamic pricing do perform well. So it is possible that over time, if the dynamic pricing model and other factors do not change significantly, drivers may be able to gradually rebuild the tacit knowledge built up in previous years to match the new dynamic pricing regime. Nevertheless, in the interim, these changes in distribution of the underlying variables pre- and post-dynamic pricing would likely be highly disruptive to any driver whose strategies and working patterns are based on previous knowledge.

## 5 DISCUSSION AND CONCLUSION

Our discussion is divided into direct commentary on the audit findings and limitations of our approach, followed by various implications for political action, policy, and research.

### 5.1 Audit Findings

Our findings suggest that post-dynamic pricing, many aspects of Uber drivers’ jobs have gotten worse. Average pay per hour on the app is stagnant, and is lower in real terms in the year following the introduction of dynamic pricing. The discrepancy in pay per hour between the Employment Tribunal’s definition of working time vs Uber’s has increased. Uber’s median take rate per driver has increased from 25% to 29%, and on some trips the take rate is over 50%. Furthermore, the higher take rates are concentrated among higher-fare trips, which explains how Uber can extract an additional 38% surplus from its driver’s labour on average (equivalent to £3.23 per hour), a phenomenon which is obscured in the average take rate. Many drivers are earning substantially less per hour. Utilisation has also changed; in some months, drivers spend more time waiting for the next trip than they do actually on trip. Finally, any tacit knowledge drivers have built up over years about how much pay a given trip is likely to yield may no longer help them; as our modelling suggests, the predictability of pay drastically changed after dynamic pricing was introduced.

### 5.2 Limitations

Our study has various limitations. First, our sample of 258 drivers, while relatively large for a research study, is still small relative

	Y	Y-1	Y-2	Y-3	Y-4	Y-5	Y-6	Y-7	Y-8	Y-9	Y-10
2014	0.873										
2015	0.628	0.485									
2016	0.799	0.592	0.211								
2017	0.849	0.817	0.134	-0.410							
2018	0.887	0.881	0.858	-0.171	-0.615						
2019	0.871	0.865	0.856	0.803	-1.583	-2.008					
2020	0.919	-0.053	0.781	0.874	0.491	-7.685	-7.766				
2021	0.821	0.801	0.283	0.509	0.772	0.517	-7.808	-6.277			
2022	0.861	0.837	0.808	0.397	0.731	0.781	0.627	-4.477	-3.678		
2023	0.858	-54.023	-20.144	-5.625	-0.739	-9.888	0.787	0.679	-3.219	-3.441	
2024	0.891	0.886	-0.219	0.434	0.711	0.833	0.630	0.828	0.774	-3.489	-3.674

**Table 1: Predicting price at year  $Y$  using previous year  $Y - n$**

to the whole population of 100,000 Uber drivers in the UK. It is also drawn from drivers who are sufficiently motivated to sign up to WIE’s DSAR scheme, and who may therefore be more likely to already have had problems with Uber’s platform than other drivers (although, as explained above, the demographics of our sample appear similar to the wider population). While some of our observations are directly attributable to dynamic pricing (e.g. those related to the variation in take rates), our study design does not enable us to isolate the causal effect of dynamic pricing on pay, utilisation and predictability. Contemporaneous changes in passenger demand, traffic flows, workforce size and composition, and even potential behavioural change among drivers in response to dynamic pricing, limit our capacity for causal attribution.

### 5.3 The Ethics of Dynamic Pay and Pricing

In its initial incarnation, Uber presented a simple payment model to drivers. Pay was based on a flat rate per mile and per minute, and Uber took a fixed commission. Since then, the payment model became increasingly complex. Additional factors were introduced into the pricing calculation; upfront fares based on Uber’s prediction of duration and distance were introduced; the commission rate increased, and then was removed; other kinds of payments and charges were introduced following the Supreme Court ruling. Finally, the entire system was replaced by dynamic pricing. As a result, the current payment model is unrecognisable from the original one. Drivers who signed up in earlier waves of recruitment, even if they didn’t know exactly what jobs they’d get, could at least know how pay is calculated (a simple function of time and distance), giving them some predictability, stability and autonomy. But this is made impossible by dynamic pricing (as our prediction models attest). As Uber’s CEO described the trajectory in a call to investors: ‘you’ve gone from just flat time and distance to now kind of point estimates for every single trip based on the driver... targeting of different trips to different drivers based on their preferences or based on behavioral patterns that they’re showing us, that really is the focus going forward, offering the right trip at the right price to the right driver’.<sup>6</sup>

A key feature of the old pricing model was that it forced Uber to be transparent about its cut, as it was a fixed commission across all

trips. This meant if it wanted to charge a customer more, the driver would always automatically benefit proportionally, thus aligning driver’s interests with those of the platform, against those of the customer. Drivers were unlikely to complain about higher customer prices if they benefit proportionally. However, post-dynamic pricing, Uber has made commission take rates variable, cutting the tie between customer price and driver pay. As our analysis of Uber’s surplus shows, Uber also take a higher percentage of the most valuable fares, thus extracting even greater surplus per hour a driver works.

For drivers, this breaks what was a mutually beneficial compact that had previously kept them from making common cause with customers. As drivers in our participatory data science sessions argued, drivers and customers are now both being ‘robbed’. One recalled: ‘You ask the customer ... how much are you being charged? And they tell you this much, and then you think well I’m only getting this much ... I’ve challenged Uber before ... What they say to me is: you saw the upfront price, you took it, that means you accept it. They just blocked me, they shut it down.’ Instead of the dependability of being paid for the time and gas spent on trip, minus a consistent and publicly known commission, drivers must now continuously gamble, never sure how much cut Uber is taking. This reflects a broader development that legal scholar Veena Dubal has called the ‘algorithmic gambification’ of work, where: ‘on-the-job data collection and algorithmic decision-making systems ... are undermining the possibility of economic stability and mobility through work by transforming the basic terms of how workers are paid’ [14]. In this model, the house always wins.

### 5.4 The Reserve Army of Platform Labour

Under capitalism, workers have the freedom to sell their labour at whatever price employers are willing to pay. However, the existence of a ‘reserve army’ of unemployed labourers, who would be willing to sell their labour for cheap, is key to resisting pressure to increase wages from those who are in employment [29]. Since the rise of zero-hours contracts, this distinction between the reserve army and the rest of the workforce was already blurred, with a large class of variably (un)employed workers. But gig work platforms have turned this dynamic between the employed and reserve army into an increasingly real-time, variable proposition. Uber has brought many new drivers onto the platform. When they get there, they are

<sup>6</sup><https://www.fool.com/earnings/call-transcripts/2024/02/07/uber-technologies-uber-q4-2023-earnings-call-trans/>

added to the waiting list of drivers on standby, looking for work. Faced with such oversupply, drivers become reluctant to turn down offered trips, even if low paid, because they fear other drivers will accept them.

Having this excess supply adds very little additional cost to Uber, but does cost the worker; as one driver put it, *'They have no cost, but we have cost'*. Such costs include: time wasted while on standby, looking for work, fuel, toll charges, as well as upfront costs of getting set up as a driver on the platform in the first place. Excess supply also has social and environmental costs; more drivers driving around looking for work means greater carbon emissions and pollution on the streets; greater congestion, slowing down other traffic and public transit; and safety impacts associated with more cars on the road. All of these costs are fully externalised by Uber, providing no incentive for them to fix them.

As drivers in our participatory data science sessions recognised, the ease with which customer requests are fulfilled might appear to be down to Uber's sophisticated AI, but may actually have more to do with the oversupply of labour: *'they're trying to create this perception so the customers think 'Oh yeah, it's good that they've used [AI], because we couldn't get a car before, now we get a car just like that ... but there's a lot more drivers out there just sitting around!'*. As drivers pointed out, there could be other, more equitable ways to use AI to better allocate the useful work of transporting passengers around, that would provide stable and dependable work. As shown in section 4.5, demand for transportation is to a large degree predictable, even if dynamic pricing isn't. A transportation service with access to such data at scale could enable better prediction and co-ordination, allowing drivers to work at pre-arranged set hours, greatly reducing unnecessary standby time and the social and environmental costs it brings to drivers, cities, and the environment [4].

## 5.5 DSARs for algorithm auditing and worker data science

To our knowledge, this is the first large-scale algorithm audit performed using DSARs, other than the aforementioned examples of the OpenSCHUFA Credit Score initiative. One of our goals was to investigate whether, despite previously identified limitations, DSARs have the potential to be effective as tools for algorithm auditing and worker organising. Our experience suggests that this can be a promising tool for research. It enables bottom-up investigations and provides ground-truth data directly from companies rather than filtered through third party APIs (e.g. Argyle.com).

However, without the sustained and long term efforts of WIE to mobilise their members to engage with DSARs, this would not have been possible. The difficulty of pressuring and negotiating with platforms themselves to respond to access requests, often requiring complaints to regulators and legal representation, cannot be underestimated. The responsibilities of handling requests, and safely stewarding the large amounts of personal data they contain, are also substantial. The use of mass DSARs as a research tool should therefore not be undertaken lightly by researchers. Researchers may need to partner with specialist organisations who can take on some of this responsibility. In addition, privacy-preserving and collectively governed research infrastructure is needed, to ensure

participants can co-determine how their data is handled, and how research is turned into action [6, 41, 51].

Given such challenges of DSAR-based audits, platform transparency should also be pursued through regulation and/or collective agreements. One example for large online platforms is the EU Digital Services Act (DSA) Article 40. More specific to rideshare platforms, related initiatives include transparency reports[36], rideshare data release efforts in cities like New York and Chicago<sup>7</sup>, and laws such as Colorado Senate Bill 24-75 requiring platforms to disclose fare, distance and direction information to drivers before they accept a ride. The latter was influenced by FairFare, a recent project similar to this one which measures take rates on rideshare platforms [10]. Similarly, WIE have recently called on UK city governments to establish better public data on ridesharing platforms impact on labour markets, communities and environments, based in part on the analysis above [15].

## 6 CONCLUSION

We undertook a 2-year long participatory action research project to put the idea of worker data science into practice. Driven by the concerns, hypotheses, and strategies of workers and organisers, we undertook an audit of Uber's algorithms, finding empirical evidence for several harmful trends. The results of this audit are already being put into action by WIE, used as evidence in its organising, advocacy and policy work (see, e.g. [15]). We conclude that DSAR-based auditing can play an important role in securing a better future for gig workers in the face of algorithmic injustice.

## REFERENCES

- [1] Jeremias Adams-Prassl. 2022. Uber BV v Aslam: '[W]ork relations... cannot safely be left to contractual regulation'. *Industrial Law Journal* 51, 4 (2022), 955–966.
- [2] Jef Ausloos and Michael Veale. 2020. Researching with data rights. *Amsterdam Law School Research Paper* 2020-30 (2020).
- [3] Tom Barratt, Caleb Goods, and Alex Veen. 2020. 'I'm my own boss...': Active intermediation and 'entrepreneurial' worker agency in the Australian gig-economy. *Environment and Planning A: Economy and Space* 52, 8 (2020), 1643–1661.
- [4] Oliver Bates, Adrian Friday, Julian Allen, Tom Cherratt, Fraser McLeod, Tolga Bektaş, ThuBa Nguyen, Maja Pieczyk, Marzena Piotrowska, Sarah Wise, et al. 2018. Transforming last-mile logistics: Opportunities for more sustainable deliveries. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*. 1–14.
- [5] Abeba Birhane, Ryan Steed, Victor Ojewale, Briana Vecchione, and Inioluwa Deborah Raji. 2024. AI auditing: The broken bus on the road to AI accountability. In *2024 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*. IEEE, 612–643.
- [6] Laura Boeschoten, Jef Ausloos, Judith E Möller, Theo Araujo, and Daniel L Oberski. 2022. A framework for privacy preserving digital trace data collection through data donation. *Computational Communication Research* 4, 2 (2022), 388–423.
- [7] Abel Brodeur and Kerry Nield. 2018. An empirical analysis of taxi, Lyft and Uber rides: Evidence from weather shocks in NYC. *Journal of Economic Behavior & Organization* 152 (2018), 1–16.
- [8] Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*. PMLR, 77–91.
- [9] Dana Calacci and Alex Pentland. 2022. Bargaining with the black-box: Designing and deploying worker-centric tools to audit algorithmic management. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (2022), 1–24.
- [10] Dana Calacci, Varun Nagaraj Rao, Samantha Dalal, Catherine Di, Kok-Wei Pua, Andrew Schwartz, Danny Spitzberg, and Andrés Monroy-Hernández. 2025. FairFare: A Tool for Crowdsourcing Rideshare Data to Empower Labor Organizers. *arXiv preprint arXiv:2502.11273* (2025).
- [11] Le Chen, Alan Mislove, and Christo Wilson. 2015. Peeking beneath the hood of uber. In *Proceedings of the 2015 internet measurement conference*. 495–508.

<sup>7</sup><https://www.nyc.gov/site/tlc/about/tlc-trip-record-data.page>, [https://data.cityofchicago.org/Transportation/Transportation-Network-Providers-Drivers/j6wf-834c/about\\_data](https://data.cityofchicago.org/Transportation/Transportation-Network-Providers-Drivers/j6wf-834c/about_data)

- [12] Lina Dencik, Jessica Brand, and Sarah Murphy. 2024. What do data rights do for workers? A critical analysis of trade union engagement with the datafied workplace. *Transfer: European Review of Labour and Research* (2024), 10242589241267006.
- [13] Kimberly Do, Maya De Los Santos, Michael Muller, and Saiph Savage. 2024. Designing Gig Worker Surveillance Tools. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–19.
- [14] Veena Dubal. 2023. The House Always Wins: The Algorithmic Gambification of Work. *Law and Political Economy Project, January 23* (2023).
- [15] Worker Info Exchange. [n. d.]. Dying for data: how the gig economy public data deficit conceals £1.9 billion in wage theft, runaway carbon emissions and a health and safety catastrophe. <https://www.workerinfoexchange.org/post/dying-for-data-how-the-gig-economy-public-data-deficit-conceals-1-9-billion-in-wage-theft-runaway>
- [16] Cailean Gallagher, Karen Gregory, and Boyan Karabaliiev. 2023. Digital worker inquiry and the critical potential of participatory worker data science for on-demand platform workers. *New Technology, Work and Employment* (2023).
- [17] Karen Gregory and Jathan Sadowski. 2021. Biopolitical platforms: the perverse virtues of digital labour. *Journal of Cultural Economy* 14, 6 (2021), 662–674.
- [18] Rie Helene, Qiuqiong Song, Yubo Kou, Xinning Gui, et al. 2024. "At the end of the day, I am accountable": Gig Workers' Self-Tracking for Multi-Dimensional Accountability Management. *arXiv preprint arXiv:2403.19436* (2024).
- [19] Jane Hsieh, Miranda Karger, Lucas Zagal, and Haiyi Zhu. 2023. Co-Designing Alternatives for the Future of Gig Worker Well-Being: Navigating Multi-Stakeholder Incentives and Preferences. In *Proceedings of the 2023 ACM Designing Interactive Systems Conference*. 664–687.
- [20] Eslam Hussein, Purna Juneja, and Tanushree Mitra. 2020. Measuring misinformation in video search platforms: An audit study on YouTube. *Proceedings of the ACM on Human-Computer Interaction* 4, CSCW1 (2020), 1–27.
- [21] Basileal Imana, Aleksandra Korolova, and John Heidemann. 2021. Auditing for discrimination in algorithms delivering job ads. In *Proceedings of the web conference 2021*. 3767–3778.
- [22] Al James. 2021. *The Gig Economy: A Critical Introduction*. By Jamie Woodcock and Mark Graham Cambridge: Polity Press, 2020.
- [23] Michael Katell, Meg Young, Dharma Dailey, Bernease Herman, Vivian Guetler, Aaron Tam, Corinne Bintz, Daniella Raz, and PM Krafft. 2020. Toward situated interventions for algorithmic equity: lessons from the field. In *Proceedings of the 2020 conference on fairness, accountability, and transparency*. 45–55.
- [24] Kalle Kusk and Claus Bossen. 2022. Working with wolt: an ethnographic study of lenient algorithmic management on a food delivery platform. *Proceedings of the ACM on Human-Computer Interaction* 6, GROUP (2022), 1–22.
- [25] Min Kyung Lee. 2018. Understanding perception of algorithmic decisions: Fairness, trust, and emotion in response to algorithmic management. *Big Data & Society* 5, 1 (2018), 2053951718756684.
- [26] Min Kyung Lee, Daniel Kusbit, Evan Metsky, and Laura Dabbish. 2015. Working with machines: The impact of algorithmic and data-driven management on human workers. In *Proceedings of the 33rd annual ACM conference on human factors in computing systems*. 1603–1612.
- [27] Shuhao Ma, John Zimmerman, Sarah E Fox, Valentina Nisi, and Nuno Jardim Nunes. 2024. "My Sense of Morality Leads to My Suffering, Battling, and Arguing": The Role of Platform Designers in (Un)Deciding Gig Worker Issues. In *Proceedings of the 2024 ACM Designing Interactive Systems Conference*. 3501–3514.
- [28] Nicholas Martindale, Alex J Wood, and Brendan J Burchell. 2024. What do platform workers in the UK gig economy want? *British Journal of Industrial Relations* (2024).
- [29] Karl Marx. 1867. The General Law of Capitalist Accumulation. In *Capital: Critique of Political Economy, Volume I*, Frederick Engels (Ed.). Progress Publishers, Moscow, Chapter 25. Original work published in 1867.
- [30] Morgan Meaker. 2024. Drivers Are Rising Up Against Uber's 'Opaque' Pay System. *Wired* (2024). <https://www.wired.com/story/drivers-are-rising-up-against-ubers-opaque-pay-system/>
- [31] Morgan Meaker. 2025. 'Impossible' to make ends meet, Uber drivers say. *BBC* (2025). <https://www.bbc.co.uk/news/articles/c1elg0267p6o>
- [32] Akshat Pandey and Aylin Caliskan. 2021. Disparate impact of artificial intelligence bias in ridehailing economy's price discrimination algorithms. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. 822–833.
- [33] Christina Purcell and Paul Brook. 2022. At least I'm my own boss! Explaining consent, coercion and resistance in platform work. *Work, Employment and Society* 36, 3 (2022), 391–406.
- [34] Rida Qadri and Catherine D'Ignazio. 2022. Seeing like a driver: How workers repair, resist, and reinforce the platform's algorithmic visions. *Big Data & Society* 9, 2 (2022), 20539517221133780.
- [35] Inioluwa Deborah Raji, SASHA COSTANZA Chock, and J Buolamwini. 2023. Change from the outside: Towards credible third-party audits of ai systems. *Missing links in AI governance* 5 (2023).
- [36] Varun Nagaraj Rao, Samantha Dalal, Eesha Agarwal, Dana Calacci, and Andrés Monroy-Hernández. 2024. Rideshare Transparency: Translating Gig Worker Insights on AI Platform Design to Policy. *arXiv preprint arXiv:2406.10768* (2024).
- [37] Alex Rosenblat and Luke Stark. 2016. Algorithmic labor and information asymmetries: A case study of Uber's drivers. *International journal of communication* 10 (2016), 27.
- [38] Trebor Scholz. 2016. Platform cooperativism. *Challenging the corporate sharing economy*. New York, NY: Rosa Luxemburg Foundation 435 (2016).
- [39] Sebastian Klovig Skelton. 2023. Uber introduces dynamic pricing algorithm in London. *Computer Weekly* (2023). <https://www.computerweekly.com/news/365531767/Uber-introduces-dynamic-pricing-algorithm-in-London>
- [40] Jake Stein and Dana Calacci. 2022. Workers Collective Data Access Rights. (2022).
- [41] Jake ML Stein, Vidminas Vizgirda, Max Van Kleek, Reuben Binns, Jun Zhao, Rui Zhao, Naman Goel, George Chalhoub, Wael S Albayaydh, and Nigel Shadbolt. 2023. 'You are you and the app. There's nobody else': Building Worker-Designed Data Institutions within Platform Hegemony. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–26.
- [42] Latanya Sweeney. 2013. Discrimination in online ad delivery. *Commun. ACM* 56, 5 (2013), 44–54.
- [43] TFL. 2023. Taxi and private hire demographic statistics Dec 2023. (2023). <https://content.tfl.gov.uk/tph-demographic-stats-dec-2023.pdf>
- [44] Alessandro Nicolò Tirapani and Hugh Willmott. 2023. Revisiting conflict: Neoliberalism at work in the gig economy. *Human Relations* 76, 1 (2023), 53–86.
- [45] Michael Veale, Reuben Binns, and Jef Ausloos. 2018. When data protection by design and data subject rights clash. *International Data Privacy Law* 8, 2 (2018), 105–123.
- [46] Michael Walker, Peter Fleming, and Marco Berti. 2021. 'You can't pick up a phone and talk to someone': How algorithms function as biopower in the gig economy. *Organization* 28, 1 (2021), 26–43.
- [47] Miranda Wei, Madison Stamos, Sophie Veys, Nathan Reitingier, Justin Goodman, Margot Herman, Dorota Filipczuk, Ben Weinschel, Michelle L Mazurek, and Blase Ur. 2020. What Twitter knows: Characterizing ad targeting practices, user perceptions, and ad explanations through users' own Twitter data. In *29th USENIX Security Symposium (USENIX Security 20)*. 145–162.
- [48] Christo Wilson, Avijit Ghosh, Shan Jiang, Alan Mislove, Lewis Baker, Janelle Szary, Kelly Trindel, and Frida Polli. 2021. Building and auditing fair algorithms: A case study in candidate screening. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. 666–677.
- [49] Angie Zhang, Alexander Boltz, Chun Wei Wang, and Min Kyung Lee. 2022. Algorithmic management reimaged for workers and by workers: Centering worker well-being in gig work. In *Proceedings of the 2022 CHI conference on human factors in computing systems*. 1–20.
- [50] Angie Zhang, Rocita Rana, Alexander Boltz, Veena Dubal, and Min Kyung Lee. 2024. Data Probes as Boundary Objects for Technology Policy Design: Demystifying Technology for Policymakers and Aligning Stakeholder Objectives in Rideshare Gig Work. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–21.
- [51] Rui Zhao, Naman Goel, Nitin Agrawal, Jun Zhao, Jake Stein, Ruben Verborgh, Reuben Binns, Tim Berners-Lee, and Nigel Shadbolt. 2023. Libertas: privacy-preserving computation for decentralised personal data stores. *arXiv preprint arXiv:2309.16365* (2023).
- [52] Zoe Zwiebelmann and Tristan Henderson. 2021. Data portability as a tool for audit. In *Adjunct proceedings of the 2021 ACM international joint conference on pervasive and ubiquitous computing and proceedings of the 2021 ACM international symposium on wearable computers*. 276–280.

## 7 ENDMATTER

### 7.1 Author Contributions

The project was led by Reuben Binns and Jake Stein who conducted workshops, interacted with workers one-on-one, and performed data transformation, programming, and analysis. Reuben Binns led preparation of the original draft while Jake Stein was on leave. Siddhartha Datta assisted with the analysis of predictability of pay. Max Van Kleek and Nigel Shadbolt provided writing review and supervision.

### 7.2 Acknowledgements

This work was supported by the Oxford Martin School Ethical Web and Data Architectures in the Age of AI project. Of course, it would not have been possible without the advocacy of James Farrar and Cansu Safak of Workers Info Exchange and all the workers who contributed. WIE were also supported by Open Society Foundation, Digital Freedom Fund, Trust for London, and Omidyar Network. Thanks to Abdurzak Hadi and other members of the App Drivers and Couriers Union for their collaboration with the workshops. Thanks to Anton Ekker of Ekker Law. Thanks to Jeremias Adams-Prassl, Six Silberman, and the iManage project at the University of Oxford for creating a vibrant and insightful community of scholars of algorithmic management which have greatly contributed towards our understanding of this field. Thanks to Abby Gilbert and the Institute for the Future of Work for supporting early development of the DSAR-based auditing concept. Thanks also to Vidminas Mikucionis, Hilde Weerts, Tony Curzon Price, and anonymous reviewers at ACM CHI'25 and FAccT'25 for feedback on various drafts.

### 7.3 Competing Interests

The authors report no competing interests.

### 7.4 Positionality Statement

Our positionality to this research is layered given our interaction with Uber, workers, and WIE. As researchers, our interest in the greater patterns of algorithmic management and pay visible in aggregate data could have led to different prioritization when compared to workers' individual interests in individual data. This was balanced with the interests of collective advocacy brought by WIE, who are motivated both by advocacy for workers on an individual basis and a collective one. As researchers from the University of Oxford, we are aware of our institution's deep ties to colonialism, while many of the primary subjects (and, we hope, beneficiaries) of this research are often immigrants to the United Kingdom from formerly colonized nations.

### 7.5 Ethical Considerations Statement

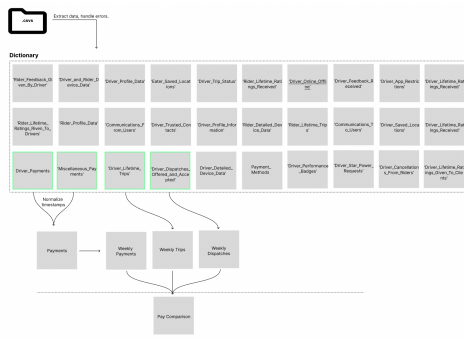
This work raised various ethical concerns. One is around the privacy of drivers whose data was included in the study. As explained in 3.2, the data is highly detailed and longitudinal, containing in some cases nearly a decade of working activities. We went to great efforts to protect such data, including: writing custom software to enable the initial data cleaning, preprocessing and extraction to take place on-premises within WIE; pseudonymising identifiers

through hashing; and stripping away any unnecessary fields. We also spent significant time discussing and formalising the governance around the data, including working with our institutional review board, drawing up a data transfer agreement, and working with WIE to update their data protection impact assessment and security controls prior to sharing.


Another concern is around the potential for retaliation by the platform against individual workers or WIE. To mitigate this, we carefully discussed the remit of the study and potential scenarios. We also consulted legal experts to assess whether there would be any grounds for legal challenge, e.g. around intellectual property / trade secrets, concluding that such risks were low.

### 7.6 Adverse Impact Statement


In addition to providing scholarly insights, an important motivation for this work is to help bring about justice for workers in the gig economy. Despite this, there could be unintended / adverse impacts of our research. One might be that if advocacy based on this research succeeds in improving conditions for workers, this comes at a cost for other stakeholders. For instance, customer fares could be raised even higher, to pay for higher driver take rates. Another possibility is that, in order to reduce the problem of oversupply identified here, some workers are terminated. Finally, if forced to improve pay and conditions, the platform might decide to remove itself from the UK market, putting jobs at risk and removing a service from customers. These risks are difficult to anticipate and mitigate. However, in this work, we have attempted to show how broader efforts towards justice in the gig economy could ultimately benefit society more generally. For instance, there can be common cause between drivers and customers whose pay and prices are both being squeezed by the platform. Regarding the potential loss of work for existing drivers, we would advocate for alternative public mobility platforms, based around human needs rather than profit, which could enable the more equitable sharing of work between existing drivers.



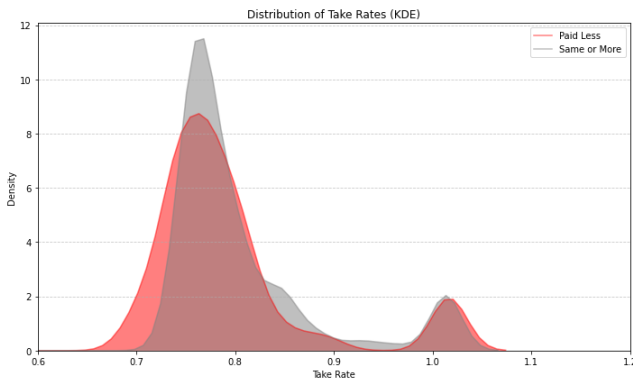
**Figure 8: Mapping data pipeline from a typical DSAR response for one aspect of our analysis (pay comparison, detailed in section 4.1)**



**Figure 11: Accept Rates for Lower-Paid vs Same/Higher-Paid**



**Figure 9: Left: Hours Worked, Lower-Paid (Orange) vs Same/Higher Paid (Blue); Right: Years of Experience for Lower-Paid (Red) vs Same/Higher-Paid (Grey) Drivers**



**Figure 10: Take Rates for Lower-Paid vs Same/Higher-Paid**

**8 APPENDICES**

	Y	Y-1	Y-2	Y-3	Y-4	Y-5	Y-6	Y-7	Y-8	Y-9	Y-10
2014	0.873										
2015	0.654	0.531									
2016	0.721	0.654	0.366								
2017	0.772	0.747	0.532	0.028							
2018	0.809	0.804	0.770	0.440	-0.189						
2019	0.821	0.820	0.801	0.747	0.235	-0.545					
2020	0.830	0.829	0.823	0.794	0.697	-0.161	-1.121				
2021	0.827	0.826	0.823	0.809	0.781	0.625	-0.707	-1.890			
2022	0.828	0.826	0.823	0.820	0.804	0.752	0.562	-0.899	-2.159		
2023	0.830	-3.133	-1.269	-0.794	-1.033	-4.499	0.690	0.472	-1.028	-2.339	
2024	0.831	0.831	-3.082	-1.243	-0.773	-1.010	-4.434	0.683	0.464	-1.033	-2.359

**Table 2: Predicting price at year  $Y$  using years leading upto the year  $1, \dots, Y - n$**