

Governing Data in the Contract State: Outsourcing, Contextual Integrity, and the Limits of Public Accountability

Costanza Fisogni

MSc Management of Information Systems and Digital Innovation
Department of Management
London School of Economics and Political Science

KEYWORDS	ABSTRACT
Keywords Data governance, Contract state, Public-sector data infrastructure, Big Tech outsourcing, Contextual Integrity, Public accountability	The outsourcing of public-sector data infrastructure to large private technology companies has become a structural feature of modern government, yet the governance frameworks through which citizens can scrutinise what happens to their data within these arrangements remain deeply inadequate. This essay argues that the problem is not merely one of legal compliance but of democratic legitimacy and governance architecture. Existing frameworks treat data handling as a procedural matter while leaving unaddressed whether the values governing a processing arrangement remain appropriate to the public sphere from which the data was drawn. Drawing on Nissenbaum's theory of contextual integrity, Solove's taxonomy of privacy harms, Cordella and Willcocks's critique of the contract state, and Morozov's account of technological solutionism, the essay develops a framework for evaluating outsourced public-sector data governance. Three case studies demonstrate how governance failures arise in practice: the Dutch childcare benefits scandal, the Carillion collapse and its Freedom of Information accountability gap, and the NHS Royal Free–DeepMind partnership, with forward reference to the Palantir Federated Data Platform. The essay concludes that adequate data governance requires public bodies to model how outsourcing alters contextual data flows, encode public-context norms into contracts and oversight mechanisms, and anchor accountability in the original public mission regardless of which private party processes the data.

1. Introduction

Public-sector services being outsourced to private enterprises, especially Big Tech companies, has been normalized across liberal democratic countries over the last 15 years. However, the encroachment of private companies into the public sector does more than reshape operational infrastructure. It alters public decision-making on policy issues and, consequently, influences how citizen data is generated, managed, and used (Sharon, 2021b). Indeed, when governments delegate both a service and the data architecture underpinning it

to a private company, governance questions of constitutional significance arise. On what legal basis does an external firm hold and process citizen data? To whom is it accountable? Are the principles that guide the private firm's management of data consistent with the public goals for which the data were initially provided?

Empirical evidence indicates that citizens living in liberal democracies are generally willing to grant governmental agencies access to their personal information in exchange for services. Belle et al. (2021) demonstrated that there exists a 'publicness effect': citizens are significantly

more willing to grant access to their personal data to public rather than private entities. This preference is grounded in structural reasons. According to Nissenbaum's (2010) Contextual Integrity framework, citizens' confidence in governmental agencies' ability to protect their personal information is predicated on a social domain where confidentiality transmission has been a fundamental principle for as long as citizens can recall. Put bluntly, the public sector is bound by inherent norms of purpose limitation and institutional accountability that no private organisation can replicate. However, as Martin (2019) illustrates, this willingness to disclose data should not be mistaken for a complete lack of ongoing expectations. Members of the community retain significant privacy expectations post-disclosure. Any recipient of the data has a positive obligation to respect those expectations, not merely to receive data and deploy it as they see fit.

As private firms take on operational functions within governmental institutions, they acquire custody of citizens' data but not the accountability structures that legitimise public bodies. Taylor (2021, p. 897) articulates the issue succinctly: the escalating participation of the private sector in public governance engenders 'the potential for power relations of domination'. These could only be averted by subjecting companies utilising public data to the same standards as the public sector. However, there are currently no governance frameworks that could identify and enforce such standards.

Outsourcing has expanded dramatically. In 2022/23, the U.K. spent about \$410 billion on purchasing goods and services from the private sector—more than any other country except the U.S. (Jozepa et al., 2025). The UK government's public data warehousing cloud infrastructure is now channelled through the Crown Commercial Service's G-Cloud framework, which facilitated £3.1 billion in cloud purchases in 2023-24 alone,

with its fifteenth iteration valued at up to £14 billion through 2030 (Register, 2025). AWS alone was awarded contracts totalling £894 million from three separate central Government departments in just one day (Computer Weekly, 2024). The exact level of reliance upon these aforementioned tech companies is unclear by design. Procurement opacity, justified as necessary for national security or commercial confidentiality, means the full picture of who processes public data and under what terms is often not in the public domain (Waugh and Hodgkinson, 2021). As Morozov (2013) noted, we have traded off all the checks put in place to monitor our elected officials for an ostensibly cleaner, more efficient technological solution. Blind faith in 'Big Tech' risks undermining the democratic ideal described by Habermas (1996), according to which legitimate political decisions require open public deliberation rather than technical optimisation alone.

This essay argues that breaking into the black-box mechanisms governing outsourcing agreements requires a data governance framework capable of publicly reporting not merely whether data processing is lawful, but whether it is legitimate. Meaning, whether the information flows it creates respect the contextual norms under which citizens originally shared their data. While some may contend that strict legal compliance is sufficient to safeguard citizens' interests in outsourced arrangements, this essay argues that this perspective fails to address the deeper issues of legitimacy and the preservation of public values in data handling. Nissenbaum's (2010) theory of contextual integrity provides that evaluative standard. It does so through three conceptual lenses—the governance of information flow, the democratic accountability of private data processors, and the crowding out of public-sector data expertise. Concluding with three case studies that make the current governance failure empirically concrete.

2. Literature review

2.1 From individual privacy to systemic responsibility

Data governance, as a field, is concerned not merely with the security of individual data records but with the institutional structures, processes, and norms through which data is stewarded. The trust relationship between data holders and the individuals or communities whose data they hold being a relevant issue (Winter and Davidson, 2019). This framing is important because it allows us to view the locus of accountability in continuing fiduciary relationships rather than in individual consent events.

Westin's notion of control—defining privacy as the right of individuals, groups, or organisations to decide for themselves when and in what manner their information will be made known to third parties—remains the normative standard (1967). However, in practical terms, this model is insufficient as a regulatory construct, as individual agency is systematically constrained. Indeed, the literature on the privacy paradox, as documented in a comprehensive review by Kokolakis (2017), reveals a consistent gap between citizens' expressed privacy attitudes and actual behaviour. There is a limitation on individual decision-making capacity resulting from the denial of access to critical services in case of reticence to share personal information. Such structural limitations form the basis for Martin's (2019) call for a paradigm shift from consent-based to consequence-based data obligations. In cases where citizens do not possess the capability to realistically withhold consent, organisations collecting data should bear the burden of ensuring proper use of that data. Martin's contention that organisations have a positive obligation to reasonably determine privacy expectations, not simply to avoid consent violations, provides the conceptual foundation for this paper. Responsibility for proper data stewardship follows the information

itself, not only the organisation that initially collects it.

Solove's (2006) typology of harms arising from privacy loss extends this insight structurally to an organisational level. His catalogue of harm includes surveillance, aggregation, identification, secondary use, exclusion, distortion, and breach of confidentiality. For Solove, the unit of harm is the pattern, not the incident. Outsourcing arrangements exemplify this because they systematically move data across contextual boundaries, generating multiple categories of privacy harm without any single technical breach. For example, data shared for one purpose may later be processed for another. Likewise, data collected within a public administrative context may be analyzed in a commercial environment governed by entirely different incentives. In both cases, the data owner is none the wiser. Thus, multiple Solovean harms arise even though no individual data point has been technically compromised.

The analytical tool that gives these concepts operational significance is Nissenbaum's (2010) Contextual Integrity. According to Nissenbaum, privacy should not be understood as secrecy or control, but as appropriate informational flow. When it occurs appropriately, it reflects the social contract associated with the context in which it was first disclosed; when it occurs inappropriately, it violates those rules. Each social context (healthcare, welfare administration, taxation, security services, etc.) can be described in terms of the characteristics of actors involved in that context (e.g., doctor/patient), the purposes (therapeutic), and the norms governing which types of information can travel across contexts (within limits). A patient providing clinical symptoms to a clinician engages in communication within a context whose norms pertain to a therapeutic goal and a duty of confidentiality. When that information is transmitted to a commercial AI

development context for other uses, it transgresses contextual integrity, even though the patient may have technically given permission for ‘data sharing’ by signing the terms and conditions few really read (Gherardi, 2021). Applied to public-sector outsourcing, Nissenbaum’s contextual integrity provides a regulatory framework for assessing whether the values and objectives underlying an arrangement respect the norms of the public domain from which the data were derived.

In their analysis of the NHS Royal Free–DeepMind case, Winter and Davidson (2019) demonstrate that the transfer of personal health information into a commercial AI development context violated contextual integrity, regardless of contractual agreements intended to regulate this activity. Specifically, they argue that the five dimensions of governance identified in this context—data domains, stakeholders, data value, governance goals and accountability mechanisms—were all affected by transferring Protected Health Information (PHI) from its initial delivery context (patient care) into a commercial data science context with fundamentally distinct norms. This five-dimensional model provides the empirical structure for analysing outsourcing in public sector contexts presented in Section IV.

To operationalise Nissenbaum’s contextual integrity to address governance issues in public-sector outsourcing, three commitments are necessary. First, to model how outsourcing affects data flows. Second, to incorporate contextual norms into contracting arrangements and oversight regimes. Third, to ground accountability in the public mission, regardless of who holds custody of the data. These commitments are detailed further in Section V.

2.2 Is the public sector able to control its ‘IT Destiny’?

Cordella & Willcocks (2010) describe a 30-year

history of creating a ‘Contract State’, in which the public sector has progressively outsourced functions to the private sector. The policy assumed that such outsourcing would lead to improved efficiency and better public services at lower cost. However, the authors argue that the Contract State was created disregarding the role of bureaucracy in delivering public value. Civil service provides more than just administrative capacity; it also embodies democratic, civic, and public service values, including fairness, impartiality, and the common good. These values, they suggest, are being undermined by the way public bodies contract out. The conclusion is that large-scale ICT outsourcing consistently undermines the public body’s ability to maintain control over its processes’ contextual integrity and, in particular, over the concrete technical arrangement, its ‘IT Destiny’ as they put it. Therefore, a review of the value of public bureaucracies and a rebuild of internal innovation and digitalisation capacities is required for any organisation wishing to retain control of its outsourcing arrangements.

This decline in in-house capability is not new. As Mazzucato (2013) clearly illustrates, the foundational technologies of the digital economy, i.e., the Internet, GPS, microprocessors, voice recognition, etc., were all developed by state investment in major R&D programmes, e.g. DARPA, and subsequently sold or licensed to private companies. The state built the infrastructure but eventually transferred ownership of both the infrastructure and the skills/knowledge needed to run it to private organisations. In the UK, the security and intelligence agencies offer a very useful historical example. In 2017, GCHQ stated that some of ‘the best’ information security solutions came from the state (Lock, 2025). However, in 1997, GCHQ’s own Innovation Power was curtailed due to a shortage of IT professionals, driven by high salaries offered by the private sector (ISC, 1997). By 2008, MI5 employed 350

contractors alongside 3,382 employees (ISC Annual Report, 2008). Here we can see one of the earliest publicised failures of private companies' operational logic when applied to public data governance: the Intelligence Services Commissioner's (ISC) 2014 report recorded two serious breaches by contractors who had made unauthorised searches of bulk personal data without proper business justification, resulting in the termination of their contracts (Report of the ISC for 2014, HC 225, 37). The pattern remains consistent: reliance on external capabilities generates dependence and, therefore, governance and government vulnerability, which enable exploitation.

Morozov (2013) identifies the ideological basis of this dependency. He describes the dominant ideology in Silicon Valley, Technological Solutionism, as the view that complex social and political issues can be redefined as neat technical problems with clear and measurable solutions. Consequently, Technological Solutionism leads to a form of systemic failure: an otherwise optimal solution is selected solely because of technical capability, irrespective of whether the values associated with the solution are congruent with the institution employing it. It fundamentally replaces institutional accountability structures with quantitative measures of technical optimisation. Some twenty-five years ago, Neil Postman (1992, p.115) noted a similar dynamic: 'how obediently people accept explanations that begin with the words 'The computer shows...' or 'The computer has determined...' It is Technopoly's equivalent of the sentence 'It is God's will,' and the effect is roughly the same'. This assertion of authority displaces democratic judgement by presenting contestable socio-political choices as objective technical facts.

2.3 Sphere transgression

The failures in governance illustrated here, as well as the transgression in contextual integrity,

are indicative of a deeper structural problem known as 'Sphere Transgression', a concept defined by Sharon (2021a) based on Walzer's (1983) work. This refers to the inappropriate transfer of benefits from one societal sphere—such as health, welfare, safety, or education—to others that operate under different principles of value and distribution. Each sphere has its own mechanisms and norms; for instance, the health system prioritises care based on need and adheres to clinical trust, professional responsibilities, and equal patient entitlement. However, when expertise developed for commercial analysis is applied to oversee digital infrastructure for health data, the digital sphere's imperatives—such as efficiency, optimisation, scalability, and profit—intrude, violating the intended rules of the health sector.

Digitalisation represents a value-laden process, rather than a neutral technological overlay. She draws upon both the Frankfurt School's critiques of instrumental reason (see Horkheimer & Adorno, 1947; Marcuse, 1968) and Heidegger's conceptualisation of technology as an attitudinal approach to the world (1977) to argue that digitalisation spreads a particular value structure across new social domains: standardisation, efficiency, control.

Stevens et al. (2024) articulate these implications into three governance-relevant types of risk associated with digitalisation. Firstly, it facilitates the creation of inequitable financial returns to the public sector from data created within public spaces as that data enhances private sector systems without creating commensurate public good. It therefore shifts the nature of public services toward commercial rather than public value orientation. Concurrently increasing structural dependency, making it unfeasible to withdraw from commercial relationships after governance deficiency becomes evident.

In addition, there is a significant implication for

public spheres: the crowding out of essential, spherical expertise (Sharon & Gellert, 2024). Lanzing (2024) has applied this framework directly to Palantir's NHS expansion. She argued that the danger is not only converting health professionals into analysts but also contaminating the values that define healthcare through applying a technology developed under an exclusive logic of the security sphere.

3. Legal Architecture and Its Structural Inadequacies

The U.K.'s General Data Protection Regulation (GDPR), along with the Data Protection Act 2018, form the relevant legal framework for the processing of data by public sector organisations. As such, public sector entities undertaking tasks performed pursuant to statute will rely primarily upon Article 6(1)(e)—‘the public task’—as their lawful basis for processing. This article allows processing when it is necessary for an entity to perform a function vested in it by virtue of the public interest or official authority. For health-related data (special categories of personal data pursuant to Article 9), the derogations available under Article 9(2)(i) allow for processing to occur if it is necessary for reasons of public health.

While these regulations do possess substance and constitute an earnest attempt to strike a balance between the operational requirements of government agencies and individual rights to privacy, they also contain a structural ambiguity that assumes particular significance in cases involving the use of third-party commercial processors. The lawful basis for data processing under Articles 6 and 9 is retained by the controller, typically a public sector organisation, even when they delegate operations to a private processor. The processor acts on the controller's request without acquiring the lawful basis itself, so the controller remains liable for compliance with the law. This delegation, while theoretically promoting public accountability, poses risks

highlighted by Taylor (2021): private entities, remaining formally subject only to commercial contract law and data protection regulation, may exercise public power over personal data in ways that create the potential for domination in the political and security spheres.

In addition to consent frameworks, the anonymisation argument (that sufficient de-identification renders the risk of harm from disclosure negligible) has been conclusively shown to be false. Specifically, Ohm (2010) demonstrates that anonymisation and utility are inherently contradictory; Narayanan & Felten (2014) show that a majority of rich datasets may be re-identified with sufficient effort. Further, the acknowledgement by NHS England that Full Disaggregation Process (‘FDP’) data would be pseudonymised (as opposed to fully anonymised), coupled with the Information Commissioner's Office's (‘ICO’) admission that pseudonymised data constitutes personal data within the scope of the UK GDPR, establishes that any such risk is not merely speculative. (See The Register, 2023)

4. Case Studies: Data Governance Failures in Practice

The three case studies that follow are not illustrative of isolated failures. Together, they demonstrate different manifestations of the same governance problem. The first case, drawn from the Netherlands, establishes that governance failures associated with Technological Solutionism and sphere transgression can arise even where no outsourcing occurs. Although it falls outside the UK legal framework considered in this paper, it demonstrates that the underlying governance risks associated with Technological Solutionism and sphere transgression are not confined to a single national context. The remaining two UK case studies then examine how these same dynamics are compounded by outsourcing arrangements operating within the UK's legal and institutional framework.

4.1 The childcare benefits case: Technological Solutionism limitation

The sphere transgression framework identifies the intruding *logic* as the unit of harm, not only the identity of the actor that carries it. This distinction allows the Dutch ‘Childcare Benefits’ case to serve as the essay’s first illustration, despite the system being developed and operated in-house. What transgressed was the logic of fraud detection, not an external commercial party. Between 2005 and 2019, the Dutch Tax and Customs Administration incorrectly claimed that nearly 26,000 parents (the majority of whom have dual nationality) committed tax fraud on their children’s daycare claims. To identify these fraudulent claims, the administration used a self-learning risk-assessment algorithm based on parental characteristics, which ultimately led to the repayment of legitimate childcare claims. While the system was implemented in-house at the Belastingdienst, Amnesty International (2021) stated that there was an institutionally racist intruding *logic* in the design of the risk assessment algorithm. It coded foreign-born citizens into the fraud-risk profile immediately upon receipt of birthplace, therefore creating a unit of harm. The Netherlands Court of Auditors (NCA) cautioned about this issue as early as 2006 in its audit report entitled ‘Tax and Childcare Allowances’, stating that the statute governing the process lacked a ‘hardship clause’ to allow for exceptions when procedures produce disproportionate results. (See Oldenhof et al., 2024).

This is exactly what Sharon (2021) describes as a type of ‘sphere-transgression’. It is one of the three types of governance risks identified by Stevens et al. (2024): restructuring public services for efficiency rather than public value. Data analysis was delegated to an anti-fraud technology that follows technological values such as efficiency, optimisation, and scalability and overrides the context-based nuances of

public welfare. This logic was built with strict variables to limit inaccuracies, to the point of excluding legitimate claims that would most likely have been picked up by a human caseworker. The harm created here is not incidental to the system’s design but part of it—the result of applying fraud detection logic to a domain it was never intended to address.

In addition to the breaches identified above, the structural limitations imposed by administrative law (e.g. obligation to provide explanations for decisions; obligation for individual discretion; right to appeal) were rendered ‘empty shells’ against automated responses that cannot be examined or overridden by civil servants. Once an applicant was marked as having a high level of risk for fraud, there would be no explanation of why they were so classified, making it practically impossible to challenge their classification (see Amnesty International, 2021).

Above all, this case illustrates a failure of modelling: the public authority failed to examine how introducing algorithmic decision-making would alter the contextual norms governing welfare data before implementing the system.

4.2 Carillion and the ‘Freedom of Information’ accountability gap

In January 2018, Carillion, one of the U.K.’s biggest government contractors, collapsed. With a combined Public-Private Partnership (PPP) capital portfolio of £59 billion and total contractual payments of £309 billion forecasted over fifty years, the firm was so deeply embedded in public service delivery that its insolvency created crises across hospitals, schools, and housing projects. However, their internal data on contract performance, structural safety, and financial liabilities was shielded from Freedom of Information (FOI) requests on grounds of commercial sensitivity (Waugh and Hodkinson, 2021). This makes a strong case for additional consequences of an unregulated

contract state: the systemic lack of transparency in how public assets and government expenditures are documented (Waugh and Hodkinson, 2021).

Data Governance is not solely about personal data. It is also about the institutional structures and procedures that determine how information on public assets, public monies, and public services is collected, distributed, and made publicly available for oversight. As illustrated by Waugh and Hodkinson (2021) in their comprehensive freedom of information (FOI) field experiment, the opacity of the Private Finance Initiative (PFI) model created a systematic barrier to accessing public outsourcing information. 42% of the 687 FOIA requests submitted to 315 governmental agencies regarding PFI Contracts resulted in absolutely no response. 32% were denied outright. Agencies were denying even simple requests for information, such as the postal codes of all public buildings contracted under PFI, or the amount of financial deduction applied due to contractor failure. All three types of responses were based on claims of commercial sensitivity, claimed cost limits, and assertions that the requested information was 'Not Held' by the requesting agency. They had delegated their institutional knowledge of their own assets to the contracting entity.

The FOIA (Freedom of Information Act) regime, which is supposed to provide a democratic counterweight to this dynamic, has been comprehensively shown to be inadequate. The ICO itself concluded that FOI is 'no longer fit for purpose' in the age of outsourcing (ICO, 2019, p. 6). For example, Carillion's internal fire safety report concerning an NHS hospital built under a Private Finance Initiative (PFI) was held not to be information held "on behalf of" the NHS Trust for the purposes of section 3(2)(b) of the Freedom of Information Act 2000. As a result, the report fell outside the scope of

FOIA and could not be obtained through an FOI request. The Information Commissioner's Office subsequently criticized this outcome as being "clearly out of step with reasonable public expectations," arguing that it exposed a significant transparency gap in the oversight of outsourced public services.

The accountability gap is structural: legislative foundations predicated on clearly demarcated public and private spheres cannot function when those boundaries have been systematically dissolved by three decades of contracting out. Carillion's 2018 collapse made this visible. Information about public money, public assets, and public safety had, through outsourcing, become private property. The principal governance failure here is one of anchoring, as accountability for information generated in the delivery of public functions was no longer tied to the public mission once responsibility shifted to a private contractor.

4.3 NHS–DeepMind and the Palantir FDP: health data as a commercial asset

The Royal Free NHS Trust's collaboration with DeepMind Health, initiated in 2015 and assessed by the ICO in 2016, has been the most extensively documented case of a contextual integrity violation in health data governance in the UK Public Sector.

Under the service agreement signed in September 2015, DeepMind was provided with vast amounts of patient medical data on approximately 1.6 million patients over 5 years for the purpose of developing an acute kidney injury monitoring application called Streams. This was done without informing patients or obtaining their consent. The fact that this project was not disclosed to the public until a Freedom of Information Request was filed in April 2016 further highlights the lack of transparency in the sharing of patients' private information. In accordance with the ICO

investigation conducted in 2017, several violations of the Data Protection Act were identified, including that patients were unaware their data would be used to develop a commercial AI application rather than to support direct clinical care. Therefore, the above represents a specific example of a contextual integrity violation: an informational exchange taking place within a clinician-patient relationship and governed by norms of therapeutic intent and professional discretion was transferred into a commercial AI development environment governed by norms of data collection and aggregation. The transfer principle had changed significantly.

DeepMind subsequently organised ‘Patient Engagement Workshops’ and committed to implementing ‘Verifiable Data Audits’, both intended to provide governance mechanisms to restore what appears to be contextual correctness after the violation has taken place, rather than proposing proactive design elements for governance.

The NHS Federated Data Platform (FDP) contract with Palantir, valued at £480 million and awarded in November 2023, exacerbates each of the issues outlined above. Vlassis (2024) describes how Palantir entered the market: During the COVID-19 pandemic, Palantir offered NHS Foundry for free (£1), intentionally creating dependence. Once the full contract was awarded, NHS England’s procurement documentation acknowledged that switching providers would have created significant disruptions and increased costs due to the migration of operational and institutional knowledge into Palantir’s architecture. Thus, the ‘lock-in’ to Palantir is not based on law but on cognition. The NHS no longer possesses a complete understanding to effectively govern the system upon which they depend and thus reproduce the same capability-erosive dynamics Cordella & Willcocks (2010) identify as the

central weakness inherent in contracting in a contract state.

The sphere transgression issue goes even deeper. Although Palantir retains no raw patient data processed via the FDP, its systems continually improve its commercial AI products through real-time learning from data flowing through them. These commercial AI products have a genealogy tied to surveillance and military intelligence, not to clinical care (Lanzing, 2024). The exact terms governing how this real-time learning takes place, and the platform’s entire operational range across 55 million patients’ records, remain hidden from public review under commercial confidentiality agreements. However, the platform’s function is clear: it serves as the point of access for NHS staff to view, understand, and act on patient data. That which is brought forward, how patients are categorised, and which risk factors are flagged by NHS staff are determined by an organisation whose institutional rationality is centred upon values from a separate sphere from healthcare.

This case demonstrates the combined failure to encode public-context norms into contractual and technical arrangements and to anchor ongoing accountability in the NHS’s public mission rather than in the operational logic of its commercial partners.

5. Towards adequate data governance: modelling, encoding, and anchoring

To ensure adequate data governance, the three cases illustrate that the bodies responsible for public data failed to utilise the three key commitments to implement Contextual Integrity governance (modelling, encoding, and anchoring) before creating dependencies. Therefore, as the conclusions drawn from these diagnoses illustrate, any changes to the institutions’ governance structures would necessarily follow those diagnoses.

Modelling necessitates that public bodies conduct a contextual integrity analysis of new data flows arising from outsourcing arrangements before engaging. This analysis, which includes a Data Protection Impact Assessment under the GDPR, evaluates the parameters established by Nissenbaum (2010) and determines if these data flows comply with the original context's norms. Specifically, it questions whether commercial AI firms are appropriate recipients of clinical information and whether algorithmic fraud-detection systems align with norms of entitlement relating to welfare data. Such inquiries should be resolved before procurement, not during litigation.

Encoding requires that the norms deciphered from the contextual mapping are enforceable both legally and technologically, ensuring that public bodies maintain control over their IT strategies. Contracts must detail what data can be processed, by whom, and for what purposes. Establish oversight mechanisms and audit rights for public funds and data. Last but not least, include an exit strategy for when services are no longer needed.

Finally, anchoring mandates in ongoing institutional oversight ensures accountability remains tied to public purpose, regardless of who manages daily operations. This is currently undermined by gaps in Freedom of Information (FOI) accountability, making it difficult for citizens and representatives to assess the fulfilment of the democratic mandate. Improvements could be made by instituting additional public accountability standards for private entities involved in public functions. However, current EU and UK digital regulations are insufficient for this purpose, as they treat healthcare and welfare sectors as market areas focused on competitive fairness rather than fully shielding them from commercial interests. To address this gap, we recommend introducing legislation that expands on the Freedom of

Information Act (FOIA) by providing statutory authority to mandate extending FOIA obligations to all private entities performing management and/or administrative functions that utilise public sector data and services.

Consequently, any information related to the execution of public functions or the use of public funds will be treated in the same way as information available from public bodies. These transparency and disclosure standards should be embedded in contractual agreements and independently regulated to ensure that the public nature of public organisations is maintained, rather than merely addressing individualised protections for personal data.

6. Conclusion

This essay argued that compliance with data protection regulations is insufficient to ensure the governance of public sector data outsourcing. This is because existing regulatory compliance frameworks do not address the critical question of whether the organisational logic underlying a particular data processing arrangement is consistent with the public domain from which the data was obtained. Citizens provide their personal information to public organisations based on an unspoken, implicit social agreement regarding care, welfare, and democratic representation. Those same expectations cannot be translated into the forms of consent typically found in commercial data regulation. When data is contracted out to commercial companies whose primary goal is efficiency and profitability, these social contracts are broken. As a result, the social contracts governing the provision of data to the public sector are severed when that data is contracted out.

Although the data governance literature provides a common language for understanding what is at issue and what constitutes good governance for data arrangements, there is no evidence that the public values inherent in bureaucratic

organisation can endure indefinitely without proactive efforts to maintain them when contracting out services. Furthermore, although there may be disagreement about the extent to which outsourcing diminishes or eliminates those values, Morozov (2013) identifies an ideological obstacle to maintaining them: the belief that optimising performance through technology is the same as, or better than, democratic governance. In addition to demonstrating the existence of this ideology, the three case studies also show that it has real-world implications for how governments make decisions and govern. Specifically, these case studies demonstrate that this ideology results in algorithmically discriminatory outcomes, opaque decision-making processes, and the transfer of valuable public health data to private corporations for profit.

While the three cases differ in their specific circumstances, all share one important element: the lack of prior examination of whether an outsourcing arrangement preserves the contextual integrity of the data flow to be governed. Together, these findings suggest that outsourcing has expanded more rapidly than the governance frameworks needed to assess its implications.

The response to the problem described above should not involve rejecting either technology or private-public partnerships. However, as Bharosa (2022) cautions, if implemented without cautionary measures, ‘GovTech’ solutions could function as a ‘Trojan Horse.’ This caution is not due to malevolent intent but stems from the growing gap in internal knowledge within legacy public sector institutions, a result of years of outsourcing. The gradual transfer of decision-making authority from public agencies to private-sector entities represents a corresponding loss of digital sovereignty. To avoid this outcome, public sector institutions must develop substantive technical knowledge and capabilities

to model violations of contextual integrity before they occur and create legally binding obligations that represent public-context norms. While these represent aspirational goals for some institutions, they constitute minimum governance standards for democratic nations committed to upholding their obligations to citizens who entrust them with their personally identifiable information.

Taken together, the case studies suggest that democratic accountability has increasingly been displaced by commercial objectives as politically significant questions are reframed as problems of technical optimisation. Regaining this lost ground will require developing governance structures as sophisticated as those they seek to regulate. The three commitments developed above. Modelling contextual flows, encoding public norms into contracts, and anchoring accountability to public mission are not technical refinements to existing procurement practice. They are preconditions for democratic legitimacy in an era where the infrastructure of governance is increasingly held by private actors operating under incompatible values. Taylor (2021) points out that developing such sophisticated regulatory bodies and rules is not a unique challenge. The institutional and legal tools required to achieve this objective already exist. The remaining challenge is to apply them consistently while resisting the siren call to substitute technological optimisation for democratic judgement.

References

- Amnesty International. (2021). Xenophobic machines: Discrimination through unregulated use of algorithms in the Dutch childcare benefits scandal. *Amnesty International*.
<https://www.amnesty.org/en/documents/eur35/4686/2021/en/>
- Belle, N., Cantarelli, P., & Battaglio, R. P. (2021). To consent, or not to consent? The publicness

- effect on citizens' willingness to grant access to personal data in the face of a health crisis. *Journal of European Public Policy*, 28(5), 782–800.
<https://doi.org/10.1080/13501763.2021.1912147>
- Bharosa, N. (2022). The rise of GovTech: Trojan horse or blessing in disguise? A research agenda. *Government Information Quarterly*, 39(3).
<https://doi.org/10.1016/j.giq.2022.101692>
- Computer Weekly. (2024, January 12). AWS secures £894m in cloud spend across three contracts with UK government on same day. <https://www.computerweekly.com/news/366566172/AWS-secures-894m-in-cloud-spend-across-three-contracts-with-UK-government-on-same-day>
- Cordella, A., & Willcocks, L. (2010). Outsourcing, bureaucracy and public value: Reappraising the notion of the "contract state". *Government Information Quarterly*, 27(1), 82–88.
<https://doi.org/10.1016/j.giq.2009.08.004>
- Gherardi, M. (2021). When terms of service are longer than a shakespeare's play: easing the understanding of legal documents through a user experience design method. *DS 110: Proceedings of the 23rd International Conference on Engineering and Product Design Education (EPDE 2021)*.
<https://doi.org/10.35199/epde.2021.54>
- Habermas, J. (1996). Between facts and norms: Contributions to a discourse theory of law and democracy (W. Rehg, Trans.). *MIT Press*.
- Heidegger, M. (1977). The question concerning technology and other essays (W. Lovitt, Trans.). *Harper & Row*.
- Horkheimer, M., & Adorno, T. W. (2002). Dialectic of enlightenment (E. Jephcott, Trans.). *Stanford University Press*. (Original work published 1947)
- ICO. (2019). Outsourcing oversight? The case for reforming access to information law: **Report of the Information Commissioner to Parliament**. *HMSO*.
- Intelligence and Security Committee. (1997). **ISC annual report 1996–97**. *HMSO*.
- Intelligence and Security Committee. (2008). **ISC annual report 2007–08**. *HMSO*.
- Intelligence Services Commissioner. (2014). **Report of the Intelligence Services Commissioner for 2014** (HC 225). *HMSO*.
- Jozepa, I., Brione, P., & Ward, M. (2025). Outsourcing by government departments (Debate Pack CDP-2025-0019). *House of Commons Library*.
<https://researchbriefings.files.parliament.uk/documents/CDP-2025-0019/CDP-2025-0019.pdf>
- Kokolakis, S. (2017). Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon. *Computers & Security*, 64, 122–134.
<https://doi.org/10.1016/j.cose.2015.07.002>
- Lanzing, M. (2024). Travelling technology and perverted logics: Conceptualising Palantir's expansion into health as sphere transgression. *Information, Communication & Society*, 27(15), 2617–2633.
<https://doi.org/10.1080/1369118X.2023.2279557>
- Lock, D. (2025). Public power and private hands: Outsourcing in UK national security law. *King's Law Journal*, 36(3), 571–602.
<https://doi.org/10.1080/09615768.2025.2593116>
- Marcuse, H. (1968). Negations: Essays in critical theory. Beacon Press.

- Martin, K. (2019). Breaking the privacy paradox: The value of privacy and associated duty of firms. *Business Ethics Quarterly*, 30(1), 65–98. <https://doi.org/10.1017/beq.2019.24>
- Mazzucato, M. (2013). The entrepreneurial state: Debunking public vs private sector myths. **Anthem Press**.
- Morozov, E. (2013). To save everything, click here: The folly of technological solutionism. **PublicAffairs**.
- Narayanan, A., & Felten, E. (2014, July 9). No silver bullet: De-identification still doesn't work. *Freedom to Tinker*. <https://www.cs.princeton.edu/~arvindn/publications/no-silver-bullet-de-identification.pdf>
- NHS England. (2023, November 24). Federated Data Platform Check and Challenge Group — minutes and action notes: 24 November 2023. <https://www.england.nhs.uk/long-read/fdp-check-and-challenge-group-minutes-24-november-2023/>
- Nissenbaum, H. (2010). Privacy in context: Technology, policy, and the integrity of social life. **Stanford University Press**.
- Ohm, P. (2010). Broken promises of privacy: Responding to the surprising failure of anonymisation. *UCLA Law Review*, 57, 1701–1777.
- Oldenhof, L., Kersing, M., & van Zoonen, L. (2024). Sphere transgressions in the Dutch digital welfare state: Causing harm to citizens when legal rules, ethical norms and quality procedures are lacking. *Information, Communication & Society*, 27(15), 2704–2720. <https://doi.org/10.1080/1369118X.2024.2358163>
- Postman, N. (1992). Technopoly: The surrender of culture to technology. **Vintage Books**.
- Racism and Technology Center. (2025, April 7). Racist technology in action: The Dutch Belastingdienst's risk analysis model should never have been used. <https://racismandtechnology.center/2025/04/07/racist-technology-in-action-the-dutch-belastingdiensts-risk-analysis-model-should-never-have-been-used/>
- Sharon, T. (2021a). From hostile worlds to multiple spheres: Towards a normative pragmatics of justice for the Googlization of health. *Medicine, Health Care and Philosophy*, 24(3), 341–352. <https://doi.org/10.1007/s11019-021-10006-7>
- Sharon, T. (2021b). Blindsided by privacy? Digital contact tracing, the Apple–Google API and Big Tech's newfound role as global health infrastructure. *Internet Policy Review*, 10(4). <https://doi.org/10.1007/s10676-020-09547-x>
- Sharon, T., & Gellert, R. (2024). Regulating Big Tech expansionism? Sphere transgressions and the limits of Europe's digital regulatory strategy. *Information, Communication & Society*, 27(15), 2651–2668. <https://doi.org/10.1080/1369118X.2023.2246526>
- Solove, D. J. (2006). A taxonomy of privacy. *University of Pennsylvania Law Review*, 154(3), 477. <https://ssrn.com/abstract=667622>
- Stevens, M., Kraaijeveld, S. R., & Sharon, T. (2024). Sphere transgressions: Reflecting on the risks of Big Tech expansionism. *Information, Communication & Society*, 27(15), 2587–2599. <https://doi.org/10.1080/1369118X.2024.2353782>
- Taylor, L. (2021). Public actors without public values: Legitimacy, domination and the regulation of the technology sector.

- Philosophy & Technology*, 34(4), 897–922.
<https://doi.org/10.1007/s13347-020-00441-4>
- The Register. (2023, March 17). NHS threatened with legal action over patient data platform. https://www.theregister.com/2023/03/17/nhs_threatened_with_legal_action/
- The Register. (2025, October 28). UK government inflates G-Cloud framework to £14B. <https://www.theregister.com/on-prem/2025/10/28/uk-government-inflates-g-cloud-framework-to-14b/1146356>
- Tucker, I. (2013, March 9). Evgeny Morozov: We are abandoning all the checks and balances. *The Guardian*. <https://www.theguardian.com/technology/2013/mar/09/evgeny-morozov-technology-solutionism-interview>
- Vlassis, V. S. (2024). More than software vendors: The peculiar case of Palantir's data integration platforms. *Information, Communication & Society*, 1–19. <https://doi.org/10.1080/1369118X.2024.2442399>
- Walzer, M. (1983). Spheres of justice: A defence of pluralism and equality. *Basic Books*.
- Wagh, M. I., & Hodkinson, S. N. (2021). Examining the effectiveness of current information laws and implementation practices for accountability of outsourced public services. *Parliamentary Affairs*, 74(2), 253–275. <https://doi.org/10.1093/pa/gsaa001>
- Westin, A. F. (1967). Privacy and freedom. *Atheneum*.
- Winter, J. S., & Davidson, E. (2019). Big data governance of personal health information and challenges to contextual integrity. *The Information Society*, 35(1), 36–51. <https://doi.org/10.1080/01972243.2018.1542648>