Socially meaningful transparency in data-based systems: reflections and proposals from practice

Abstract

- Purpose

The purpose is to present proposals to foster what we call a *socially meaningful transparency* practice that aims to enhance public understanding of data-based systems through the production of accounts that are relevant and useful to diverse publics, and society more broadly.

- Design/methodology/approach

Our proposals emerge from reflections on challenges we experienced producing written and visual accounts of specific public sector data-based systems for research purposes. Following Ananny and Crawford's call to see limits to transparency practice as "openings", we put our experience into dialogue with the literature to think about how we might chart a way through the challenges. Based on these reflections, we outline seven proposals for fostering socially meaningful transparency.

- Findings

We identify three transparency challenges from our practice: information asymmetry, uncertainty and resourcing. We also present seven proposals related to reduction of information asymmetries between organisations and non-commercial external actors, enhanced legal rights to access information, shared decision making about what gets made transparent, making visible social impacts and uncertainties of data-systems, clear and accessible communication, timing of transparency practices and adequate resourcing.

- Originality

The paper contributes to existing debates on meaningful transparency by arguing for a more social, rather than individual, approach to imagining how to make transparency practice more meaningful. We do this through our empirical reflection on our experience of doing transparency, conceptually through our notion of socially meaningful transparency, and practically through our seven proposals.

- Social implications

Socially meaningful transparency aims to enhance public understanding of data-based systems. It is therefore a necessary condition not only for informed use of data-based products, but crucially for democratic engagement in the development of datafied societies.

Transparency; data systems; algorithmic systems; information asymmetry; uncertainty; resourcing

Introduction

The adoption of new forms of data-based systems in different contexts, frequently for the purpose of algorithmic processing and decision making, is often accompanied by confusion or controversy about their inner workings and societal consequences. To overcome confusion, commentators often call for greater transparency. Clear information about how data-based systems work, it is proposed, will enable scrutiny of them, and individuals can then make informed decisions about their participation in them (European Parliament, Directorate-General for Parliamentary Research Services, 2019; Pasquale, 2015). This is also what members of the public appear to want: surveys exploring attitudes of the UK public regularly find that people want to know who has access to data about them and where it is stored (e.g. Aitken et al., 2016; Hartman et al., 2020; Kennedy et al., 2021). In other words, they want data systems to be more transparent.

However, there is growing criticism of efforts to make data systems transparent. Ananny and Crawford (2018), amongst others, have highlighted the limitations of transparency as a solution to the controversies, harms and problems that can result from the introduction of data-based systems in different domains. A central point in their argument is that the ideal of transparency assumes that knowing is possible as a result of seeing inside systems, thus it privileges seeing over understanding. Furthermore, not only is knowing a data system challenging – systems developers themselves do not always fully understand how systems work, and they change over time – but more importantly, there is a danger that transparency becomes confused with, or is assumed to lead directly to, the ability of individuals to hold organisations to account for their data practices. As such, transparency is equated with giving power to individuals to contest systems, a view that Obar (2020) challenges when he argues that transparency is of little value if access to information is not packaged with "tools for turning that access to agency" (2020: 3). Yet, rather than enhancing agency, Draper and Turow (2019) argue that existing transparency practices are actively encouraging digital resignation – a term they use to describe the "feelings of futility" many people have regarding corporate data uses which are cultivated through firms' use of "placation, diversion, jargon and misnaming" (2019: 1830) as a means of obfuscation in their transparency practices.

These critiques of existing transparency efforts raise the question of whether there is a form – or various forms – of transparency practice that moves beyond individualist approaches to positively contribute to enhancing public understanding and democratic governance of data-based systems. Public understanding is important, we suggest, because it is a necessary condition for informed use of data-based products, but also for democratic debate and decision making about the development of datafied societies. In this paper, we argue that the answer to this question is yes. To work towards this end, we develop a number of proposals for what we call a *socially meaningful transparency* practice that aims to enable the production of accounts of data-based systems that are relevant and useful - in other words, meaningful - to multiple and diverse publics. We choose the term data-based systems, rather than e.g. algorithmic systems, because while algorithmic systems are all data-based, there are data-based systems that are socially relevant which are not algorithmic in nature. The target of socially meaningful transparency is therefore the social relevance of the system, rather than its technical functioning. By 'meaningful transparency' we mean a transparency practice that foregrounds the needs and interests of those who require information to be transparent for them to understand data-based systems, rather than centring the interests of data-systems developers or others who are engaged in transparency for the purpose of compliance or public relations. By 'socially meaningful', we mean a form of transparency that decentres individual users and focuses instead on the societal aspects of transparency practice. That is, a form of transparency which enables diverse publics to "understand and become comfortable with the

strengths and limitations of the system [and] overcome a reasonable fear of the unknown" (Felzmann et al, 2019, drawing on Weller (2017)), and which is embedded within democratically informed policymaking.

Our argument for socially meaningful transparency builds on reflections on research that we undertook on the *[anonymised]* project. In this research, we produced written and visual accounts of specific data-based systems at the Department for Work and Pensions (DWP), British Broadcasting Corporation (BBC) and National Health Service (NHS) - all UK public bodies. These accounts were used for elicitation purposes in later stages of the project in discussions with members of the public, which we report elsewhere ([anonymised]). These discussions with the public are not the focus of this paper. Rather, in this paper we focus on our production of these written and visual accounts of data-based systems, in which making data systems transparent and understandable for participants was one of our aims. We put this experience of 'doing transparency' into dialogue with literature about data system transparency in order to develop our proposals for a socially meaningful transparency practice.

We start by outlining debates about transparency as an ideal and as a practice in the context of data-based systems, including situating this contribution in the wider literature on meaningful transparency. Next, we describe our own transparency practice as part of the research we undertook, prior to discussing three key challenges that we encountered when trying to produce meaningful accounts of data-based systems: information asymmetry, uncertainty, and resourcing. Building on these reflections about our efforts, we then introduce seven proposals that aim to address these challenges and which together aim to foster socially meaningful transparency of data-based systems. We envisage that our proposals will be of value to people working in data policy, practitioners involved in efforts to make data-based systems transparent and understandable, and people seeking to foster the engagement of diverse publics and their representatives in debate about datafied societies.

Questioning transparency practices

When new data-related controversies surface, the media and privacy campaigners often call for more transparency about data uses. However, what is often missing from such calls is detail about what transparency means in practice. Often there is advocacy for public documentation about the functioning of systems (e.g. Gebru et al., 2021; Mitchell et al., 2019; Pasquale, 2015). The expectation is that putting such documentation into circulation will facilitate scrutiny of it and enable individuals to make informed decisions about their participation in a given data system, taking advantage of opt-in and opt-out mechanisms, if they exist (e.g. Kwok and Chan, 2021; Turilli and Floridi, 2009). It is also claimed that transparent processes can build trust in an organisation's data systems and practices (e.g. Fournier-Tombs, 2021). This assumption can mean that the purpose of transparency becomes building public trust in data systems and the institutions that develop and use them. This expectation of transparency as an enabler of trust is often a simplification of what is required to foster trust, particularly amongst diverse publics whose relationships to societal institutions and technologies are shaped by their experiences of discrimination (Benjamin, 2016).

The ideal of transparency is grounded in a liberal democratic political framework that has its roots largely in the USA (Crain, 2018; Obar, 2020), as well as the social democratic political frameworks of Nordic countries. US Associate Justice of the Supreme Court (1916 to 1939) Brandeis' frequently cited assertion that "sunlight is the best disinfectant", by which he meant making government practices visible to the public is the best means of rooting out corruption, has become something of a clarion call for transparency advocates worldwide. Calls for greater government transparency began to take hold in many countries in the late 20th century, followed in the early 21st century with

3

the expansion of a variety of open government practices, including opening up government datasets for public scrutiny and commercial re-use (Bates, 2014). In much the same way as calls to make the inner workings of data systems transparent can assume this will result in increased accountability for those developing and implementing systems, these 'open government' initiatives assumed making administrative data transparent and re-usable would fuel an army of "armchair auditors" that would hold public authorities to account (Gov.uk, 2011).

More recently, transparency advocacy has expanded into the domain of the "black boxed" data and algorithmic systems that are now a core component of our everyday lives. The term "black box" has its roots in Science and Technology Studies (Pinch and Bijker, 1987), and in this context refers to the opacity of many algorithmic and data systems. Calls to open algorithmic black boxes have been led by prominent US legal scholars such as Frank Pasquale (2015). They have also been embedded into new regulations such as the EU's GDPR, which gives individuals the right to information about uses of their personal data in automated decision-making about and profiling of them. Similar calls have emerged from within the fields of data systems and data science. The Netherlands-based Responsible Data Science initiative, for example, coined the acronym FACT (fairness, accuracy, confidentiality, and transparency) to refer to what it perceives as key principles for responsible data science. Transparency here is understood to address the question: "how to clarify answers [produced by data systems] so that they become indisputable?" (van der Aalst et al., 2017). Around the same time in 2018, what is now the ACM Conference on Fairness, Accountability, and Transparency took place for the first time in the USA, providing a platform for researchers addressing, among other things, issues of algorithmic explainability and audit (Facctconference.org, 2018).

However, the belief that transparency can address the challenges posed by opaque datafication and algorithmic processes is increasingly called into question. In contrast to Pasquale's (2015) proposition that transparency is a "foundational normative value" (Brevini and Pasquale, 2020: 2), commentators such as Heald (2006) and Bannister and Connolly (2011: 5) have long argued that transparency is not an "intrinsic value". Even Pasquale has more recently questioned whether positioning transparency as a "first step towards a more emancipatory" algorithmic governance may be "an easily deflected demand, or actually worsens matters by rationalizing the algorithmic ordering of human affairs" (Brevini and Pasquale, 2020: 2). Here Brevini and Pasquale are pointing to two things. The first is that calls for transparency can be met with what Crain (2018) describes as an appropriation of transparency values as part of public relations efforts. The second is that, in some cases, calling for transparency – rather than e.g. bans on some types of systems – can implicitly signal acceptance of the use of algorithmic applications in high-risk contexts, such as sentencing.

There are further concerns about the merits of transparency. Actually opening 'black boxes' is difficult in practice (Ananny and Crawford, 2018; Kemper and Kolkman, 2019), and different forms of transparency may be required to foster particular ends, such as accountability (Ananny and Crawford, 2018) or meaningful consent (Obar, 2020). People's capacity - in terms of time, expertise and inclination - to make sense of existing transparency documentation ought also to be taken into account (Ananny and Crawford, 2018; Kemper and Kolkman, 2019; Obar, 2020). So should the risks of commercial exploitation of information made accessible as part of transparency practices (Bates, 2014; Mulinari & Ozieranski, 2022). These challenges echo long-standing concerns about the idealism that underlies expectations of a transparency-enabled citizenry, for example within the US political system (Obar, 2020), and have led some to call for a reprioritisation away from transparency towards decommodification i.e. dismantling the treatment of personal information as a commodity (Crain, 2018).

Responding to this challenge, some scholars have suggested ways of reimagining the communicative aspects of transparency practice beyond commonplace information and explanation approaches such as seen in the GDPR.. Felzmann et al (2019), argue that more engagement is needed with the performative factors that influence transparency in practice. By performative they mean, the multitude of "tensions, struggles, and discourses inherent in transparency projects" (p. 8). They bring together the dominant "transparency as -information or -explanation" approaches with the "performative" approach in what they call a "critical contextual approach" (p. 1) that they believe can help foster trust in algorithmic systems. Their approach is based on a relational understanding in which "information provision is conceptualized as communication between technology providers and users" (p. 1).

Within the legal field, researchers have considered how to reduce the burden of making sense of complex information on individuals through delegation processes. For example, Obar (2020) suggests the development of intermediary organisations similar to trustees, lawyers and accountants. People could then delegate to these intermediaries the responsibility to interpret complex information and protect their personal information, he argues. On the other hand, Kaminski (2020) argues that in a context where most proposals for transparency initiatives defer some responsibility to the private sector via "collaborative governance" arrangements, "external input and oversight" by third-party actors to avoid "regulatory capture" by the private sector is necessary. A form of "second-order transparency" that targets the governance regime thus becomes necessary, she argues.

Others have called into question when transparency ought to occur. For example, prospective transparency describes upfront how a system works in general, whereas retrospective transparency provides explanations for automated decisions after they have been made e.g. in the case of deep learning applications (Felzmann et al, 2019). The timing of transparency is also an issue in relation to the process of system design. Often, developers design systems and make varying levels of detail about the system 'transparent' to users after the design is finalised, as part of privacy policies and consent processes. Some may engage users earlier in the design process through user research that may or may not address expectations around data practices. Ananny and Crawford (2018) suggest that an alternative form of transparency practice focused on acceptable data uses could happen earlier in the design process. They point to the US National Environmental Policy Act, which requires that the public is engaged at the outset of any proposed action that may have environmental impact, as an example.

All of this adds up to the need for transparency to be meaningful. As Obar (2020) argues, it is easy to call for more transparency, but much more challenging to do transparency in a way that is meaningful for those it is imagined to enable or empower. The notion of "meaningful transparency" has recently gained traction, and it tends to be defined in one of two ways. The first definition focuses on what types of information should be made transparent, moving beyond basic accounts of systems to also include e.g. the performance of an algorithm, the decision paths that training and trained algorithms took, the training data and its provenance, among other things (Brauneis & Goodman, 2018; Chouldechova, 2020). Other research focuses more on the human aspects of meaningful transparency, particularly how appropriate and useful the transparency practice is for achieving stated aims for a target audience in a given context e.g. Schor et al (2022) and Obar (2022). As Sloane et al (2023) identify, this latter approach defines meaningfulness in the pragmatic tradition as "the property of conveying information that is receivable and useful to a recipient, and that has consequences in that it makes a difference to practice". Pasquale's (2015: 217) suggestion that clear and actionable information that allows citizens to judge the safety of a system, rather than detail about how it functions, is needed can be understood as an example of this type of more

meaningful transparency practice. Existing conceptualisations of meaningful transparency are therefore orientated towards consideration of individual users in context, with only Ramesh et al (2022) beginning to address the question of what collective forms of meaningful transparency may look like.

In the following sections, we build on these discussions of the challenges and limitations of transparency through reflection on our own practice of trying to make specific UK public sector data-based systems transparent. Reflecting on the challenges of doing transparency that we experienced in practice, we follow Ananny and Crawford's (2018) call to see these "limits as openings" towards imagining different forms of transparency practices. We do so by presenting a number of proposals that aim to reduce barriers to socially meaningful transparency of data-based systems. These proposals emerge from us putting the challenges we experienced into dialogue with the literature discussed above. Our paper thus contributes both empirically – through our discussion of our experience of doing and troubling transparency – and conceptually – through our notion of socially meaningful transparency – to current debates about transparency in data-based systems.

Making data-based systems transparent on *[anonymised]*

Our efforts to produce accounts of data-based systems were the first phase of a larger project ([anonymised]) that aimed to understand people's perceptions of how data about them is collected, processed, analysed, shared and used. These accounts were used for elicitation purposes in later stages of the project on which we report extensively elsewhere ([anonymised]). We chose to focus on systems from public sector organisations because their data-based systems and practices increasingly shape everyday life experiences, and yet they had received less attention than high profile commercial data systems at the time of our research. Interested in people's views on the pervasive datafication of everyday life, we identified welfare, media and health as three relevant domains. To produce accounts of data-based systems in the first two domains, we partnered with the UK government Department for Welfare and Pensions (DWP) and the British Broadcasting Company (BBC). For the latter domain, health, we produced two accounts of data systems, one based on information in the public domain, and the other based on research that one of us had undertaken to map flows of the UK's National Health Service (NHS) data (*[anonymised]*). We say more about how these partnerships shaped our transparency practice below.

Our contacts within the BBC and DWP selected the data-based systems about which we produced accounts. At the BBC, our focus was on two projects about personal control over data. The first of these was BBC Box, a prototype device which pulls together data about what users watch or listen to and gives them control over who has access to this data to generate recommendations. The second was BBC Own It, a free app designed by the BBC to support, help and advise children when they use their phones to chat and explore the online world, without adult supervision. The two DWP data-based systems both focused on ways of making it possible to verify identity online. The first was Confirm Your Identity, an identity verification process for welfare payments which enables online identity confirmation. The second, Dynamic Trust Hub, explored a range of issues relating to identity verification, including attribute-based approaches, technology integration and possible security checks. We also produced an account of the NHS COVID-19 Data Store, a national data store to hold data in one place to help organisations responsible for coordinating the UK's COVID-19 response, by drawing on information in the public domain, on government web pages and elsewhere. Finally, we produced an account of a data-based system in an NHS antibiotic prescribing research project which drew on prior research that one of us had undertaken (*[anonymised]*). In the algorithmic transparency literature, there is a distinction made between prospective and retrospective

transparency (e.g. Felzmann et al., 2019). Our practice was prospective - it describes how the system works in general, not a post hoc explanation/rationale for a decision.

As critical scholars of information and communications, we know that it is not possible to produce neutral or objective accounts of phenomena – interpretation takes place in the act of writing, of choosing what words to use, and what aspects of a data-based system to highlight. As intermediaries between partner organisations and research participants, we made decisions about what was more or less important to include in our accounts of data systems to enable understanding and facilitate conversation. We also made decisions on how best to document the data systems in a way that could be meaningful, i.e. relevant and useful, for diverse publics. This is a part of the performativity that Felzmann et al (2019) argue characterises transparency efforts. We aspired to produce simple, but detailed enough, accounts of data systems, to avoid the obfuscation that can occur when providing too much or too little information about data-based systems and processes. We also aspired to be as accurate as possible in our accounts. To do this, we deployed some of the techniques used by Bates et al (2016) to map data flows and frictions, as part of a mobile ethnographic approach they call Data Journeys. This process involved developing detailed knowledge of each data system, through interviews and textual analysis of partner organisation and publicly available documentation. It also involved multiple iterations across the research team, and our contacts at the DWP and BBC checked our accounts of their data-based systems, as did the project advisory board.

For each data-based system we examined, we produced visual representations and written accounts that we then used in the later phases of the research project. Figure 1 shows an example of our visualisation of BBC Own It. The visualisations were produced by a member of the project team, a data studies researcher who is also a data visualisation and information designer (*[anonymised]*).

[Insert figure 1]

Ananny and Crawford (2018) argue that transparency is needed in relation to data-based systems' effects in the world. They describe this as an "actor-network theory of truth" (2018: 984) that prioritises looking across, as opposed to only inside, systems to consider how assemblages of human and nonhuman actors come together to work as a system. As Kemper and Kolkman (2019) similarly argue, algorithms are embedded in social, political and economic settings from which they cannot be separated. Recognising the value of looking across systems, when we presented each account to people in the later stages of the project, we verbally highlighted an alleged potential social harm alongside an alleged potential social benefit for each data-based system, drawing on the opinions of experts in each case. This approach can also be seen in Figure 2, in which we highlight the risks related to uncertainty about who has access to data in the NHS COVID-19 data store and for how long. We did this because we cannot expect people to assess the potential benefits and harms of data uses if they do not know what they are. We recognised that if our approach to transparency was adopted outside of a research context, benefit-harm ratios and weightings should be more accurately reflected.

 [insert Figure 2]

Making public sector data-based systems transparent in a socially meaningful way was challenging. As Felzmann et al (2019) suggest, transparency in practice is performative. It involves a multitude of

"tensions, struggles, and discourses" as well as unintended consequences (p. 7). In what follows, we highlight three particular challenges, which link back to some of the criticisms of the ideal of transparency that we discussed in the previous section, and forward to our proposals for socially meaningful transparency that we present later in the paper. While our research was conducted exclusively with UK public bodies as academic researchers, we believe that insights from our reflections apply to the challenges external third parties, including researchers, are likely to experience when working with organisations to make their data-systems transparent. The three challenges we identified in our practice relate to information asymmetry, uncertainty, and resourcing.

Three transparency challenges – information asymmetry, uncertainty and resourcing

Information Asymmetry

Working in collaboration with partner organisations shaped our transparency practice, in part because there was a significant information asymmetry between us as researchers and our partners. This information asymmetry played out in a number of ways, and shaped which systems we accessed and what we could say about them.

Partner organisations BBC and DWP selected which systems we would focus on, following a process of negotiation in which we requested systems that met the following criteria:

1. the data is about users;
2. the case organisation aims to use the data to enhance users' experiences (individually or collectively);
3. there is some form of algorithmic processing of data;
4. there is a possibility that the actual or proposed use of the data could impact upon different groups of people in different ways (e.g. there is potential for algorithmic bias).

We had little knowledge of the range of data systems in use and development within the partner organisations, which made it difficult to know which specific systems we should ask to focus on. Our efforts to 'do transparency' were therefore driven by our partner organisations' interest in sharing particular data systems with us, rather than being driven by our beliefs about what diverse publics might find relevant or noteworthy. We hoped we might gain access to some major data systems that were significant in reputation, impact, scale and functionality, and therefore socially meaningful for diverse publics. However, most of the systems that we eventually accessed were small in scale and in most cases still under development.

While we had initially intended to examine algorithmic processing (criteria 3) at the DWP, we were informed that there was very little taking place. DWP Data Science teams' own access to DWP data was not straightforward, and they were not commonly using advanced – potentially controversial - data science techniques as is often suggested in the media. The only algorithmic processing in the DWP systems we examined was a potential future development for checking claimants' identities when they log into online systems.

At first glance, it seemed we had been given access to mundane or 'safe' data systems. Making data uses and systems transparent is not without risk, and organisations may feel they need to control access to particular data systems in order to mitigate such risks, which might relate to e.g. privacy, security, and reputation (Ananny and Crawford, 2018; Bannister and Connolly, 2011). An exchange with one of our partner organisations about whether a particular detail about a possible future data use could be described in a report was evidence of concern about such risks. Here we seemed to run into the long-standing challenge of transparency that Bannister and Connolly (2011) identify, that making certain information transparent is handled with excessive caution by the organisation

because it is seen as risky. Yet, in some cases this potentially riskier information could be what is most relevant to meet the information needs of diverse publics who are impacted by systems that may cause harm. All these factors point to the complex considerations at play in relation to what is made transparent by organisations and who decides.

Another factor that influenced data system selection by partners was the willingness of employees within the organisations to help us with our research. This was a project following standard ethical guidelines that required participation to be on the basis of full consent, and it is well recognised that professionals can resist transparency in order to protect their expertise (Ananny and Crawford, 2018; Bannister and Connolly, 2011). However, in the case of both the DWP and BBC teams we spoke to, there was enthusiasm to engage with our research. In the BBC teams, there was clear commitment to embedding ethical and safeguarding practices into the development of their data systems, and to making them transparent by blogging about their development processes. The systems that we investigated showcased what might be understood as these organisations' efforts at state-of-the-art 'responsible data systems' thinking in relation to recommender systems, personal data stores, and secure identity verification. The BBC and DWP's selection of data systems was likely informed, at least in part, by this enthusiasm and what might have been perceived as an alignment of values with our research.

Partner organisations' control over data system selection, resulting from information asymmetry, therefore shaped our transparency practice in significant ways. Nonetheless, all the selected data systems addressed issues that are of growing social significance for people in the UK: digital security and identity, young people's media use and mental health, how personal data are used for recommendation systems, and how National Health Service patient data are shared and used for different purposes including and beyond public health.

In the case of the NHS COVID-19 Data store, we did not work with a partner organisation, but we still experienced information asymmetry. We researched this data system entirely from documents available online, observing obfuscation of particular details and uncertainty generated through partial and changing information. As can be seen in figure 2 above, there was an absence of detail about commercial companies' access to personal health data. In fact, a lawsuit drawn up by OpenDemocracy forced the government to make clearer the nature of private sector companies' involvement and the extent of their access to data, but even then it remained difficult to understand fully these firms' access to data and its duration. The official descriptions of the NHS COVID-19 Data Store can be seen as a superficial transparency practice, which Crain (2018) identifies are sometimes undertaken by data brokers and tech companies as part of a public relations exercise. As above, the information that is withheld in these cases is often that which is of most significance.

Uncertainty

A significant challenge in making data-based systems transparent related to various forms of uncertainty. On our project, some things changed in the short time between collecting information about data systems and checking with partners to ensure we had understood and represented the data system accurately. Some changes were in response to the pandemic. For example, DWP Confirm Your Identity was released sooner than originally intended, because COVID-19 lockdowns made the need for it more pressing. As one informant at the BBC commented related to the project they were working on: "The concrete plan is less concrete because of lockdown".

It was not only the changing systems themselves that caused uncertainty. In the case of the NHS COVID-19 Data Store, it was also the documentation published about the system. Information on government websites about the use and storage of the data in the NHS COVID-19 Data Store was sometimes inaccurate, sometimes incomplete, and changed a number of times while we were trying to produce an account of it. For example, in April 2020 shortly after the store was announced, the message on the NHSX website was that data would be used only for COVID-19 purposes. In June 2020 the narrative changed, and a new message stated that future uses of stored data for the benefit of the population were possible, but it was not clear how, when, and under what circumstances this might happen. Changes continued to emerge in subsequent months, and some of the web pages that we used as sources of information ceased to exist when we returned to them for additional details. These uncertainties were often significant in terms of societal implications.

Conflicting information from different informants was another source of uncertainty. At both the DWP and the BBC, some aspects of the systems we examined were prototypes or plans, rather than existing systems. We encountered inconsistent, and sometimes contradictory, understanding of potential data uses within these systems in documentation and interviews, which is not unexpected, given the differing professional roles and remits of informants in relation to data and the complexity of the data system. This contributed to our uncertainty about how to describe and visualise the data systems accurately. This was not only an issue with planned and prototype systems, but also some existing systems. This is what Ananny and Crawford (2018: 981) refer to as the "technical limitations" of transparency, resulting from a variety of computational and organisational factors. People within organisations often only know a small part of their systems; staff and processes can change regularly; versions are updated; and what informants say about data uses reflects their standpoint and subjective perspective. Kemper and Kolkman (2019) argue that this is one of the challenges of making complex algorithmic models transparent, yet our own research indicated that such challenges also exist in simpler systems that do not include hard-to-explain algorithmic processes.

The involvement of different organisations in DWP's Confirm Your Identity also resulted in complexity and our uncertainty describing the system. In this identity verification system, data was anonymously matched across the DWP and HMRC (Her Majesty's Revenue and Customs) systems, and with a range of private sector organisations, including financial organisations (based on formal data sharing agreements to protect people's data). For a fuller picture of this data system, we would have needed to examine the workings of the system from the perspective of all of these different actors, but there were significant barriers to doing so, including resourcing the additional work and gaining access to the field, especially during a pandemic. As a result, our account of this data system was from the perspective of a single organisation. While our DWP informants were confident that the system was secure, anonymous and legal, they were not always able to provide any detail on what the system looked like from the perspective of third party organisations, due to the limited remit of their professional roles.

All of these points confirm Kemper and Kolkman (2019) and Crain's (2018) claims that data-based systems can be indecipherable even to those who create them, and this is a transparency challenge. We tried to reflect some of these uncertainties in our visual accounts, adopting evolving norms for doing so, such as blurring aspects about which uncertainty exists, as can be seen in Figure 2, our visualisation of the NHS COVID-19 Data Store. In essence, rather than aiming to present a previously opaque 'black boxed' data system with full clarity, and as a stable phenomenon, we were engaged in a practice of greying through the black boxes. That is, we were shining a light with the aim of finding some middle ground between opacity and full clarity.

Resourcing

Producing transparent accounts of data-based systems requires extensive resourcing (Bannister and Connolly, 2011). Kemper and Kolkman (2019) argue that the time required to map the functionality of a given algorithmic model – or, in our case, a data system – is significant and requires domain expertise. Further, how to make processes transparent is not always obvious (Bannister and Connolly, 2011). In order to make our selected data systems transparent, we needed to engage our varied expertise in information management, information science, data visualisation and critical data studies.

In our case, this meant a highly iterative process involving detailed discussions within the team about how best to interpret and communicate specific aspects of the data systems we were observing. After conducting interviews with data scientists, technical architects and project leads at DWP and BBC, team members with expertise in information management and information science interpreted and synthesised interview data and technical documents provided by key informants to produce written descriptions and draft diagrams of the data systems and uses. This was not a straightforward process, particularly when it came to developing a shared understanding across the broader interdisciplinary research team. We had to pause to share and discuss definitions of technical terms, work through misunderstandings about particular technical processes, engage in conversations as a team to try to make shared sense of descriptions of data systems provided to us, decide which details about a data system were important to communicate, and in some cases go back to our key informants for clarification.

We also had to understand the debates around such systems and their potential societal consequences, so that we could include potential social implications of data systems in our accounts, and decide how to present the data use in context. This point also links to uncertainty and information asymmetry. There is often uncertainty about the actual societal consequences of data systems; one reason for this is because their black-boxed character can make it difficult to establish a causal relationship between data system and societal outcome. This uncertainty could, in turn, be seen as a kind of information asymmetry. However, as noted above, looking in detail at a system's inner workings is not always needed in order to know that it has effects in the world. To look across systems, as Ananny and Crawford (2018) put it, to surface their power relations, effects and politics requires the kind of expertise that some of us have as critical data studies scholars.

Alongside, we had to work out a way to create accessible visualisations and descriptions of complex data systems that would engage non-expert, diverse publics in conversation about them. Our data visualiser used our written accounts and rough diagrams to inform the production of visual representations. Making sense of the data systems in order to communicate them visually for different publics was hard, iterative and resource-intensive work and our data visualiser's dual role as a data studies researcher helped her through this process. We also noted that what was technically accurate terminology, according to the interview data or the data systems mapping processes, was not always the most accessible. Here, understanding how people interpret data and how to create descriptions in plain English was a necessary expertise resource. In other words, accounts of data systems translate complex data uses into something relatable to diverse members of the public.

In sum, producing transparent accounts of data systems which facilitate critical and engaged reflection and dialogue, as socially meaningful transparency necessitates, is resource intensive in multiple ways.

Challenges as openings: proposals for meaningful transparency

Ananny and Crawford (2018) argue that the limits of current transparency practices can be seen as openings, or opportunities, for rethinking transparency so that it has the potential to achieve more "meaningful social effects" (2018: 978). Following their proposal, below we put our experience of trying to do transparency into dialogue with the literature discussed earlier in this paper, to think about how we might chart a way through the performative challenges of doing transparency in practice that we and others (e.g. Felzmann, 2019) have noted.

We outline seven proposals for socially meaningful transparency practices, which we see as important steps towards enhancing public understanding of the use of data-based systems, which is necessary to inform consumer choice, but also for democratic debate and decision making about the development of datafied societies. We argue that these proposals would make it easier for organisations to make transparent their data practices in a way that is relevant and useful for diverse publics, as well as society more broadly. Our first three and our final two proposals articulate our emphasis on the social in 'socially meaningful transparency', while proposals four and five build on existing individual in context definitions of meaningful transparency in the literature (e.g. Sloane et al, 2023; Schor et al, 2022).

The first proposal is to **reduce information asymmetries** between organisations and non-commercial third-parties (e.g. researchers, policy makers, journalists, political representatives, service users and members of the public) about the data-based systems that are in operation or proposed within a given organisation. Alongside enhancement of existing rights to access information about data systems, this will enable informed discussions and shared decision making among different groups and communities of people about which systems are of societal relevance and therefore ought to be made transparent, as is arguably appropriate under "collaborative governance" (Kaminski, 2020) of transparency practice (as defined above).

This suggestion builds on our experience of negotiating access to data systems with partners in order to make them transparent. Given the information asymmetry between partner organisations and our team about their data systems, discussions about system selection were directed by partners. This reflects a wider context in which organisations have more power than external third parties in deciding what systems are subject to transparency efforts, beyond basic compliance with existing legal frameworks. If, on the other hand, organisations were required by law to publish lists of data systems in use and under development and legal rights to access relevant information about particular systems were enhanced, this would allow for more public debate and decision making about which systems ought to be made visible as part of a socially meaningful transparency practice. As Bannister and Connolly (2011) observe in relation to process transparency, there are many data processes within organisations and their implications have varying levels of significance for society. It is therefore important to prioritise where resource intensive transparency practices are directed. Reducing information asymmetries about what systems are in use would enable a better targeting of meaningful transparency practices towards systems that are of most relevance to diverse publics and society more broadly.

This proposal is not far-fetched. In 2018, the House of Commons Select Committee on Science and Technology recommended that the UK Government "should produce, publish, and maintain a list of where algorithms with significant impacts are being used within Central Government, along with projects underway or planned for public service algorithms" (Parliament. House of Commons, 2018). Furthermore, in 2021, the UK government published an algorithmic transparency standard for government departments and public sector bodies. The standard, which was being piloted in early 2022, requires "a short description of the algorithmic tool, including how and why it is being used",

12

as well as "more detailed information about how the tool works, the dataset/s that have been used to train the model and the level of human oversight" (Gov.uk, 2021). Beyond the UK, similar algorithm registers are being set up in cities such as Amsterdam and Helsinki (City of Amsterdam, n.d.; City of Helsinki, n.d.), and have been understood as a form of meaningful transparency (e.g. by Murad (2021)). Expansion of such practices to take into account that not all socially relevant data-based systems are algorithmic in nature is a vital next step, as are enhanced legal rights to access relevant information about selected systems. We envisage such mandatory documentation requirements being of value not only for public bodies, which are the focus of this paper, but also for private sector organisations whose data-systems are equally socially meaningful to diverse publics, in line with other research in the field (e.g. Selbst and Barocas, 2018), and this could be examined in further research.

The second and third proposals are closely related and are tied to the aspects of data systems that are made transparent and understandable, and who decides.

The second proposal is to enhance "collaborative governance" (Kaminiski, 2020) of transparency practices by **fostering discussion between organisations and third parties about what aspects of selected data systems should be made transparent,** and use these insights to inform decision making about where to focus socially meaningful transparency efforts**.** The proposal builds on our experience of engaging in dialogue with partner organisations, advisory board members and within our project team about what aspects of systems are likely to be meaningful to people and therefore should be made transparent. This includes the level of technical and system related detail an account ought to go into, balancing the need to make social consequences transparent (the third principle, discussed below) with recognition of when it is valuable to communicate what happens inside data systems. These highly iterative discussions were necessary to agree on what to include in our accounts, and reflect user-oriented definitions of meaningful transparency in the literature. A scaled-up, more socially orientated, version of our approach aimed at enhancing public understanding of data systems could involve other social actors, including representatives of diverse publics, to ensure that related decision making was not undertaken solely by powerful actors. Such an approach could function as a form of second-order transparency, targeting not only the technology and organisations developing it, but also the transparency governance regime (Kaminski, 2020). Again, enhanced legal rights to access relevant information about selected systems would be necessary to support this proposal.

The third proposal brings together, on the one hand, Ananny and Crawford's (2018) argument that it can be more important to look across data systems than examine their inner technical workings, and, on the other, our strategy of highlighting the possible social consequences of data uses: **recognise the potential and evolving societal impacts of a data system** and find ways to enhance understanding of them. This principle builds on our experience of working through the challenge of how to communicate the implications of the systems for diverse publics; an important information need if people are to engage critically with data-based systems. Crucially, this necessitated communicating the potential harms, as well as the benefits, of data systems. As we note above, this requires expertise not just in data science, but in the kind of critical thinking that characterises much of the work in data studies, and methodological expertise in communicating socio-technical phenomena to non-experts.

The fourth proposal is to **avoid obfuscation when communicating about data systems** through provision of too much, too detailed, too complex, or too little information. This proposal builds on

our experience of making sense of and synthesising vast amounts of information about each data system and deciding how to present our accounts accessibly in written and visual form. This requires significant time and appropriate expertise, and is contingent on context. In other words, information might need to be communicated in different ways for different data uses and audiences. This principle also emerges from some of the literature discussed earlier in the paper (e.g. Pasquale, 2015; Sloane et al, 2023; Schor et al, 2022), which argues that what information and how much information to include in transparency accounts are decisions which are as important as deciding which data systems to make transparent.

Our fifth proposal is to **acknowledge and foster understanding of the uncertainty inherent in data systems**. This proposal builds on our experiences of trying to make transparent systems that were still in development, the future uses of which were not stabilised, as well as uncertainty over future social implications. More importantly perhaps, it acknowledges that few data systems are ever stable – rather, they are more commonly in a state of constant interpretative flexibility, a term used within Science and Technology Studies (STS) to characterise socio-technical assemblages for which a range of meanings exist, whose definition and use are still under negotiation (Law, 1987; Wyatt, 1998). This is an essential feature of data systems, identified in the literature we discuss above (e.g. Ananny and Crawford, 2018; Kemper and Kolkman, 2019) and something that needs to be explained in and through transparency practices.

How to represent uncertainty is a subject of considerable debate in literature about data visualisation (Dasgupta et al., 2012; MacEachren et al., 2012; Hullman, 2020), where techniques such as blurring are proposed as solutions. Meaningful transparency practices could engage with these debates. The uncertainty inherent in data systems means it is often more appropriate to aim to 'grey through' rather than fully open the black boxes of some data uses, a term that acknowledges that attempts to make transparent dynamic data systems will only ever be partially successful. We need to embed talking about uncertainty in our transparency practices.

The sixth proposal for socially meaningful transparency is to recognise that **transparency practices can take place at various stages of data system design and implementation**. This point builds on the insights we gained from trying to make transparent systems at different stages of development – from proposals, through to prototypes and fully operational systems. Through this process, we recognised that the timing of transparency acts can bring different challenges, but also potential opportunities for fostering public understanding about and potentially engagement in data system development. This observation brings to mind Ananny and Crawford's argument that "different moments in time may require or produce different kinds of system accountability" (2018: 982).

Finally, the seventh proposal is to **recognise the resources needed for socially meaningful transparency and commit to ensuring they are available.** All of the above commitments would require significant resourcing if they were to be implemented at scale as part of a sustainable data policy intervention. The availability of what resources and how these are valued should also be recognised as inherently political.

Conclusion: enacting socially meaningful transparency in context

The proliferation of data-based systems into more and more aspects of everyday life has led to calls for increased transparency. Yet, data system transparency in practice is often critiqued as

obfuscating (Draper and Turow, 2019), de-contextualised (Ananny and Crawford, 2018) and driven by the needs of powerful actors (Crain, 2018). As such, much existing transparency practice can alienate diverse publics and foster "digital resignation" (Draper and Turow, 2019). While we agree with these critiques of much existing practice, we also maintain that there is nonetheless value in a critically informed transparency practice in relation to data-based systems. Rather than dismissing transparency as always limited (Crain, 2018), we instead argue that a democratic datafied society needs socially meaningful transparency about data-based systems and uses that is integrated into wider efforts to enhance public understanding and debate, and ultimately more democratic ownership of decisions regarding the use of data-based systems by organisations.

Building on our transparency practice on *[anonymised]*, above we set out a series of seven proposals that aim to reduce barriers to socially meaningful transparency practice. As Ananny and Crawford (2018) conclude, it is necessary "to ask to what ends, exactly, transparency is in service". We see our socially meaningful transparency proposals as one among various starting points for working towards more public understanding of, and ultimately more democratic control over, the development of datafied societies. Socially meaningful transparency moves beyond considering meaningfulness only in relation to individuals' specific transparency needs (e.g Sloane et al, 2023; Schor et al, 2022), to focus attention on societal needs in terms of what is made transparent, for whom, how, when and in what ways, and, crucially, who decides. Our intention is to tip the balance of power away from powerful organisations and towards non-commercial third parties including researchers, journalists, policy makers and diverse publics, so that transparency efforts focus on what matters to these publics. In doing so, we need to acknowledge that there is not one public and that transparency must, therefore, be contingent across contexts, publics, and data systems and uses. In this way, socially meaningful transparency departs from other suggestions for addressing the transparency challenge. It is more socially grounded in nature than, for example, individualised approaches to meaningful transparency (Sloane et al, 2023; Schor et al, 2022), or legal approaches such as Obar's (2020) suggestion that individuals could delegate responsibilities to expert practitioners or Pasquale's (2015) idea to make available information that enables individuals to assess the safety of a system.

The policy implications of socially meaningful transparency are varied. The first proposal suggests that an initial step that is needed is to substantially strengthen recent policy efforts in the UK and elsewhere to publish information about the range of data-based systems in use or proposed by organisations across all sectors. Beyond this, additional regulatory interventions could aim to rebalance power in transparency practice in ways that align with the proposals described above. Potential interventions may address, for example, how to foster "collaborative governance" (Kaminiski, 2020) of transparency efforts by engaging diverse actors in decision making about which data systems – and what about them – ought to be made transparent; developing guidance and standards for the production of visual and written accounts of data systems and uses that communicate uncertainty, avoid obfuscation and ensure accessibility for diverse publics; fostering public dialogue and debate about the potential harms of some types of data-based systems; exploring the potential for informed deliberation and decision making with diverse publics about data uses at different stages in the development of data-based systems; and, investing sufficient resources in socially meaningful transparency and related practices that aim to enhance the agency of diverse publics in shaping data-based systems and the future of datafied societies.

References

Aitken M, de St Jorre J, Pagliari C, et al. (2016) Public responses to the sharing and linkage of healthdata for research purposes: a systematic review and thematic synthesis of qualitative studies. *BMC medical ethics* 17(1): 73. DOI: 10.1186/s12910-016-0153-x.

Ananny M and Crawford K (2018) Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *New Media & Society* 20(3): 973–989. DOI: 10.1177/1461444816676645.

Bannister F and Connolly R (2011) The Trouble with Transparency: A Critical Review of Openness in e-Government. *Policy & Internet* 3(1): 1–30. DOI: 10.2202/1944-2866.1076.

Bates J (2014) The strategic importance of information policy for the contemporary neoliberal state: The case of Open Government Data in the United Kingdom. *Government Information Quarterly* 31(3): 388–395. DOI: 10.1016/j.giq.2014.02.009.

Bates J, Lin Y-W and Goodale P (2016) Data journeys: Capturing the socio-material constitution of data objects and flows. *Big Data & Society* 3(2): 2053951716654502. DOI: 10.1177/2053951716654502.

Benjamin R (2016) Informed Refusal: Toward a Justice-based Bioethics. *Science, Technology, & Human Values* 41(6): 967–990. DOI: 10.1177/0162243916656059.

Brauneis R, & Goodman, E (2018). Algorithmic Transparency for the Smart City. *Yale Journal of Law and Technology* 20: 103-176.

Brevini B and Pasquale F (2020) Revisiting the Black Box Society by rethinking the political economy of big data. *Big Data & Society* 7(2): 2053951720935146. DOI: 10.1177/2053951720935146.

Chouldechova, A (2020). Transparency and Simplicity in Criminal Risk Assessment. *Harvard Data Science Review* 2(1).

City of Amsterdam (n.d.) City of Amsterdam Algorithm Register. Available at: https://algoritmeregister.amsterdam.nl/en/ai-register/ (accessed 12 December 2022)
City of Helsinki (n.d.). City of Helsiniki Algorithm Register. Available at: https://ai.hel.fi/en/get-to-know-ai-register/ (accessed 12 December 2022)
Crain M (2018) The limits of transparency: Data brokers and commodification. *New Media & Society* 20(1): 88–104. DOI: 10.1177/1461444816657096.

Dasgupta A, Chen M and Kosara R (2012) Conceptualizing Visual Uncertainty in Parallel Coordinates. *Computer Graphics Forum* 31(3pt2): 1015–1024. DOI: 10.1111/j.1467-8659.2012.03094.x.

Draper NA and Turow J (2019) The corporate cultivation of digital resignation. *New Media & Society* 21(8): 1824–1839. DOI: 10.1177/1461444819833331.

European Parliament, Directorate-General for Parliamentary Research Services (2019) *A Governance Framework for Algorithmic Accountability and Transparency*. LU: Publications Office of the European Union. Available at: https://data.europa.eu/doi/10.2861/59990 (accessed 5 June 2022).

Facctconference.org (2018) 2018 Call for papers. Available at: https://facctconference.org/2018/cfp.html (accessed 5 June 2022).

Felzmann, H, Villaronga, E F, & Lutz, C, et al (2019). Transparency you can trust: Transparency requirements for artificial intelligence between legal norms and contextual concerns. *Big Data & Society* 6(1). https://doi.org/10.1177/2053951719860542

Fournier-Tombs E (2021) Towards a United Nations Internal Regulation for Artificial Intelligence. *Big Data & Society* 8(2): 2053951721039492. DOI: 10.1177/20539517211039493.

Gebru T, Morgenstern J, Vecchione B, et al. (2021) Datasheets for Datasets. *Communications of the ACM* 64(No. 12): 86–92.

Gov.uk (2011) Armchair auditors should act on their rights to inspect town hall books. Available at: https://www.gov.uk/government/news/armchair-auditors-should-act-on-their-rights-to-inspect-town-hall-books (accessed 5 June 2022).

Gov.uk (2021) UK government publishes pioneering standard for algorithmic transparency. Available at: https://www.gov.uk/government/news/uk-government-publishes-pioneering-standard-for-algorithmic-transparency (accessed 5 June 2022).

Hartman T, Kennedy H, Steedman R, et al. (2020) Public perceptions of good data management: Findings from a UK-based survey. *Big Data & Society* 7(1). SAGE Publications Ltd: 2053951720935616. DOI: 10.1177/2053951720935616.

Heald D (2006) Transparency as an instrumental value. In: Hood C and Heald D (eds) *Transparency: The Key to Better Governance?* Proceedings of the British Academy 135. Oxford University Press for The British Academy, pp. 59–73. Available at: https://global.oup.com/academic/product/transparency-the-key-to-better-governance-9780197263839?q=9780197263839&lang=en&cc=gb (accessed 5 June 2022).

Hullman J (2020) Why Authors Don't Visualize Uncertainty. *IEEE Transactions on Visualization and Computer Graphics* 26(1): 130–139. DOI: 10.1109/TVCG.2019.2934287

Kaminski, M E (2020). Understanding Transparency in Algorithmic Accountability. In: Barfield, W. (Ed.) *Cambridge Handbook of the Law of Algorithms.* Cambridge University Press. Available at SSRN: https://ssrn.com/abstract=3622657

Kemper J and Kolkman D (2019) Transparent to whom? No algorithmic accountability without a critical audience. *Information, Communication & Society* 22(14): 2081–2096. DOI: 10.1080/1369118X.2018.1477967.

Kennedy H, Taylor M, Oman S, et al. (2021) *Living with Data Survey Report*. Available at: https://livingwithdata.org/project/wp-content/uploads/2021/10/living-with-data-2020-survey-full-report-final-v2.pdf.

Kwok C and Chan NK (2021) Towards a political theory of data justice: a public good perspective. *Journal of Information, Communication and Ethics in Society* 19(3): 374–390. DOI: 10.1108/JICES-11-2020-0117.

Law J (1987) Technology and heterogenous engineering: the case of Portuguese expansion. In: Bijker W, Hughes T, and Pinch T (eds) *The Social Construction of Technological Systems*. Cambridge: MIT Press, pp. 105–127.

Livingwithdata.org (2022) Data uses. Available at: https://livingwithdata.org/resources/data-uses/ (accessed 13 June 2022).

MacEachren AM, Roth RE, O'Brien J, et al. (2012) Visual Semiotics & Uncertainty Visualization: An Empirical Study. *IEEE Transactions on Visualization and Computer Graphics* 18(12): 2496–2505. DOI: 10.1109/TVCG.2012.279.

Medina Perea I (2021) *Socio-material factors shaping patient data journeys in the United Kingdom*. PhD Thesis. University of Sheffield. Available at: https://etheses.whiterose.ac.uk/30111/.

Mitchell M, Wu S, Zaldivar A, et al. (2019) Model Cards for Model Reporting. In: *Proceedings of the Conference on Fairness, Accountability, and Transparency*, New York, NY, USA, 29 January 2019, pp. 220–229. FAT* '19. Association for Computing Machinery. DOI: 10.1145/3287560.3287596.

Mulinari S and Ozieranski P (2022) Capitalizing on transparency: Commercial surveillance and pharmaceutical marketing after the Physician Sunshine Act. *Big Data & Society*, *9*(1). DOI:10.1177/20539517211069631

Murad, M (2021). Beyond the "Black Box": Enabling Meaningful Transparency of Algorithmic Decision-Making Systems through Public Registers. MSc dissertation. Massachusetts Institute of Technology

Obar JA (2020) Sunlight alone is not a disinfectant: Consent and the futility of opening Big Data black boxes (without assistance). *Big Data & Society* 7(1): 2053951720935615. DOI: 10.1177/2053951720935615.

Obar, J (2022). Defining and Assessing Data Privacy Transparency: A Third Study of Canadian Internet Carriers. *International Journal of Communication* 16: 1688–1712

Parliament. House of Commons (2018) *Algorithms in Decision-Making -Conclusions and Recommendations* (HC 351). London: House of Commons. Available at: https://publications.parliament.uk/pa/cm201719/cmselect/cmsctech/351/35108.htm (accessed 5 June 2022).

Pasquale F (2015) *The Black Box Society: The Secret Algorithms That Control Money and Information*. Cambridge, Massachusetts; London, England: Harvard University Press. Available at: https://www.jstor.org/stable/j.ctt13x0hch (accessed 5 June 2022).

Pinch T and Bijker, W. E. (1987) The Social Construction of Facts and Artefacts: or How the Sociology of Science and the Sociology of Technology might Benefit Each Other. In: Weibe E B, Hughes T, and Pinch T (eds) *The Social Construction of Technological Systems: New Directions in the Sociology and History of Technology*. Cambridge, Massachusetts: The MIT Press.

Ramesh D, Kameswaran, V & Wang, D et al (2022). How Platform-User Power Relations Shape Algorithmic Accountability: A Case Study of Instant Loan Platforms and Financially Stressed Users in India. *Proceedings of the Conference on Fairness, Accountability, and Transparency*, Seoul, Republic of Korea, 21-24 June 2022. DOI: 10.1145/3531146.3533237

Schor, B, Blackwell, A, Singh, J (2022). Making Algorithmic Systems More Meaningfully Transparent for Users Through Visualisation. In: *Designing Interactive Systems Conference*, Virtual Event, Australia, 13-17 June 2022. Association for Computing Machinery. DOI: 10.1145/3532107.3532875

Selbst, A. and Barocas, S. (2018) The Intuitive Appeal of Explainable Machines. *87 Fordham Law Review* 1085.

Sloane M, Solano-Kamaiko, I, & Yuan, J et al (2023). Introducing contextual transparency for automated decision systems. *Nature Machine Intelligence* 5: 187-195. DOI: 10.1038/s42256-023-00623-7

Turilli M and Floridi L (2009) The ethics of information transparency. *Ethics and Information Technology* 11(2): 105–112. DOI: 10.1007/s10676-009-9187-9.

van der Aalst WMP, Bichler M and Heinzl A (2017) Responsible Data Science. *Business & Information Systems Engineering* 59(5): 311–313. DOI: 10.1007/s12599-017-0487-z.

Weller A (2017) Challenges for transparency. Available at: https://arxiv.org/pdf/1708.01870.pdf (accessed 26 April 2023).

Wyatt S (1998) *Technology's arrow : developing information networks for public administration in Britain and the United States*. PhD Thesis. Datawyse / Universitaire Pers Maastricht, Maastricht.