Safeguarding Data

The Data Consensus and the Public Good in Children's Social Services

Jennifer Cearns and Hannah Knox

Abstract

Over the past decade, data-driven systems have transformed how UK public services engage with the population. This ethnographic study, conducted in one of the United Kingdom's poorest boroughs, investigates the implementation of a machine learning data system for identifying at-risk children in need of safeguarding. It examines the concept of a 'data consensus' amongst council workers pursuing the 'public good'. Within this consensus, the article explores contrasting stances on data's potential, namely that of data scientists, who use it as a tool to predict human behaviour and prevent harm, and social workers, who resist the notion that a person is predictable and therefore incapable of change. This reveals the political implications of these opaque algorithmic systems as various groups harness data in the name of public benefit.

Keywords: algorithm, bureaucracy, data, local government, neoliberalisation, public good, social welfare

In the summer of 2020, a controversy erupted in the United Kingdom regarding the use of an algorithm to predict A-Level grades for students who were unable to sit their exams due to the COVID lockdown. Rather than relying on teachers' grade predictions, the government decided to delegate grade decisions to a standardised algorithm, the effect of which was to downgrade the predictions of nearly forty per cent of teachers, with dramatically deleterious effects for some students. The algorithm disproportionately affected high-achieving students from underperforming schools, whose predicted results were sometimes lowered by two or three grades. Ultimately, the algorithm used to restore grading balance and determine students' final grades was widely regarded to have been both inconsistent and unfair (Smith 2020).[1] By the end of August 2020, the algorithm was discontinued, and teacher-assessed grades were reinstated.

Whilst this particular use of algorithms resulted in a highly publicised government U-Turn, the A-Level algorithm fiasco also served to shed light on the increasing use of algorithms by an array of government functions in the United Kingdom from visa processing to policing and welfare governance (Bright et al. 2019). Following the A-Level controversy, reports and newspaper articles surfaced describing a quiet technological revolution that had been underway in UK local authorities.[2] These reports indicated that algorithms and machine-learning (ML) technologies were being implemented for various functions, such as predicting the likelihood of debt non-payment, categorising social housing applicants, identifying potential harm to children and detecting fraudulent activity in housing and benefits claims.[3]

As the introduction to this special issue demonstrates, the use of statistics, measurement and quantification in government is nothing new. However, the rise of algorithmic decision-making in government appears to be unsettling the existing uses of data by the state, raising new questions about what constitutes appropriate or inappropriate uses of citizens' data. One of the findings of the reports on local government use of algorithms following the A-Level saga was that some local authorities were disbanding their use of algorithms, as they were concerned about their cost, their effectiveness and their unforeseen consequences.[4] Others were still enthusiastically embracing algorithms as a tool to help them resolve profound challenges facing local government whilst also remaining concerned that their use of algorithms did not create ethical problems or unforeseen consequences like those generated by the A-Level algorithm (Vogl 2021).

Just as the A-Level debacle was making headlines in the United Kingdom, our own ethnographic research within one council was also underway, exploring how new approaches to datafication are (re)shaping bureaucratic approaches to public sector service provision. This article delves into this parallel moment of digital destabilisation of governance to consider the causes and fortunes of the embrace of datafication in public services. There have already been some powerful critiques of overly celebratory approaches to data systems (Amoore 2020; Benjamin 2020; Eubanks 2018; González 2017; Kotliar 2020; Rouvroy and Stiegler 2016; Ruppert et al. 2017; Sapignoli 2021; Verran 2012). These have drawn attention to the exclusionary, divisive and biased effects of algorithmic governance tools, particularly in the context of welfare, policing and border control, topics which are being actively debated in public discourse about the benefits and dangers of algorithmic governance (Ada Lovelace Institute et al. 2021; Bright et al. 2019; Floridi et al. 2018). And yet the public use of algorithms and data continues to be embraced by local and national governments. Why, we ask, is this the case?

Whilst algorithms have been interrogated ethnographically in terms of their use in corporations (Seaver 2021), medicine (Ruckenstein and Schull 2017), the music industry (Born 2022), the platform economy (del Nido 2022; Timko and Van Melik 2021) and within national government projects such as India's Aadhar (Chaudhuri 2019; Nair 2019; Rao and Nair 2019), less consideration has been afforded to the reasons why local bureaucrats are reframing their work through algorithms and the

kinds of ethical or moralistic discourses that underpin such a shift. If US technology professionals treat algorithms as 'traps' capable of capturing the attention of users (Seaver 2018), Argentinian citizens interpret the algorithms of new taxi services as indices of a political philosophy of freedom (del Nido 2022), and India's Aadhar has demanded a reinterrogation of what constitutes the individual in a country historically preoccupied with the politics of social structure (Nair 2021), what are the particular understandings of algorithms circulating within the setting of local government in the UK? Our answer to this question lies in the way algorithms in local authorities have become tied to concerns about the public good and notions of public interest, which are sustaining the pursuit of new technologies and framing the question of what good data practices look like.

Drawing on ethnographic material on the implementation of a ML system in one local authority in the United Kingdom, our aim with this article diverges from broader critical studies of digital welfare that have sought to assess the impact of artificial intelligence (AI) and machine learning on citizens. Instead, we seek to understand why data is still embraced in public services despite these critiques by recounting the hopes, the negotiations, the ambivalences and the slippages that enable ML systems to be sustained as tools of contemporary governance. We look at the value that data is attributed in the context of public service reform and situate these data discussions within a broader set of conversations about the ongoing challenge of defining what it means to act in the public good (Ballestero 2012; Bear and Mathur 2015; Bernstein and Mertz 2011; Graeber 2015; Herzfeld 1992; Lea 2021). When algorithms and ML systems are introduced into bureaucratic settings, we find that they are necessarily caught up within a broader concern with maintaining and extending good governance and transparency around decision-making. Moreover, understanding their use in local government highlights how they become a medium for potentially divergent ideas about how the public good is enacted (Bear and Mathur 2015; see also Burnyeat and Sheild Johansson 2022).

One approach of recent literature on contemporary forms of bureaucracy has been to draw attention to the neoliberalisation of public services (Collier 2017; Morgen and Maskovsky 2003; von Schnitzler 2015; Wacquant 2012). Whilst discourses of austerity government and cost-cutting form a background to the story we want to tell, we also see data as adding a further dimension to the question of what the public good looks like and how it is enacted. Here, what we find is not just the neoliberalisation of bureaucracy and welfare, but also attempts to find new techniques and tools that can resolve the challenges of addressing the public good in times of austerity politics and limited budgets. Rather than seeing datafication as a manifestation of an audit culture infused with critiques of managerialism, what we found was a turn to data as an answer to assaults on public services, leading to what we term the 'data consensus'.

While we use the term 'consensus', we do not mean that this is a space of total agreement. Rather, we find 'data consensus' to be a useful term that points both to data's embrace as a possible answer to challenging circumstances and to its role as a container for negotiated and provisional conversation about the role of data in

relation to the public good. What we found was that data was accepted as being necessary, even as its specific form and uses were not always agreed upon (or even broadly understood). Just as Hannah Arendt's (2006) focus on how the banality of documentation can alienate bureaucrats from the ultimate goal of their labour through record-keeping regimes, our work points to a mutual commitment to produce a bureaucracy for good rather than evil whilst also recognising the challenges bureaucrats face in attempting to enact this through algorithmic systems which by definition also seek to dehumanise and alienate the citizens whose data they process. This article does not tell a tale of villains and victims, then, but rather describes how a diversity of good intentions can serve to create a data consensus where complex technical systems unfold and eventually become embedded within governance processes.

As we show below, discussions about the data consensus were characterised by what we call a 'yes-but' approach. What we mean by this is that, whilst there was agreement amongst our interlocutors that data is a necessary and important tool of governance, this was accompanied by qualifications that sought to demarcate the boundaries between appropriate and inappropriate uses of data. By delving ethnographically into the data consensus and yes-but narratives, we trace how equivocal positions and subtle ambivalences are swept up and along by an active commitment to using data for good. While an enthusiasm for data sustains investment in novel projects and new data uses, we suggest that the 'but' side of the yes-but discussions about data reveals traces of a more unstable terrain, wherein potential transformations in governance may be taking place.

The Research

This article draws on in-depth research that we conducted over a six-month period from May to October 2020. The research focussed empirically on the experiences of one local authority – that we have pseudonymised as Summertown Borough Council or SBC – based in an economically depressed urban borough in the United Kingdom. Our research explored the local authority's use of AI and machine learning to support three key areas: children's services, housing and their COVID-19 response. Our core focus in this article is primarily on material that was collected regarding the use of the system in children's services.

As outlined above, existing studies of the use of ML and AI in governance processes have often approached these technologies through a critical lens. This has drawn attention to the inequalities, misclassifications, exclusions and forms of discrimination that algorithmic systems seem to exacerbate. Virginia Eubanks' (2018) analysis of the use of algorithms in the United States' welfare system is one such approach. Eubanks' work illustrates how the design of algorithmic systems for assessing risk profiles of welfare applicants replicates the ideology of the American poorhouse, which historically differentiated between the deserving and undeserving poor. The people who are denied healthcare and welfare support due to opaque decision-making processes have no means of challenging them. As a result, those

subject to these systems are both subjugated by them and unable to express the injustices and inequalities that these new algorithms appear to be creating. As with the A-Level case summarised at the outset of this article, the problem that algorithms pose for scholars like Eubanks is that a blunt technical tool is being introduced with the aim of producing neutral decisions but it does so with effects that are discriminatory and therefore unethical. In her argument, the problem is that the justification for the use of algorithms is technocratic; meanwhile, their effects are political.

What interests us about the case that we discuss in this article, however, is that the introduction of algorithms within SBC was, from the outset, framed as an explicitly ethical intervention. Our research came about because SBC, aware of public concerns over the use of algorithms in public services actively sought out a critical assessment of their use increasing use of ML and AI driven data analytics. Some of those closely involved in the council's data unit were concerned that the implementation of a new ML system be done in a robust and ethical way and so they actively sought out research partnerships that would critically evaluate their work. This was more than just a rubber-stamping exercise. Rather, the aim for all involved was that research would help identify any areas or issues of concern around data use, derived from a nuanced, ethnographic analysis of the implications of automated technologies of decision-making on local government practice. It was clear to all involved that these technical systems were also political technologies, both in terms of the promises they implied and the ethical challenges they embodied.

In what follows, we consider this narrative of attempting to create an ethical intervention through AI/ML systems. We start by exploring the reasons given for the need for such a system in local authority settings and unpack the milieu of 'data consensus', wherein everyone we spoke to saw a data-driven system as beneficial for the council, and therefore the public. We then move to the 'but' side of the yes-but discussion to explore the caveats and questions that arose around the use and role of data in the council. We show how qualms about data revolved not around their exclusionary or discriminatory effects, but rather around the question of whether they were overstepping certain boundaries of what data should be doing in public sector work.

The Data Consensus

Our conversations with various council staff unfolded over the summer months of 2020. As they did so, it soon became clear that discussions about data were happening everywhere – in every virtual Microsoft Teams corridor we might have wandered down, data was being mentioned. It also gradually became clear that what people referred to when they talked about 'data' varied. For some, data was spreadsheets of statistics, while for others it was their own working knowledge of the borough and its residents. The one thing upon which everyone seemed to agree, however, was that where data is concerned, more is *more*. Most of our interlocutors

were united in a social imaginary that positioned data as omnipotent, agentive and inherently *a good thing*. The rhetoric of a technological sublime in the ethnographic work we present in this article 'involves hymns to progress that rise like froth on a tide of exuberant self-regard sweeping over all misgivings, problems, and contradictions' (Marx 1964: 207), and in this way we suggest that the data consensus acts socially much like mythology, in so far as it animates individuals on a path towards transcendence, or towards the 'good' (see also Ames 2019; Mosco 1999; Nye 1996). Its use might even be seen as inherently moral; indeed, *not* making 'proper' use of the abundant data the council had in its possession was seen as a missed opportunity, and perhaps even as failure in the council's duty to serve its community in the most effective yet efficient way possible.

When our research began, SBC was exploring the possibility of becoming a pioneer in the use of data to support council work. In early interviews with senior managers, we were told how the council was enthusiastic about data use, evidenced by the presence of an in-house data team that could support different aspects of the council's work. More data meant a richer picture, a clearer and more accurate story and a better-informed decision by the council when a decision needed to be taken. Moreover, amongst all our council participants, more data was unanimously seen as 'the direction of travel'. As one frontline staff member put it: 'We are always in the race for getting as much information as we can . . . The more information we can get in a timely manner, the better'. Data promised the ability to 'really see' what was happening, but more than that, data was considered as able to *tell* what was likely to happen as well. Data did not just *exist* for many of our participants, it had the ability to foresee, to interpret, to go beyond the realm of insight accessible to individual council workers 'on the ground'. Therein lay data's exciting potential to help pressured council workers do more with less.

Intrigued by this enthusiasm for data, one of the first things we sought to understand was why data was deemed so important. There was a tacit agreement amongst everyone that data was simply an obvious part of the milieu or environment that they were operating within. Most discussions around data were underpinned by an awareness that the council was facing significant challenges in its work to support residents. Data was seen as an available resource which could help the council tackle endemic poverty and the challenge of improving the life chances of the borough's residents.

The borough where this fieldwork took place was recognised by council workers as being economically deprived, a fact that was often cited by our interlocutors as a justification for why drastic measures and shifts in the role of local government were required in the first place. The borough's Corporate Plan for 2020–2022, in which councillors and policymakers set out their vision for the area over the coming years, introduced the main themes of the report by setting the context. It highlighted how residents often did not reach their educational or social potential, how they faced problems with crime and how many residents had serious and ongoing health problems that were not being effectively addressed. The plan to galvanise the public provision of services through efficient use of data hinged on the

language of aspiration, acknowledging that the current situation was both unfair and unnecessary and that people in the borough should expect more. Such discussions served as a starting point for all subsequent conversations we had about data and its uses within local government.

For many of our interlocutors, this context also served as justification that the old way of governing was no longer fit for purpose. Desperate times called for desperate measures, and the borough's strategic plan explicitly positioned the centrality of innovative uses of data as a way of modernising the relationship between state and citizen for the public good. It used the language of computation to describe a need to 'reboot' the council, and find new ways of working. This was justified as necessary to prevent the escalation of problems in the borough, and avoid them reaching a tipping point. The promise of a data driven transformation lay in its ability to tackle looming problems related to poverty and health by first anticipating them, and then as a result enabling them to be prevented.

In this council, as in many others, it was clear that the innovative and forward-looking utilisation of big data and ML in government systems and data sets was being positioned as a direct response to austerity measures, localised poverty, and deprivation. There was an understanding that data had, in Jennifer Gabrys' (2014) terms, already become 'environmental'. What we mean by this is that years of data collection, administration, spreadsheets, databases and charts had served to position data as an important part of the context of everyone's working environments. The challenge now was to take this data and to use it to build preventative measures that could minimise poverty and inequality. In this sense, the use of data was itself an ethical stance taken by the council: it was already there, collected and passively stored, and so its mobilisation was a natural step in harnessing existing material and using it in a better way as part of this 'system reboot'.

If SBC was unusual in having quite such a public embrace of data and its promise for public service transformation, the challenges that the council was facing were not unique. Local authorities in the United Kingdom have come under significant pressure in recent years with large budget cuts since 2013 and the introduction of 'austerity' politics (Koch and James 2022). Such challenges are shared with local governments elsewhere, where many countries have seen a widespread neoliberalisation of welfare and local government (Bear and Mathur 2015; Elychar 2012; Forbess 2022). Anthropologists of these changes have charted a move from state-centred provision of welfare during the second half of the twentieth century to an increasing marketisation and managerialisation of public services (Strathern 2000). This has in turn ushered in new techniques of measurement and accountability and, in their wake, more and more data. Sometimes captured under the term New Public Management, the story of the shift in state welfare provision has also been one of increasing marketisation of public services, evidenced by the rise in public–private partnerships, consultancies, accountability regimes, calls for governance through transparency, and the decentralisation and redistribution of government functions into diverse agencies and actors (Graham and Marvin 2001; Morgen and Maskovsky 2003).

Those that we interviewed echoed this characterisation of the circumstances in which they were working. They faced pressures from funding cuts alongside an increased demand for services as a result of austerity, as residents in the borough also struggled to make ends meet. Specifically, social workers in Children's Services were facing mounting pressure due to heavier caseloads, along with concerns that the pandemic would lead to a greater need for support but with fewer resources available to provide it. With financial pressures driving the need for new techniques of cost-cutting, data was no longer just a side effect of accountability processes, but appeared as a potential resource and tool that could be used to help make financial savings.

The data systems being implemented at SBC were thus understood as a pragmatic (and perhaps even ethical) response to cuts in public services. But it was not council employees who were charged with the task of creating such tools, but rather external companies with the expertise at building information technology (IT) systems. At SBC, as in other local authorities, the AI/ML system that they were exploring was to be built through a partnership arrangement with a private technology provider. The AI/ML system itself had been developed by one of several niche firms providing predictive decision-making technologies to local authorities in the United Kingdom. In this case, the provider built a software system on top of existing databases of information from various agencies and departments including the National Health Service (NHS), the local authority, schools and housing associations. The software implemented by the private technology firm was designed to 'read' information from these public databases, anonymising the data with markers and cross-referencing it with similarly marked data from other sources to generate a risk profile for each resident. The aim of the analysis was to generate a 'flag' denoting the risk profile of individuals. The basis upon which the flag was arrived at was not visible to the users of the system (or to us, despite persistent questioning), but it was also not meant to be used as a definitive categorisation. Instead, it was meant as an additional piece of information that would help experts – in this case social workers – to navigate caseloads and to prioritise the most risky or vulnerable cases (see also Vogl 2021; Vogl et al. 2020).

In the case of the children's social care system, the system would generate a report every fortnight. These reports would highlight the twenty cases most in need of review, which would likely be due to changes in risk factors (such as a change in housing, debt, parent(s)' mental health, and so on), and could signal a case that needed more attention ('stepped up') or less resources ('stepped down'). Social workers would use these reports and alerts to prioritise their workload but also to add context that might not always be disclosed by the individuals themselves. For instance, a family might not tell their case worker they had fallen into arrears on housing payments, yet this might be valuable context to the social worker in determining the stability of the child's environment.

The impetus for this enthusiasm for data-sharing was a few high-profile cases where children in need had fallen through the gaps in safeguarding systems. Social workers we spoke to often cited the case of 'Baby P' as an example of what can

happen when critical information is not available or is overlooked. The 'Baby P' case was widely discussed in the British media, as well as in Parliament in 2008–2009, as an example of how children's welfare services failed to save a child suffering abuse, despite numerous interactions with social workers and healthcare staff over a period of months. Following the child's death, a serious case review found that the child might have been saved had the relevant authorities shared information more effectively (Department for Education 2010).

The promise of SBC's system, then, was that it would bring a newfound efficiency and rationality to local government services, solving some of the problems that social workers faced of poor access to data distributed across a range of databases. The council itself could not afford to build such a system, but through a partnership with technological companies, who would learn from this implementation and be able to sell the system to other councils, they had found an opportunity to grapple with fundamental issues of miscommunication, division and lack of connection. As one interviewee working in safeguarding put it:

> If a child known to a local authority is seriously injured or dies, we have a serious case review . . . One of the things that comes up is sharing of info, and communication between professionals . . . And so if [this] is able to help us with that communication, sharing of info, making sure we don't have gaps, I see that as a positive.

Much of the internal support for the system came from the hope that a properly designed and integrated system that could effectively monitor and flag vulnerable children could also protect social workers from increasing caseloads and the risks of mistakes being made. A blog about the data system explained that the number of cases a social worker would need to review would be much more manageable. In addition, more effective referrals were expected to directly impact the council's finances, as child protection through data-sharing would reduce the number entering the costly care system. As one manager put it:

> The insight is worth it on a purely economic basis alone. Every time we take a child into care, that's £125k per year. If you have a few hundred children in care, that liability extends well beyond their 18th birthday. If you take a young child away from someone, you're looking at bills in the millions. Times a couple of hundred, that's huge. If that data gets in front of it and lets us recover the situation, on a pure hard-nosed cost-savings basis, it's a success.

If the council had financial and social reasons to support the development of a data system they saw as being in line with doing a 'public good' with less resources, then the IT companies designing such systems had their own ideas about why the technology that they were building was beneficial. Their focus was less on how it would help this borough and more on how it could become a model for council work that then could be replicated elsewhere. Companies developing algorithms like those being deployed at SBC generally maintain proprietary control over the system's design and function, with a view to selling the same software to other local authorities. The rise of COVID seemed to be amplifying and extending the growth

of 'data-powered' governance that was already underway. One of the operational team members at the council characterised this moment as 'system development on steroids', explaining that 'we have to acknowledge that what we've seen, especially around March and April,[5] has been system development on steroids. Let's not huff about. I've seen how people have reacted to this; it's been really impressive. I was sitting back and enjoying the fireworks'.

All the reasons discussed above culminated in what we experienced as an over-whelming sense of positivity, or *consensus*, towards the idea that digital data should be central to local authority work. This was surprising to us, given the critical litera-ture we were familiar with. We had expected to encounter a more critical appraisal of data analytics, such as concerns about bias or the privatisation of resources that such projects often entail. While there were evidently economic and political reasons for implementing the algorithmic system, the key observation from our ethnographic data was nevertheless that council workers all united around this shared 'data consensus' that positioned (and justified) the introduction of such systems in terms of the 'public good'. Indeed, we found that more council workers voiced concerns about the ethics of not utilising data sets 'properly', and of the missed opportunities that might arise from such oversights, than about potential misuses or implications of surveillance that such data-driven systems might also entail. In this light, the discussions summarised here arguably add a new dimen-sion to the literature on how datafication is used and understood in the context of governance processes.

Although there appeared to be a consensus that 'more is more' where data is concerned, there were still underlying differences in the motivations for this support, as well as some reservations about what constitutes useful data and how one ought to ethically use it. In the following section, we shift our attention from the question of why people were supportive of these novel data systems, to the caveats, ambivalences or 'but' narratives which often accompanied broad support for more and better data. To illustrate this, we look at the concerns and considera-tions of different groups involved in the implementation of the new AI/ML system including the social workers and their managers, and the internal data and techni-cal development team, which was responsible for overseeing the implementation of the new system.

The 'But' of Data

Supportive Data

Amongst the frontline social workers and their managers, discussions increasingly consolidated around the degree to which data might inform and direct decisions regarding the provision of care to vulnerable people. The data system employed within Social Care seemed a particularly high-stakes space of interaction between data, citizens, policymakers and care providers, given the anticipated role of the new system to highlight or filter out families requiring increased attention. There

were particular concerns amongst senior managers in Social Care, who felt that using data to predict and profile the vulnerable was 'overstepping the mark'. 'There's a high risk that we are damaging our relationship with our communities', one senior manager told us, 'and I like to think that there are better ways to get an early indication of people that require help'.

As the council had developed its thinking around data, it had shifted towards emphasising the role of data as a 'support' to social workers in their own professional practice, rather than a displacement of their expertise and professional instincts. In the words of one policy officer interested in data ethics, there was a clear delineation between data as 'decision-supporting' and 'decision-making': 'That's one of the real advantages of [the system we've put in place]', she told us during one meeting; 'it doesn't make decisions, we're not using algorithms to make decisions, we're using them to pull information together'. This 'full picture' of a citizen (or, in the parlance of the local authority, a 'member' or 'customer'), as painted through the joining of multiple data sets, could then support a council worker in making a better-informed decision. 'It's not a decision-making tool, it's a decision-supporting tool. So, it's always a person rather than a machine that would be deciding the outcome. It tells us who's likely to be at risk, and then when that comes to me, it's up to me as a practitioner to look at this and say, "do we need to take action?"', a social worker explained to us. In this regard, the algorithmic system was acting more as a filter to accelerate the rate at which workers could process caseloads, but agency remained with a 'professional' human who would take on the work of adjudicating any decisions.

In the eyes of Jon,[6] another manager whose work was concerned with social care, the ultimate distinction within this delineation of decision-supporting and decision-making came down to a stance on morality. In a service that 'should be all about people', he told us, there's often not much openness towards the use of 'modern technology', but he firmly believed social workers could 'benefit from it if we're still using our moral compass. But there's a really important point in there . . . We are using that data to inform our assessments. The basis for any support we give has to be . . . a well-trained and qualified worker assess[ing] the data and the information to make a professional decision'. Ultimately for the managers of the social work teams, data's ideal role within social care would be to increase the efficiency of caseload management, providing staff with broader information to draw on and helping them navigate massive caseloads by giving them an indication of where to start. Data was fine if it remained the milieu, but not if it overstepped the line into active decision-making and the role of adjudicator.

Contextual Data

Perhaps a more cynical take on the council's emphasis on data's 'supportive' rather than 'decisive' capacities might be that this was arguably the only way in which the utilisation of such systems would be palatable to the social workers on the frontline, many of whom had worked in the area for years and slowly developed

trusting relationships with the families they interacted with. The core of their work was empathic practice, building human relationships and slowly building up a 'rich picture' of a person within the context of their wider life, which came from years of commitment and experience.

'The more we can understand about what's going on for that family, the more helpful it is and the more we're likely to make the correct decision for that child. It's not just data for us, it's about understanding what is actually happening and what that's telling us . . . it's about understanding how that [data] came to be', argued Linda, another manager working in social care. For Linda, who had years of prior experience as a social worker amongst some of the most deprived families in the country, her job wasn't about data, it was about people, and the abstraction of vulnerable people into data was hard to reconcile with her own professional experience. 'We are not here to analyse data . . . [we're here] to listen to people . . . [in] a human interaction that's based on lived experience', she stated in one of our team meetings.

The tension that seemed to arise from this articulation of data versus human interaction seemed to be one of fluidity versus stasis, of predictability versus the messiness of 'actual life'. It was not that social workers were against data per se, but that they were frustrated at the suggestion that they should use data-driven systems in ways that misunderstood the way that they actually worked. As one team member complained in a broader discussion about using the data system to review alerts: 'We're looking at static figures that require an intervention, whereas in social care . . . people don't work that way'. Many of the social workers agreed that this spoke to a more fundamental gulf between how they saw the people they worked with, embedded within a richer contextuality of changeable 'life', and the inevitably stripped back picture of an individual stitched together across various council data sets and systems. 'Part of what we do is about hypothesising situations and saying if things don't change, here is the future risk', newly qualified social worker Yara[7] told us, 'but we have to believe that people are capable of change, and if we don't believe that there's no point. It's not a linear process. You have to believe that you can support families to make change [and] I don't know how much [the predictive system] would recognise that'.

Here the 'but' of data being articulated was not an inherent tendency for it to be biased or reductive, but rather a concern that the expectation that data could flag dangers or risks failed to recognise the fundamentally changeable quality of persons. The very rationale for engaging families over time was to help people improve their fortunes; therefore, a system which seemed to classify the present status of a person on the basis of their past activities or deeds fundamentally misunderstood what the social workers saw as the purpose of their interventions, namely, to improve people's life chances and, in doing so, to enable people to create a new future for themselves that did not simply replicate their past. Here, the problem was not just that there were dangers in data making decisions, but that the very understanding of what constituted intervention and judgement was different for data analysts and social workers. While data systems implied that judgement was a necessary part

of decision-making and would follow seamlessly from the information messages of collated information, the social workers understood their decisions to entail qualities of empathy, memory, interpretation and understanding of people as they changed across the life course.

A few weeks later, on another team call with social workers, the topic of data use in decision-making came up in conversation again. The team was reviewing some cases that had been alerted to them through the system, pulling in risk factors and signalling changes over the past fortnight that would require a 'step-up' or 'step-down' intervention. In this instance, a child had been flagged due to showing 'undesirable behaviour', and the team was discussing their resistance to this term, which they felt was judgemental and not in line with the ethos of their practice. As they looked more closely into the file, Yara pointed out several discrepancies. The file displayed elevated risk flags for the child due to their being from a single parent household (which in fact was not accurate) and reported 'criminal behaviour'. Yara had been in touch with the family only a week before, as they had been victims (not perpetrators) of crime: context which had not been captured within the system alert. This has echoes of Ruha Benjamin's (2019) work, which has demonstrated how moralistic frameworks of what constitutes 'criminality' may become encoded within new technologies. In this case, in addition to illustrating that data can lead to erroneous classifications, it also illustrates how the very idea of data-driven decision-making implies that data should be a tool for judgement – something that the social workers did not subscribe to. In this case, the social workers at SBC were perplexed (and concerned) as to where this language might even have come from. 'It feels quite judgy, this alert', said Yara, peering at her screen. 'It says "criminality is appearing to become an issue"; well actually, the opposite is happening, and the family has totally bossed lockdown, they've been amazing. Can you imagine if I went round there and knocked on their door on the back of this?!' The team agreed and discussed how the predictive risk alerts might introduce new challenges to be navigated. 'We would want to be working in partnership with everyone who's surrounding the family; the whole professional network around that family would need to be part of our work. We're not just going to accept a piece of data and run with it', summarised Linda, 'and that's the buffer that anything needs to get through before it becomes social work practice'.

Abstracting Data

Iain, a data scientist working on 'cleaning data' for the new system, had a different perspective on the potentiality of predictive data within local governance. For him, the benefits of data analysis lay in an emerging capacity to 'see trends . . . so that when someone presents, we could shorten that initial conversation to get to the nitty gritty'. For Iain the exciting opportunity afforded by a joined-up data system was essentially that more data would be . . . more. If citizens are an amalgamation of interactions with different council systems, and therefore different data points, joining these together would inevitably produce a 'richer picture' of an individual.

Equipped with this knowledge, council workers could then respond to vulnerabilities before they escalated further, and council resources could be assigned more efficiently in response to predicted need.

Such data models articulate citizens and families as relatively static, with risk assuming an almost linear (and inherently predictable) form whereby pre-defined risk thresholds are catalogued and compared to updated data sets to produce flags and alerts to social workers. Yet here, data is also lauded for its capacity to add context and richness to the view the council has of any single individual. The Data Team was operating within a context where local authorities have been criticised in the past for not sharing data between safeguarding services, resulting in harm to vulnerable people, as highlighted above in the famous 'Baby P' case.

In this light, the development of the data system at SBC might be seen as an attempt to join up the formerly fractured data fragments of people, scattered across different services and systems, and instead to create a unified picture of a family and their relationships that could be conjured at the click of a button. This process involves abstracting people into data points, stripping away much of the lived contextuality as bemoaned by the social workers described above, but then also a rehumanisation of the data as the system uses natural language processing to produce a report or flag and interpret this pattern. The striking distinction between how the data scientists and the social workers we spoke to viewed individuals can be summed up as a disagreement about how changes in people's circumstances come about. While those responsible for data processing and analysis emphasised people primarily as data patterns that could steer intervention and thus bring about change, the social workers were trained to view people in the context of their lives and circumstances and consider the role of individual agency in effecting historical and future change. In this way, these two groups of council workers had different professional focuses, and the contrast between these become particularly stark in their approaches towards harnessing 'data'.

Conclusion

This article has sought to move away from a 'villains and victims' framework for thinking about the uses of data systems within local government and welfare provision, and has instead advocated a view of the roles of data that does not privilege a critical stance. By attending to the subtle ambivalences, resistances, distinctions and distortions set out in the ethnographic material above, we attempted to build a picture of how such data systems are developed and promoted in practice.

Our research within SBC showed an overwhelming consensus that data is *good* and that more is *more*. This was sustained by a sense conveyed by our interlocutors that data was the milieu in which they found themselves, a context in which they were operating and a resource they were expected to use. This inherent promise of data was underpinned by numerous wider circumstances: the rapid amplification of remote governance at the outset of the pandemic, a decade of austerity meas-

ures, high-profile public debates about local authority safeguarding failures due to poor information-sharing, specific fiscal and demographic challenges faced by the borough, and grander desires to carve out a new relationship between state and citizens for the twenty-first century. Interestingly, many groups within the council, whilst adamant that 'data was good', also admitted that they personally did not have much need of it whilst assuring us that other teams 'over there' could definitely make use of it. In this regard, the ethnography here resembles observations previously made about bureaucratic organisations, where the locus of responsibility and accountability shifts to always be just over the horizon (Petrakaki 2018).

Crucially, for all our interlocutors each of these justifications for utilising data analytics at scale in local governance was intended to advance the 'public good'. We have sought in this article to acknowledge the enthusiasm and dedication shown by all our interlocutors to do the best they could for the people they represented, often despite trying circumstances and working conditions. Part of the enthusiasm that we have described here may in fact have derived from an ambition harboured by many of our interlocutors to make things better alongside a hope that modern technologies might address some of the weaknesses and challenges so typical of 'old-school' bureaucracies. Data was seen as having the agency and capacity to effect a much needed 'system reboot' of the council in the wake of a global pandemic.

Yet, as we have also seen, what this looked like to different parts of the council varied. This is what we have referred to as the 'yes, but' problematic, whereby different articulations of what data *is*, what it *should do* and indeed what it *can do* come into tension with one another. In particular, we have drawn attention to the way in which many of the tensions we have described in this article arose from ambiguity over the appropriate form of data's agentiveness. When just existing as part of the milieu of social work, and treated as a mute or interpretable resource, many of the dangers associated with data seemed to be downplayed. Interlocutors repeatedly lauded data's ability to not only capture, but also *tell* a story, *paint* a richer picture and *inform* decision-making. Where concern arose, however, was when data seemed to take on a socially inappropriate role as decision-maker, judge or adjudicator of families and their needs. The ethnography we have offered within this article has captured those discussions in their infancy in one such bureaucracy, and followed individuals as they navigated and negotiated these tensions whilst always working to enact 'the public good'.

Jennifer Cearns is a Postdoctoral Fellow in the Department of Anthropology at University College London and an Associate Research Fellow of the Alan Turing Institute. Her current work focusses on artificial intelligence and subjectivity of emotion, particularly within the development of digital technologies to support mental health. Her first monograph, *Circulating Culture*, was published in 2023 with University of Florida Press. E-mail: jennifer.cearns@ucl.ac.uk

Hannah Knox is Professor of Anthropology at UCL. She joined UCL in 2014 after a decade working for the ESRC Centre for Research on Socio-Cultural Change at the University of Manchester. Her book publications include *Roads: An Anthropology of Infrastructure and Expertise*; *Ethnography for a Data Saturated World*; *Thinking like a Climate: Governing a City in Times of Environmental Change*; *Speaking for the Social: A Catalogue of Methods* and *Digital Anthropology*. She was Director of the UCL Centre for Digital Anthropology from 2018 to 2021 and is currently co-editor of the *Journal of the Royal Anthropological Institute*. E-mail: h.knox@ucl.ac.uk

Notes

1. https://blogs.lse.ac.uk/impactofsocialsciences/2020/08/26/fk-the-algorithm-what-the-world-can-learn-from-the-uks-a-level-grading-fiasco/ (accessed 8 March 2023).
2. https://www.theguardian.com/education/2020/aug/17/a-levels-gcse-results-england-based-teacher-assessments-government-u-turn (accessed 9 March 2023).
3. https://www.govx.digital/data/local-authorities-achieving-results-with-ai-roll-outs (accessed 9 March 2023); https://www.local.gov.uk/publications/using-predictive-analytics-local-public-services (accessed 9 March 2023); https://www.theguardian.com/society/2020/oct/28/nearly-half-of-councils-in-great-britain-use-algorithms-to-help-make-claims-decisions (accessed 9 March 2023).
4. See, for example, https://www.theguardian.com/society/2020/aug/24/councils-scrapping-algorithms-benefit-welfare-decisions-concerns-bias (accessed 9 March 2023).
5. The first national lockdown in the United Kingdom at the beginning of the COVID pandemic was March–April 2020.
6. All personal names used in this article are pseudonyms.
7. All personal names used in this article are pseudonyms.

References

Ada Lovelace Institute, AI Now Institute and Open Government Partnership, *Algorithmic Accountability for the Public Sector*. 2021. August. https://www.opengovpartnership.org/documents/algorithmic-accountability-public-sector/.

Ames, M. G. 2019. *The Charisma Machine: The Life, Death, and Legacy of One Laptop per Child*. Cambridge, MA: MIT Press.

Amoore, L. 2020. *Cloud Ethics: Algorithms and the Attributes of Ourselves and Others*. Durham, NC: Duke University Press.

Arendt, H. 2006. *Eichmann in Jerusalem: A Report on the Banality of Evil*. London: Penguin.

Ballestero, A. 2012. 'Transparency in Triads'. *Political and Legal Anthropology Review (PoLAR)* 35 (2): 160–166. https://doi.org/10.1111/j.1555-2934.2012.01196.x.

Bear, L. and N. Mathur. 2015. 'Introduction: Remaking the Public Good'. *The Cambridge Journal of Anthropology* 33 (1): 18–34. https://doi.org/10.3167/ca.2015.330103.

Benjamin, R. 2019. 'Assessing Risk, Automating Racism'. *Science* 366 (6464): 421–422.

Benjamin, R. 2020. *Race after Technology: Abolitionist Tools for the New Jim Code*. Oxford: Oxford University Press. https://oar.princeton.edu/bitstream/88435/pr1901zf44/1/AssessingRisk.pdf.

Bernstein, A. and E. Mertz. 2011. 'Bureaucracy: Ethnography of the State in Everyday Life'. *Political and Legal Anthropology Review (PoLAR)* 34 (1): 6–10. https://doi.org/10.1111/j.1555-2934.2011.01135.x.

Born, G. 2022. *Music and Digital Media: A Planetary Anthropology*. London: UCL Press.

Bright, J., B. Ganesh, C. Seidelin and T. Vogl. 2019. *Data Science for Local Government: Challenges and Opportunities*. April. Oxford: Oxford Internet Institute. http://dx.doi.org/10.2139/ssrn.3370217

Burnyeat, G. and M. Sheild Johansson. 2022. 'An Anthropology of the Social Contract: The Political Power of an Idea'. *Critique of Anthropology* 42 (3), 221–237. https://doi.org/10.1177/0308275X221120168.

Chaudhuri, B. 2019. 'Paradoxes of Intermediation in Aadhaar: Human Making of a Digital Infrastructure'. *South Asia: Journal of South Asian Studies* 42 (3): 572–587. https://doi.org/10.1080/00856401.2019.1598671.

Collier, S. J. 2017. 'Neoliberalism and Rule by Experts'. In H. Vaughan and L. Wendy (eds), *Assembling Neoliberalism: Expertise, Practices, Subjects*. London: Springer Link, 23–43.

Department for Education. 2010. *Serious Case Review 'Child A'*. https://www.gov.uk/government/publications/haringey-local-safeguarding-children-board-first-serious-case-review-child-a (accessed 5 September 2023).

Elyachar, J. 2012. 'Next practices: Knowledge, infrastructure, and public goods at the bottom of the pyramid.' *Public Culture* 24 (1) 109–129.

Eubanks, V. 2018. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. New York: St. Martin's Press.

Floridi, L., J. Cowls, . . . and E. Vayena. 2018. 'AI4People – An Ethical Framework for a Good AI Society: Opportunities, Risks, Principles, and Recommendations'. *Minds and Machines* 28: 689–707. https://doi.org/10.1007/s11023-018-9482-5.

Forbess, A. 2022. 'Redistribution dilemmas and ethical commitments: Advisers in austerity Britain's local welfare state.' *Ethnos* 87 (1) 42–58.

Gabrys, J. 2014. 'Programming Environments: Environmentality and Citizen Sensing in the Smart City'. *Environment and Planning D: Society and Space* 32 (1): 30–48. https://doi.org/10.1068/d16812.

González, R. J. 2017. 'Hacking the Citizenry? Personality Profiling, "Big Data" and the Election of Donald Trump'. *Anthropology Today* 33 (3): 9–12. https://doi.org/10.1111/1467-8322.12348.

Graeber, D. 2015. *The Utopia of Rules: On Technology, Stupidity, and the Secret Joys of Bureaucracy*. Brantford, UK: Melville House.

Graham, S. and S. Marvin. 2001. *Splintering Urbanism*. London: Routledge.

Herzfeld, M. 1992. *The Social Production of Indifference*. Chicago: University of Chicago Press.

Koch, I. and D. James. 2022. 'The State of the Welfare State: Advice, Governance and Care in Settings of Austerity'. *Ethnos* 87 (1): 1–21. https://doi.org/10.1080/00141844.2019.1688371.

Kotliar, D. M. 2020. 'Data Orientalism: On the Algorithmic Construction of the Non-Western Other'. *Theory and Society* 49 (5–6): 919–939. https://doi.org/10.1007/s11186-020-09404-2.

Lea, T. 2021. 'Desiring Bureaucracy'. *Annual Review of Anthropology* 50: 59–74. https://doi.org/10.1146/annurev-anthro-101819-110147.

Marx, L. 1964. *The Machine in the Garden: Technology and the Pastoral Ideal in America*. Oxford: Oxford University Press.

Morgen, S. and J. Maskovsky. 2003. 'The Anthropology of Welfare "Reform": New Perspectives on US Urban Poverty in the Post-Welfare Era'. *Annual Review of Anthropology* 32 (1): 315–338. https://doi.org/10.1146/annurev.anthro.32.061002.093431.

Mosco, V. 1999. 'Cyber-Monopoly: A Web of Techno-Myths'. *Science as Culture* 8 (1): 5–22. https://doi.org/10.1080/09505439909526528.

Nair, V. 2019. 'Governing India in Cybertime: Biometric IDs, Start-Ups and the Temporalised State. *South Asia: Journal of South Asian Studies* 42 (3): 519–536. https://doi.org/10.1080/00856401.2019.1598122.

Nair, V. 2021. 'Becoming Data: Biometric IDs and the Individual in "Digital India"'. *Journal of the Royal Anthropological Institute* 27 (S1): 26–42. https://doi.org/10.1111/1467-9655.13478.

del Nido, J. M. 2022. 'Uber Mobilities, Algorithms, and Consumption: Politicizing Ethical Reflection'. *Mobilities* 17 (5): 729–744. https://doi.org/10.1080/17450101.2022.2114843.

Nye, D. E. 1996. *American Technological Sublime*. Cambridge, MA: MIT Press.

Petrakaki, D. 2018. 'Re-Locating Accountability through Technology: From Bureaucratic to Electronic Ways of Governing Public Sector Work'. *International Journal of Public Sector Management* 31 (1): 31–45. https://doi.org/10.1108/IJPSM-02-2017-0043.

Rao, U. and V. Nair. 2019. 'Aadhaar: Governing with Biometrics'. *South Asia: Journal of South Asian Studies* 42 (3): 469–481. https://doi.org/10.1080/00856401.2019.1595343.

Rouvroy, A. and B. Stiegler. 2016. 'The Digital Regime of Truth: From the Algorithmic Governmentality to a New Rule of Law'. *La Deleuziana* 3: 6–29. http://www.ladeleuziana.org/wp-content/uploads/2016/12/Rouvroy-Stiegler_eng.pdf.

Ruckenstein, M. and N. D. Schüll. 2017. 'The Datafication of Health'. *Annual Review of Anthropology* 46: 261–278. https://doi.org/10.1146/annurev-anthro-102116-041244.

Ruppert, E., E. Isin and D. Bigo. 2017. 'Data Politics'. *Big Data and Society* 4 (2): 1–7. https://doi.org/10.1177/2053951717717749.

Sapignoli, M. 2021. 'The Mismeasure of the Human: Big Data and the "AI Turn" in Global Governance'. *Anthropology Today* 37: 4–8. https://doi.org/10.1111/1467-8322.12627.

Schnitzler, A. V. (2016). 'Democracy's Infrastructure: techno-politics and protest after apartheid'. Princeton, NJ, Princeton University Press.

Seaver, N. 2021. 'Seeing Like an Infrastructure: Avidity and Difference in Algorithmic Recommendation'. *Cultural Studies* 35 (4–5): 771–791. https://doi.org/10.1080/09502386.2021.1895248.

Smith, H. 2020. 'Algorithmic Bias: Should Students Pay the Price?' *AI and Society* 35: 1077–1078. https://doi.org/10.1007/s00146-020-01054-3.

Strathern, M. 2000. *Audit Cultures: Anthropological Studies in Accountability and Ethics and the Academy*. London: Routledge.

Timko, P. and R. van Melik. 2021. 'Being a Deliveroo Rider: Practices of Platform Labor in Nijmegen and Berlin'. *Journal of Contemporary Ethnography* 50 (4): 497–523. https://doi.org/10.1177/0891241621994670.

Verran, H. 2012. 'The Changing Lives of Measures and Values: From Centre Stage in the Fading "Disciplinary" Society to Pervasive Background Instrument in the Emergent "Control" Society'. *The Sociological Review* 59: S2: 60–72. https://doi.org/10.1111/j.1467-954X.2012.02059.x.

Vogl, T. M. 2021. *Artificial Intelligence in Local Government: Enabling Artificial Intelligence for Good Governance in UK Local Authorities*. Oxford: Oxford Internet Institute. https://dx.doi.org/10.2139/ssrn.3840222.

Vogl, T. M., C. Seidelin, B. Ganesh and J. Bright. 2020. 'Smart Technology and the Emergence of Algorithmic Bureaucracy: Artificial Intelligence in UK Local Authorities'. *Public Administration Review* 80: 946–961. https://doi.org/10.1111/puar.13286.

Wacquant, L. 2012. 'Three Steps to a Historical Anthropology of Actually Existing Neo-liberalism'. *Social Anthropology / Anthropologie Sociale* 20 (1): 66–79. https://doi.org/10.1111/j.1469-8676.2011.00189.x.

Zuboff, S. 2019. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. New York: Public Affairs.