

Article

Untangling the Algorithmic Leviathan: Palantir, the Post-Factual Polity, and the Infrastructural Crisis of AI Governance in Public Administration

Haris Alibašić 

Department of Business Administration, Lewis Bear Jr. College of Business, University of West Florida,
Pensacola, FL 32514, USA; halibasic@uwf.edu

Abstract

This paper extends the post-factual polity framework into AI infrastructure and public administration systems theory. It asks how proprietary analytical platforms alter the state's capacity to produce, audit, and contest the categories through which risk, threat, eligibility, fraud, and deviance become actionable. Using a structured documentary case analysis of Palantir Technologies across United States agencies and allied jurisdictions, the study applies three diagnostic markers—categorical opacity, contestation displacement, and substitutive dependency—to examine the migration of sovereign classification into vendor-controlled infrastructure. The research gap was identified through an integrative review of public administration, AI governance, algorithmic accountability, systems theory, surveillance studies, and Palantir scholarship. The analysis distinguishes AI epistemic capture from ordinary IT vendor lock-in: the former concerns not merely technical dependence or high exit costs but the loss of public capacity to define and contest consequential administrative categories. The paper argues that administrative law, procurement reform, and algorithmic impact assessment remain necessary but insufficient when agencies lack substitutive capacity. It specifies untangling as a systems-level task involving capacity reconstruction, categorical repatriation, contractual restructuring, and procurement reorientation. Hybrid intelligence is advanced as a post-untangling architecture that embeds machine processing within contestable, accountable, and legally governed human judgment. The contribution is diagnostic, methodological, and design-oriented for AI systems governance.

Keywords: post-factual polity; AI governance; algorithmic accountability; ethical AI frameworks; Palantir; hybrid intelligence; public administration; categorization infrastructure; epistemic capture; responsible AI deployment



Academic Editor: Yanyan Xu

Received: 10 June 2026

Revised: 12 July 2026

Accepted: 21 July 2026

Published: 22 July 2026

Copyright: © 2026 by the author.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

The algorithmic Leviathan has arrived [1]. What has not yet been adequately specified, in either the public administration or the AI ethics literature, is that the Leviathan is not principally a public agency deploying AI tools. It is a configuration in which the categorization layer of public decision-making—the institutional apparatus by which risk, threat, eligibility, fraud, and deviance are rendered legible to state action—has been substantially relocated into proprietary analytical platforms whose internal logics are protected by trade-secret law, whose recalibration is governed by vendor processes rather than agency procedures, and, in at least one currently dominant case, whose normative

commitments have been publicly articulated as hostile to the procedural constraints of liberal-democratic governance.

The question is not whether the legal decision-maker formally disappears. In many deployments, the public official remains visible as the nominal actor who approves, denies, investigates, targets, or allocates. The more consequential shift is positional: the platform reorganizes the informational environment before the official deliberates. It structures the classification space, the evidentiary hierarchy, the risk narrative, the menu of plausible actions, and the default path of institutional attention. A tool that appears merely to assist the decision-maker may therefore alter the decision-maker's position inside the decision system. This is the systems problem at the center of the article.

The dominant current instantiation of this configuration in the United States and several allied jurisdictions is the operational footprint of Palantir Technologies. The company's platforms—Gotham, Foundry, AIP, the Maven Smart System, and ImmigrationOS (Palantir Technologies Inc., Denver, CO, USA)—perform categorization, risk-assessment, and enforcement-targeting functions inside the United States Department of Defense, the Federal Bureau of Investigation, the Internal Revenue Service, Immigration and Customs Enforcement, the United Kingdom National Health Service, the police services of multiple German federal states, and a growing roster of additional public administrations [2–4]. The April 2026 publication by Palantir of a 22-point condensation of CEO Alex Karp and head of corporate affairs Nichol as Zamiska's volume *The Technological Republic* [5] supplies, in the company's own register, the ideological self-description of the configuration: a public articulation of the normative commitments under which the platforms are designed and operated, articulated by scholars of authoritarianism and public administration as constituting a coherent vision of governance grounded in surveillance, militarization, and the rejection of democratic constraint [6–9].

The argument of this paper is that this configuration constitutes the current infrastructural form of the *post-factual polity*—the structural condition, developed in the author's prior work [10], in which the institutional substrates that historically stabilized empirical contestability for purposes of public decision-making have been displaced or hollowed. The post-factual polity framework, as originally specified, identified four such substrates: independent statistical infrastructure, procedurally accountable categorization within public administration, deliberative institutions capable of metabolizing contested evidence, and an information ecology in which empirical claims are subject to corrigibility [10]. The Palantir case requires the framework to specify a fifth substrate, logically prior to the others—the *categorization infrastructure itself*. When categorization migrates into proprietary platforms operating under explicit normative commitments not subject to deliberative scrutiny by the receiving agencies, the post-factual condition is no longer principally a property of discourse, statistics, or institutions. It is a property of the technical and contractual substrate beneath them.

The Palantir case matters not as one example among many but because it is the clearest, best-documented, and ideologically most explicit instantiation of this configuration currently in public view. The April 2026 manifesto did not produce the configuration; the configuration has been operationally in place for years [2–4]. What the manifesto produced was the loss of plausible deniability about the configuration's normative orientation. Public administration scholarship now has both the empirical record and the ideological self-disclosure required to take the problem seriously.

The paper makes three claims. First, that *untangling*—the deliberate, institutionally specified separation of proprietary categorization infrastructure from the categorization functions of the state—is now a first-order public administration problem, not an adjacent one. Second, that existing governance instruments (administrative law, procurement re-

form, and algorithmic impact assessment) are individually and collectively insufficient for the problem because none addresses the substitutive-capacity dimension on which entanglement ultimately rests. Third, that *hybrid intelligence* [1] supplies the constructive architecture that should replace Palantir-class infrastructure once untangling is institutionally underway: a configuration in which human judgment and machine processing are structured as complementary forces under specific institutional conditions, rather than the current arrangement in which categorization has been substantively delegated to proprietary platforms with stated normative commitments hostile to public oversight.

The argument proceeds in six steps. Section 2 specifies the materials and methods of the analysis, including the case selection criteria, the documentary evidence base, and the operational definitions of the diagnostic instruments—making the conceptual analysis replicable by other scholars applying the same instruments to other vendors and jurisdictions. Section 3 develops the infrastructural extension of the post-factual polity framework. Section 4 examines the Palantir case across five layers of ethics and AI governance. Section 5 specifies why existing governance instruments are insufficient. Section 6 develops the untangling problem as a specific institutional task. Section 7 sketches hybrid intelligence as the post-untangling architecture. Section 8 concludes.

The contribution is constructive in its conclusion but uncompromising in its diagnosis. The Special Issue this paper joins calls for research that operationalizes ethical principles in AI systems [11]. Operationalization, in the present moment, requires first naming what is in place, second specifying its incompatibility with the ethical principles in question, and third charting the institutional path to a different configuration. The paper attempts all three.

2. Materials and Methods

No chemicals, reagents, devices, instruments, commercial cell lines, samples, or laboratory materials were used in this study. No separate analytical software was used for data collection or analysis; the named Palantir platforms are treated as objects of documentary analysis rather than as research instruments.

2.1. Research Design

The paper employs a theoretically grounded structured documentary case analysis design, integrating conceptual framework development with documentary case examination and systems-level boundary critique. The design is appropriate to the research question—whether the post-factual polity framework adequately specifies the infrastructural conditions of contemporary AI governance—because the question requires theoretical extension, empirical anchoring, and specification of the feedback and boundary conditions through which administrative systems remain contestable. The design follows the conventions of conceptual public administration scholarship that uses a focal case to test, extend, and operationalize a theoretical framework [12,13].

The research gap was identified through an integrative literature review and structured theoretical mapping rather than through an exhaustive bibliometric study. The review proceeded across six intersecting bodies of work: post-truth and post-factual governance; automated administration and algorithmic accountability; public procurement and vendor accountability; surveillance studies and proprietary platforms; systems theory and socio-technical boundary critique; and the peer-reviewed scholarship on Palantir. The gap criterion was whether a body of literature specified how administrative categories themselves—not only data, models, outputs, or decisions—migrate into proprietary infrastructure and how that migration changes the state's own epistemic capacity. This procedure follows integrative-review conventions for theory-building work, where the objective is

to synthesize dispersed conceptual traditions and identify an undertheorized relationship among them [14,15].

Three components define the design. First, *framework extension*: the paper develops the infrastructural dimension of the post-factual polity framework [10] by specifying a fifth substrate (the categorization infrastructure) and three diagnostic markers (categorical opacity, contestation displacement, substitutive dependency) operationalized as analytical instruments applicable to AI vendor relationships in public administration. Second, *case examination*: the paper applies the extended framework to the focal case of Palantir Technologies' operational footprint in United States and allied public administrations, examining five layers of the case—operational footprint, contractual configuration, ideological self-disclosure, diagnostic profile, and public administration response. Third, *governance specification*: the paper translates the analytical findings into a specified institutional task (untangling) with four constitutive moves and a constructive architecture (hybrid intelligence) for the post-untangling configuration.

2.2. Case Selection

The Palantir case was selected against three pre-specified criteria. The first is *operational scale*: the case must involve a vendor whose platforms are integrated into multiple categorization-intensive functions of the state across multiple jurisdictions, sufficient to permit examination of the infrastructural dimension rather than a single deployment. Palantir's contracts across the United States Department of Defense, Federal Bureau of Investigation, Internal Revenue Service, Immigration and Customs Enforcement, the United Kingdom National Health Service, and police services in multiple German federal states satisfy this criterion [2–4,16,17]. The second is *documentation availability*: the case must possess a documentary record sufficient to support analysis without requiring proprietary disclosure or empirical access that would not be replicable by other scholars. The combination of peer-reviewed scholarly literature on Palantir [16–20], investigative journalism in mainstream outlets, public corporate communications, and government contracting records satisfies this criterion. The third is *ideological articulation*: the case must include a publicly articulated normative orientation by the vendor itself, sufficient to permit examination of the third layer (ideological self-disclosure) without speculative attribution. The April 2026 publication of the Karp–Zamiska 22-point thread, drawn from a published Crown volume [5], satisfies this criterion uniquely among current AI vendors in public administration.

The case is treated as illustrative of a class, not as singular. The diagnostic instruments developed in Section 3 are explicitly portable to other vendors and configurations. Section 7 identifies Anduril, Clearview AI, ShotSpotter, and Thomson Reuters CLEAR as candidates for parallel analysis using the same instruments.

2.3. Documentary Evidence Base

The paper draws on five categories of documentary evidence, each with explicit selection criteria.

Peer-reviewed scholarly literature on Palantir: Brayne's American Sociological Review study of Los Angeles Police Department use of Palantir's Gotham platform [16] and her subsequent monograph [18]; Iliadis and Acker's computational topic modeling of Palantir's surveillance patents [19]; Ulbricht and Egbert's analysis of Palantir regulation in public security [17]. These constitute the core peer-reviewed scholarly base on Palantir as of the manuscript date.

Peer-reviewed scholarly literature on AI governance and the automated administrative state: Engstrom et al.'s Administrative Conference of the United States report [21]; Citron and Calo on the automated administrative state [22]; Coglianese and Lampmann on contracting

for algorithmic accountability [23]; Brauneis and Goodman on algorithmic transparency for the smart city [24]; Jobin, Ienca and Vayena on the global landscape of AI ethics guidelines [25]; Mittelstadt et al. on the ethics of algorithms [26]; Floridi and Cowls on principles for AI in society [27]; Margetts and Dorobantu on rethinking government with AI [28].

Peer-reviewed scholarly literature on comparative public-sector AI governance: Dreyling et al. on Estonia's Bürokratt [29]; the Algorithmic State Architecture comparative framework [30]; the author's own comparative evidence base [1].

Mainstream investigative and news reporting on the focal case (Fortune, TechCrunch, Al Jazeera, Tech Policy Press, Common Dreams, MS NOW, Gazetteer SF), used for empirical claims about the April 2026 thread, the operational footprint, and the public reception. Reporting was selected against two criteria: editorial accountability (mainstream outlets with named authors and standing correction policies) and corroboration (claims used in the analysis appear in two or more independent outlets).

The documentary evidence was also screened for source-position bias. Investigative journalism was used only for claims concerning documented events, contracts, public statements, parliamentary debates, and public controversy, and not as a substitute for legal or technical proof of internal platform operation. Where a claim rested on reporting rather than a public record, it was included only when corroborated by at least one additional independent source or by a primary document. Corporate communications were treated as evidence of vendor self-description, not as neutral evidence of performance. Scholarly commentary was used to map the interpretive controversy surrounding the case, not to establish empirical facts. This hierarchy of evidentiary use reduces, without eliminating, the bias risks associated with a documentary record shaped by litigation risk, corporate strategy, political contestation, and journalistic selection.

Scholarly commentary by named scholars on the focal case, including Mudde, Moynihan, Coeckelbergh, Karpf, and Ganz. Commentary is treated as scholarly opinion in an open conversation, not as empirical evidence. The paper cites the commentary to position itself within the scholarly conversation, not to substitute commentary for analysis.

2.4. Operationalization of Diagnostic Instruments

The three diagnostic markers are operationalized as follows. *Categorical opacity* is present when one or more of the following conditions obtain: (i) the categories used by an agency to act on its population are produced by systems whose internal logic is protected by trade secret; (ii) model weights are not retained by the agency; (iii) recalibration over the life of the deployment is documented to the agency only at vendor discretion. *Contestation displacement* is present when one or more of the following obtain: (i) practical avenues for correcting an erroneous classification require engagement with vendor processes rather than agency processes; (ii) vendor-imposed terms of service condition access to the relevant evidentiary record; (iii) agency-internal contestation procedures operate on categorical outputs rather than on the categorization process. *Substitutive dependency* is present when the agency lacks the in-house technical and analytic capacity to perform the categorization function independently of the vendor, evidenced by the absence of trained personnel, documentation, or infrastructure sufficient to permit substitutive performance within a defined transition horizon (twelve to twenty-four months).

These operational definitions permit replication. A scholar examining a different vendor or jurisdiction can apply the same definitions to documentary evidence in that case and produce a comparable diagnostic profile. Replicability for conceptual work of this kind consists in the operationalizability of the framework's instruments, not in the re-running of a unique focal case.

2.5. Limitations

The design has three principal limitations. First, the analysis depends on documentary evidence and does not include direct agency interviews, internal model documentation, or proprietary disclosure; the diagnostic profile is therefore a documented-state assessment rather than an internal-state assessment. This limitation is also analytically meaningful. The fact that researchers must reconstruct the configuration from public contracts, parliamentary proceedings, investigative reporting, corporate self-description, and peer-reviewed studies is an administrative symptom of the very opacity diagnosed in this article. The hiddenness of model documentation, recalibration practice, training-data lineage, and agency-vendor deliberation is not merely a barrier to research; it is part of the institutional condition under examination. Second, the focal case examines United States and allied public administrations; generalization to other regulatory environments, including the European Union under the AI Act, China, and Gulf states, requires parallel analysis under those jurisdictions' specific institutional conditions. Third, the paper develops a conceptual and design-oriented contribution; empirical validation of the four-move untangling task and the hybrid intelligence post-untangling architecture across multiple agencies and jurisdictions is a research agenda the paper opens but does not itself complete.

3. Results: The Post-Factual Polity, Infrastructural Extension

3.1. The Original Framework

The author's prior work in *The Post-Factual Polity* [10] developed a framework for analyzing the structural conditions under which empirical contestability—the capacity of a polity to distinguish accurate from inaccurate claims for purposes of public decision-making—is preserved or eroded. The framework distinguishes the post-factual polity from the more familiar concept of “post-truth” politics. Post-truth, as developed in the existing literature, centers on the prevalence of falsehood in public discourse and the affective conditions under which falsehood becomes resistant to correction [31]. The post-factual polity framework, in contrast, is structural rather than rhetorical. It asks not whether actors lie but whether the institutional conditions under which lies can be reliably identified and corrected remain operative.

The original formulation identified four institutional substrates whose integrity is necessary for a polity to function on a factual basis: (i) independent statistical infrastructure, capable of producing population-level data not controlled by partisan actors; (ii) procedurally accountable categorization within public administration, in which the categories used by the state to act on its population are produced under administrative-law constraints; (iii) deliberative institutions capable of metabolizing contested evidence into binding decisions; and (iv) an information ecology in which empirical claims are subject to corrigibility through accessible, trusted processes of verification and correction [10]. Erosion of any single substrate does not by itself produce a post-factual polity. Their simultaneous degradation does. The 2024 volume traced the simultaneous degradation across the United States and several other jurisdictions and argued that the resulting condition is properly characterized as post-factual rather than merely polarized.

3.2. The Missing Substrate

The Palantir case requires the framework to specify a fifth substrate, logically prior to the other four: the *categorization infrastructure itself*. The four substrates of the original framework all presuppose a settled answer to a prior question: where is categorization performed, by whom, and under what constraints. Statistics aggregate over categories. Administrative decision-making applies categories. Deliberation contests categories. Verification corrects errors in categories. None of these operations can be performed coherently if the categories

themselves are produced inside infrastructure that the polity cannot reconstruct, audit, or modify.

The fifth substrate emerges path-dependently from the degradation of the original four. When independent statistical infrastructure is distrusted, when agency categorization is weakened by outsourcing and downsizing, when deliberative institutions struggle to process technical evidence, and when the public information ecology loses corrigibility, proprietary categorization platforms become attractive precisely because they promise speed, integration, and operational certainty. The vacuum created by weakened public substrates becomes a market opportunity for private analytical infrastructure. The fifth substrate therefore does not replace the original four; it explains how their erosion makes vendor-mediated classification appear administratively necessary and politically convenient.

For most of the history of the modern administrative state, the categorization function was performed substantially inside public agencies under administrative-law constraints. Categorical decisions—what counts as fraud, what triggers enforcement priority, what constitutes eligibility, what defines risk—were produced by agency personnel applying agency-developed criteria, subject to rule-making procedures, judicial review, and freedom-of-information disclosure [21–23]. The categories were contestable in principle and frequently contested in practice. The administrative-law apparatus was developed precisely to make this contestation procedurally tractable.

The migration of categorization into proprietary analytical platforms changes the configuration in three structurally significant ways. First, the construction process is no longer subject to the procedural constraints that govern agency action; it is governed instead by vendor contracts and trade-secret protections [21,23]. Second, the categories produced are presented to receiving agencies as *inputs* rather than as outputs of a contestable process; their constructed character is occluded by the technical interface and the receiving agency's limited capacity to interrogate it [22]. Third, the conditions under which a category could be challenged, audited, or recalibrated are themselves controlled by the vendor, frequently through terms-of-service mechanisms that lie outside the administrative-law domain [21,23].

The result is that the categorization layer of public decision-making moves outside the institutional domain in which the post-factual polity framework's original four substrates operate. The locus of empirical contestability shifts from the public sphere to the contractual and proprietary sphere. The public sphere often retains no procedurally adequate means of recovering it. This is the *infrastructural form of the post-factual polity*: a configuration in which the prior question—where categorization is performed and under what constraints—has been answered, in increasing portions of the administrative state, in a manner that the existing post-factual polity framework's substrates cannot govern.

3.3. Why the Extension Matters

The extension is not a refinement. It changes what the framework is for. The original framework was diagnostic of conditions that public administration scholarship and democratic theory had the conceptual resources to address, even where the political conditions were unfavorable. The infrastructural extension identifies a class of problems that cannot be addressed within the existing conceptual resources, because the operative substrate sits outside the institutional domain those resources were developed to govern. The extension converges with related theoretical developments in surveillance studies—Zuboff's specification of *instrumentarian power* as a distinct form of power that “knows and shapes human behavior toward others' ends” through ubiquitous computational infrastructure [32], Crawford's analysis of AI as an *infrastructural* and *extractive* configuration rather than a discrete technology [33], and Pasquale's earlier specification of the “black box society” in

which categorical determinations are produced inside opaque proprietary systems [34]. The contribution of the present paper is to specify how this configuration interacts with the post-factual polity framework's existing substrates, and to identify the institutional task—untangling—that the convergence requires.

This distinction also separates AI epistemic capture from ordinary legacy IT vendor lock-in. Legacy IT dependency typically concerns technical migration costs, data-format incompatibilities, licensing terms, and organizational inertia. Those problems are serious, but they usually leave the agency's substantive categories intact. AI infrastructure dependency reaches a different layer. It concerns the outsourcing of epistemic classification: the administrative capacity to define what counts as a risk, a threat, a fraud signal, an eligibility condition, a compliance anomaly, or an enforcement priority. The state is not merely renting software plumbing; it is importing an organized way of seeing the population. That qualitative difference is why substitutive dependency becomes a constitutional and administrative problem rather than only a procurement problem.

The Palantir case, examined in the next section, makes the abstract argument concrete.

3.4. Contribution to Systems Theory

The study contributes to systems theory by treating AI governance as a socio-technical boundary problem rather than as compliance oversight over discrete tools. Systems theory has long emphasized that complex social systems are constituted through boundaries, feedback loops, classification schemes, control mechanisms, and observer-dependent definitions of the problem situation [35–38]. The Palantir-class case shows that AI governance failures arise when the boundary around a public decision system is redrawn through contracting so that the transformative function—the conversion of raw data into actionable administrative categories—is no longer located inside the accountable public system.

The three diagnostic markers operate as systems-theoretic instruments. Categorical opacity identifies hidden transformations inside the system. Contestation displacement identifies the relocation or weakening of corrective feedback. Substitutive dependency identifies the loss of redundancy, resilience, and internal capacity needed for viable system performance. Untangling is therefore a form of boundary redesign: the reconstruction of public control over the categories, feedback loops, and exit capacities that make an administrative system governable. This is the article's central systems-theory contribution.

4. The Palantir Case: Multiple Layers of Ethics and AI Governance

The Palantir case is significant for the present argument across at least five distinguishable layers: the operational footprint, the contractual configuration, the ideological self-disclosure, the diagnostic profile, and the public administration response (or non-response). Each layer raises a distinct set of ethics and AI governance questions, and the layers compound. The case is not a hypothetical illustration of the post-factual polity in code; it is the configuration's current operative form in the jurisdictions where it is most fully developed.

4.1. Operational Footprint

Palantir Technologies was founded in 2003 with investment from In-Q-Tel, the venture capital arm of the United States Central Intelligence Agency, and built its early business on post-9/11 intelligence integration work [4,16]. Sarah Brayne's ethnographic study of the Los Angeles Police Department's adoption of Palantir's Gotham platform, published in *American Sociological Review* in 2017 and elaborated in her 2020 monograph *Predict and Surveil*, established the foundational scholarly account of how the company's analytical infrastructure transforms police categorization practices [16,18]. Iliadis and

Acker's 2022 computational topic modeling of Palantir's surveillance patents extended the analysis to the company's broader sociotechnical aspirations [19]. Ulbricht and Egbert's 2024 analysis of Palantir regulation in public security examined the structural difficulty of governing the company's integration into European policing infrastructures [17]. Its current operational footprint includes contracts with the United States Department of Defense (including the Maven Smart System used in targeting infrastructure), the Federal Bureau of Investigation, Immigration and Customs Enforcement (under which the company holds a 2025 no-bid contract for ImmigrationOS, an AI platform used to identify noncitizens and track deportations), the Internal Revenue Service, multiple components of the intelligence community, the United Kingdom National Health Service (under the £330 million Federated Data Platform contract awarded in November 2023 to a Palantir-led consortium), and the police services of multiple German federal states including Hesse, North Rhine-Westphalia, and Bavaria [2–4,16–18,39]. Additional contracts span healthcare systems, financial regulators, and municipal governments across multiple jurisdictions.

The United Kingdom NHS Federated Data Platform (FDP) deployment merits particular attention as a case in which the contractual configuration and the political contestation are both unusually well-documented in the primary parliamentary record. The £330 million seven-year contract, formally part of the UK Government Major Projects Portfolio, has been the subject of a formal Westminster Hall debate on 16 April 2026, in which parliamentarians from at least seven political parties spoke against the contract and called for the activation of its 2027 break clause [39,40]. The debate's lead speaker, Liberal Democrat MP Martin Wrigley, presented evidence that the FDP "is awful to use," benefits only a quarter of its user organizations, and "leaves no deliverables after the subscription—no software, no improvements and no intellectual property after spending more than £330 million" [39,40]. The Department of Health and Social Care junior minister Zubir Ahmed responded that the government would evaluate alternative providers at the 2027 break clause point [39]. The debate constitutes formal parliamentary recognition of substitutive dependency as a structural feature of the deployment rather than as a contingent failure.

The footprint matters analytically because of its distribution across categorization-intensive functions of the state. Palantir is not principally a back-office productivity vendor. Its platforms are integrated into precisely the functions in which categorization, risk-assessment, and enforcement-targeting operate as the state's coercive substrate: defense and intelligence operations, immigration enforcement, tax and benefits administration, policing, healthcare resource allocation. The categorization layer of these functions has been substantially relocated into a single vendor's analytical platforms across multiple jurisdictions.

The cross-jurisdictional footprint does not imply that the United States, the United Kingdom, and Germany share identical administrative-law structures or contestation pathways. They do not. A United States immigration classification within the executive branch, a defense-intelligence targeting support system, a United Kingdom NHS resource-allocation platform, and a German federal-state policing system sit within different constitutional, statutory, procurement, data-protection, and judicial-review environments. The case is therefore not used to make a single doctrinal claim about the state as a monolith. It is used to test whether the same systems-level markers—categorical opacity, contestation displacement, and substitutive dependency—appear across jurisdictions despite different legal cultures. The answer is diagnostic, not doctrinal: the markers travel, while the available legal remedies vary significantly.

4.2. Contractual Configuration

The contractual configuration compounds the operational concern. A non-trivial portion of Palantir's federal contracting is conducted under sole-source or no-bid ar-

rangements [3]. Where competitive procurement has occurred, the company's incumbent operational integration frequently structures the competition in ways that limit substitutive capacity in receiving agencies. Once Palantir's platforms are integrated into an agency's data infrastructure, exit from the relationship requires not merely the substitution of a different vendor but the reconstruction of the agency's analytical capacity, frequently from a state of substantial atrophy [4].

The company's frequent characterization of itself as a data processor, and of public agencies as data controllers or owners, is legally important but analytically insufficient. The processor label may allocate formal responsibilities under a data-protection regime, but it does not settle whether the agency retains epistemic control over category production, auditability, contestation, or exit. Cross-border data governance also complicates the processor framing. The United States CLOUD Act clarified the reach of certain U.S. legal process to data in the possession, custody, or control of covered service providers regardless of storage location [41]. This does not mean that every dataset processed by a U.S.-connected provider is automatically disclosed to U.S. authorities. It does mean that ownership language alone is an inadequate governance safeguard when a critical public infrastructure relationship depends on a vendor embedded within a different sovereign legal order.

This configuration has been articulated with particular clarity by former United Kingdom prisons minister Rory Stewart, drawing on his direct ministerial experience of Palantir's integration into UK government services. In the 30 April 2026 episode of *The Rest Is Politics* devoted to the company's expansion in British government [42], Stewart characterized Palantir's market entry strategy as the "plumbing" pitch—the company offers to integrate fragmented, incompatible government databases at minimal initial cost, presenting itself as the connective infrastructure that allows disparate departmental systems to interoperate. The pitch is operationally compelling for ministers facing the chronic problem of database fragmentation. The structural consequence Stewart identifies, however, is not principally a privacy concern. It is what he terms the *dependency trap*: once the integration is in place, the cost and technical difficulty of "rooting them out" become prohibitive, and the company acquires significant leverage over government operations [42]. The configuration shifts effective control over backend infrastructure from public officials to the private firm managing it.

Stewart's account is significant for the present analysis for three reasons. First, it provides a credentialed, named, government-internal articulation of substitutive dependency as an operational reality rather than a theoretical risk—the framing emerges from ministerial experience rather than from academic critique. Second, it specifies the mechanism by which entanglement deepens: the convenience of low-cost initial trials produces long-term arrangements in which providers cannot be switched. Third, it identifies the consequence as a shift in the locus of operational control, which is precisely the structural inversion the post-factual polity in code framework specifies. Stewart's "plumbing" and "dependency trap" framings map almost exactly onto the substitutive dependency marker developed in Section 3 of this paper, and supply an in-government counterpart to the academic specification.

The vendor's own defense of this configuration warrants direct analytical engagement, since reviewers of the present framework would reasonably ask whether the diagnostic markers survive the company's stated position. The clearest extant articulation of that position is the oral evidence of Louis Mosley, Executive Vice-President of Palantir Technologies, before the United Kingdom House of Commons Science, Innovation and Technology Committee on 8 July 2025 [43]. In that testimony, Mosley defended Palantir's role in the NHS by reframing the company's function as patient-care optimization: "the most important thing,

when serving the NHS, is to ensure that we help the NHS to treat its patients faster, more effectively and ultimately more efficiently” [43]. This framing—that Palantir functions as a neutral efficiency tool rather than as ideologically inflected infrastructure—is the central rhetorical move in the company’s public defense and corresponds to the “plumbing” pitch Stewart identifies from the procurement side [42].

The framing does not, however, address the three diagnostic markers developed in Section 3, and its inadequacy on each is analytically clarifying. On *categorical opacity*: the question is not whether the categorizations Palantir’s systems produce are intended to optimize patient care, but whether the receiving agency can independently reconstruct, audit, and modify the categorization process. The Mosley testimony does not address this question, and the same parliamentary committee record documents Mosley’s failure to directly answer when asked by Kit Malthouse MP what protections would be put in place to ensure UK programmes are not complicit in alleged war crimes through the company’s interoperable civil-military product architecture [43,44]. On *contestation displacement*: the question is whether NHS patients and clinicians subject to FDP-mediated resource allocation possess procedurally adequate, agency-controlled avenues for challenging categorizations. The patient-care framing does not address contestation accessibility at all. On *substitutive dependency*: Wrigley’s parliamentary statement that the FDP “leaves no deliverables after the subscription—no software, no improvements and no intellectual property” [39,40] is a direct empirical refutation of the framing that the company is providing a neutral efficiency tool. A neutral efficiency tool would leave the receiving institution capable of continuing the function after the relationship ends. The FDP configuration, by the contracting authority’s own admission in parliamentary record, does not.

The analytical point is not that the Mosley framing is offered in bad faith. It is that the framing operates on a category of analysis—the intended optimization function of the deployed tool—that is logically distinct from the category of analysis on which the post-factual polity in code framework operates: the institutional conditions under which the categorization function is performed and remains contestable. Vendor defenses pitched in the first register do not address questions raised in the second.

The efficiency framing also deserves a sharper public administration critique. Optimization is never a neutral systems objective in public administration. A platform may improve localized operational velocity while displacing procedural due process, professional discretion, equity review, transparency, and public contestability. In healthcare, patient-care optimization sounds benign; in practice, the interface may privilege throughput, scheduling efficiency, waiting-list movement, or resource allocation metrics over values that remain harder to encode. The governance problem is not that efficiency lacks value. The problem is that efficiency becomes ideological camouflage when the system accelerates local action while moving the reasons for action outside the accessible institutional record. In that configuration, public officials trade systemic accountability for localized operational speed.

The contractual configuration is also governed by trade-secret protections and proprietary data licensing arrangements that limit the receiving agency’s ability to retain model documentation, audit recalibration over the life of the deployment, or transfer learned configurations to a successor system [21,23]. The receiving agency, in many cases, does not possess the model weights, the training data lineage, or the technical documentation required to reconstitute the analytical function independently. This is not a secondary contractual detail; it is the structural feature that converts a procurement relationship into an infrastructural one.

4.3. Ideological Self-Disclosure: The April 2026 Manifesto

On 19 April 2026, the official corporate account of Palantir Technologies posted a 22-point thread on the X platform presenting itself as a condensed version of the 2025 Crown volume *The Technological Republic: Hard Power, Soft Belief, and the Future of the West*, co-authored by Karp and Zamiska [5,45,46]. Within forty-eight hours, the thread had registered more than thirty million views and provoked sustained commentary from scholars of authoritarianism, public administration, and technology ethics [6–9].

The thread is significant for the present argument not because it endorses any particular political position—corporate political speech is not, by itself, an AI governance problem—but because it constitutes a public articulation, in the vendor’s own register, of the normative commitments under which the platforms identified in Section 3.1 are designed and operated. The document explicitly rejects what it characterizes as “vacant and hollow pluralism,” asserts that “certain cultures and indeed subcultures” have proved “middling, and worse, regressive and harmful,” advocates the reinstatement of mandatory national service, calls for the reversal of postwar German and Japanese demilitarization, and frames Silicon Valley’s relationship to the United States defense and intelligence apparatus as a matter of moral debt requiring expanded private-sector integration into state coercive functions [45–47].

Scholars of authoritarianism and public administration have read the document as constituting a coherent ideological program. Cas Mudde, one of the leading scholars of the contemporary radical right, characterized the document as “Technofascism pure” and called for European institutions to divest from the company [6,7]. The public administration scholar Donald Moynihan published an analysis concluding that the manifesto’s vision is one in which the United States government and its private technology partners operate “as dominant players, unconstrained by accountability” [8]. The philosopher Mark Coeckelbergh framed the document as a case study in technofascism understood as the fusion of expansive surveillance, predictive analytics, and state power [9]. The day after the thread was posted, the journalist Gil Durán was permanently suspended from the X platform after reposting it with the caption “TLDR: Fascism” [48].

For the present argument, the relevant point is not whether “technofascism” is the correct scholarly label for the ideological program articulated in the document. It is that the categorization layer of multiple public administrations now incorporates infrastructure produced under publicly stated normative commitments that the receiving agencies did not deliberately scrutinize before adopting. The ethics and AI governance question is not whether any individual receiving agency endorses the commitments. It is that the categorization layer of state action has been delegated, by procurement decision, to infrastructure operating under stated commitments that are incompatible with the procedural constraints of liberal-democratic public administration.

4.4. Diagnostic Profile

Applied to the Palantir case, the diagnostic profile is concise. Categorical opacity appears where risk scores, enforcement priorities, eligibility classifications, targeting supports, or resource-allocation indicators are generated through vendor systems whose internal logic, recalibration, and documentation remain only partly available to the receiving agency [3,16–19,21,23]. Contestation displacement appears where affected individuals, clinicians, officers, or agency personnel may challenge an outcome but lack a procedurally adequate path to challenge the categorization process that produced it [3,4,22]. Substitutive dependency appears where the agency cannot reproduce the categorization function within a reasonable transition period without the vendor’s infrastructure, personnel, documentation, or technical cooperation [4,39,40,42].

The Palantir case is especially significant because the three markers reinforce one another. Opacity makes error difficult to locate. Displaced contestation makes error difficult to correct. Dependency makes exit difficult even after error, controversy, or normative incompatibility is recognized. This compound profile is what distinguishes the case from ordinary technology outsourcing. It describes a public decision system in which the state retains formal authority while losing parts of the epistemic, procedural, and technical capacity that make formal authority publicly governable.

4.5. Public Administration Response

The fifth layer is the response—or, more accurately, the structural conditions that have made adequate public administration response difficult. The administrative-law apparatus developed to govern agency decision-making does not map cleanly onto vendor-mediated categorization; the procurement frameworks under which the relationships were entered did not anticipate the substitutive-capacity problem; the algorithmic accountability literature, while increasingly sophisticated, has been oriented to transparency and impact assessment rather than to disentanglement [21–26]. The result is that even agencies and oversight bodies inclined to address the configuration have lacked the institutional vocabulary and the operational template to do so. Section 5 examines why.

5. The Insufficiency of Existing Governance Instruments

5.1. Administrative Law

Administrative law governs the processes by which public agencies make decisions affecting rights and interests. Its core mechanisms—notice-and-comment rule-making, reasoned decision-making, judicial review of arbitrary action, freedom-of-information disclosure—were developed under the assumption that agency decision-making is performed by agency personnel using agency-controlled procedures [21,23,24]. When decision-making is partly performed by external categorization infrastructure, several of these mechanisms are circumvented without being formally repealed. Vendor-produced categorizations may not constitute “agency action” in the doctrinal sense, may fall outside standard administrative-record requirements, and are frequently shielded from disclosure under trade-secret exceptions to freedom-of-information regimes [21,23]. The administrative-law apparatus operates on a unit of analysis (agency action) that does not map cleanly onto the locus of categorization (the vendor-controlled platform).

5.2. Procurement Reform

Procurement reform has been proposed as an alternative regulatory site. The argument is that public agencies can use procurement contracts to require algorithmic transparency, model documentation, bias auditing, and continuing disclosure as conditions of vendor selection [23]. The proposal is constructive but, applied to the Palantir-class case, faces three structural limits.

First, procurement leverage presupposes a competitive market. In jurisdictions where a small number of vendors hold dominant positions in defense, intelligence, and enforcement analytics, the conditions for genuine procurement leverage are attenuated [4,49]. Sole-source contracting, incumbent operational integration, and the structural dynamics of vendor consolidation all reduce the agency’s *ex ante* negotiating position.

Second, procurement requirements operate at the moment of contract formation; they have limited purchase on the *ongoing* recalibration of model behavior over the life of a deployment, which is where many of the most consequential decisions are made [25,26].

Third, and most fundamentally, procurement frameworks address transparency of the vendor system but do not by themselves restore the agency’s *substitutive capacity*. An

agency that requires its vendor to disclose model documentation has acquired a measure of categorical visibility, but it has not necessarily acquired the in-house capacity to perform the categorization function independently. Substitutive dependency, the most consequential marker of the post-factual polity in code, is largely outside the reach of procurement instruments operating on incumbent vendors.

5.3. Algorithmic Impact Assessment

Algorithmic impact assessment (AIA) regimes, modeled in part on environmental impact assessment, have emerged as a third governance approach [24,50]. AIAs require agencies to document and publicly justify the deployment of algorithmic systems, with provision for public comment and continuing review. AIAs are a meaningful contribution but their effectiveness depends on prior conditions—categorical visibility, agency interpretive capacity, substantive avenues for contestation—that are precisely what the post-factual polity in code erodes. An AIA regime applied to a substantively entangled infrastructure documents the entanglement rather than reverses it.

5.4. The Joint Insufficiency

The implication is not that administrative law, procurement, and impact assessment are without value. They are necessary components of any serious governance response. The implication is that they are *individually and collectively insufficient* for the problem identified in Section 4. None addresses the substitutive-capacity dimension on which entanglement ultimately rests. None operates on the institutional locus where the categorization function has been relocated. None engages the question of how to reverse the configuration once it is in place. They are instruments developed for a different problem.

The problem the present configuration poses is not the regulation of a deployed AI system. It is the *disentanglement of state categorization functions from a vendor whose ideological self-description is incompatible with the procedural constraints that public administration is institutionally obligated to preserve*. This is a different problem and requires different instruments.

6. Untangling: Specifying the Institutional Task

6.1. The Untangling Problem

Untangling, as the present paper uses the term, refers to the deliberate, institutionally specified separation of proprietary categorization infrastructure from the categorization functions of the state. It is not synonymous with vendor exclusion, contract termination, or technological prohibition. It is the rebuilding of the institutional conditions under which the categorization functions of the state are performed by the state, with whatever AI tools are appropriate, under the procedural constraints that the post-factual polity framework's original four substrates were developed to preserve.

The institutional task is not theoretical. The 16 April 2026 Westminster Hall debate on the NHS Federated Data Platform constitutes the first formal parliamentary deliberation in a major Western democracy on whether to invoke a contractual break clause to begin untangling a Palantir-class deployment [39,40]. The debate convened parliamentarians from at least seven political parties, generated direct ministerial commitment to evaluate alternative providers at the 2027 break clause, and articulated the substitutive dependency problem in the parliamentary record [39,40]. The Welsh NHS decision to build its own analytical platform rather than adopt the FDP [51] supplies a parallel jurisdictional case in which substitutive capacity is being preserved rather than recovered. These are instances of untangling at varying stages—formal contestation in England, parallel construction in Wales—and they indicate that the institutional task is being attempted, however incompletely, by actors operating without the analytical vocabulary the present

paper supplies. The framework developed here is intended to support, clarify, and extend that emerging practice.

Untangling is institutionally demanding. It requires four sustained moves.

6.2. The Four Moves

Capacity reconstruction: Public agencies must rebuild the in-house technical and analytic capacity to perform the categorization functions on which their authority depends, independent of vendor relationships. This requires sustained investment in personnel, training, infrastructure, and institutional knowledge across a time horizon measured in years rather than budget cycles. It requires that public-sector technical roles be compensated and protected at levels that support recruitment and retention against private-sector alternatives. It requires statutory protection against the kind of mass dismissals that characterized the 2025 DOGE initiative [1,52–54]. If full capacity reconstruction is politically impossible in the short term, agencies should at minimum retain non-negotiable sovereign classification functions: ownership of category definitions; control over feature selection and validation criteria; access to training-data lineage and decision logs; internal capacity to reproduce and audit sampled outputs; authority to suspend categories that fail audit; and an exit plan that preserves operational continuity. Anything less leaves the agency formally in charge but substantively dependent.

Categorical repatriation: The categorization functions currently performed by vendor platforms must be rebuilt inside public agencies under administrative-law constraints. This does not require that agencies stop using AI tools, cloud infrastructure, or vendor-provided components. The boundary is legal and technical at the same time: the agency must control the policy taxonomy, feature definitions, validation thresholds, fine-tuning or configuration choices that shape public categories, decision logs, appealable explanations, and audit/retraining authority. The state need not own every foundation model or every server used in an AI workflow; it must own and govern the consequential administrative categories on which it acts. The distinction is between using a vendor tool as one component of an agency-controlled categorization process and consuming categorical outputs from a vendor process the agency cannot reconstruct. Repatriation requires the former and rejects the latter.

Contractual restructuring: Existing vendor relationships must be restructured to support, rather than impede, capacity reconstruction and categorical repatriation. This requires the renegotiation or termination of contractual provisions that prevent agency retention of model documentation, training data lineage, and technical knowledge required for substitutive capacity. It requires statutory limits on trade-secret protections that operate against the agency's institutional interest in disentanglement. It requires, in some cases, the structured wind-down of incumbent vendor relationships under conditions that preserve operational continuity during the transition.

Procurement reorientation: Future procurement must be reoriented around the prevention of re-entanglement. This requires structural diversity across vendors, methodological approaches, and ideological orientations; explicit assessment of vendor normative commitments where those commitments have been publicly articulated; and the integration of categorical visibility, contestation accessibility, and substitutive capacity as conditions of award rather than as post-hoc accountability instruments. The goal is not to extract better terms from incumbents but to maintain the conditions under which alternative configurations remain possible across future procurement cycles.

6.3. The Political Economy

Each of the four moves cuts against the political economy that produced the current configuration. Capacity reconstruction requires sustained public investment at a moment when many jurisdictions are dismantling rather than rebuilding public-sector technical capacity [1,52–54]. Categorical repatriation requires agencies to perform functions they have, in many cases, structurally forgotten how to perform. Contractual restructuring requires the unwinding of incumbent vendor relationships in which significant institutional and political-economic interests are invested. Procurement reorientation requires the development of alternative vendor ecosystems that may not currently exist at the required scale.

The paper does not minimize these difficulties. It argues that they are the actual contours of the problem. The post-factual polity in code is not a configuration that can be governed in place by adjusting the existing instruments. It is a configuration that must be untangled, and the untangling must be specified institutionally before it can be undertaken politically.

7. Hybrid Intelligence as the Post-Untangling Architecture

The paper closes with the constructive question: once untangling is institutionally underway, what should replace the Palantir-class infrastructure in the categorization-intensive functions of the state? The answer, developed at length in the author's recent volume *Hybrid Intelligence for Effective Digital Governance* [1], is *hybrid intelligence*—a governance architecture in which human judgment and machine processing are structured as complementary forces under specific institutional conditions, rather than substitutable inputs to be optimized against each other.

Hybrid intelligence is defined by three properties [1]. *Complementarity*: machine processing handles scale, pattern recognition across large datasets, and consistent application of explicit rules; human judgment handles contextual interpretation, ethical assessment, novel situations, and the integration of considerations that resist formalization. *Institutional embedding*: hybrid intelligence is not a property of individual decisions but of the institutional arrangements within which decisions are produced; the conditions that make human judgment substantive rather than ornamental are part of the architecture, not external to it. *Outcome orientation*: the configuration is evaluated by whether it produces effective governance outcomes—accountability, transparency, ethical integrity, service quality, public trust—not by automation metrics.

Hybrid intelligence also requires protection against automation bias. Human oversight does not become meaningful merely because a person appears somewhere in the workflow. Empirical work on human-AI interaction in public-sector decision-making, together with the broader automation-bias literature, shows that decision-makers may over-rely on automated advice, defer to algorithmic outputs under workload pressure, or selectively adopt algorithmic recommendations that confirm prior assumptions [55–57]. In administrative settings, this risk is intensified by caseload pressure, fear of personal accountability, managerial performance metrics, and the rhetorical authority of quantitative scores. A human-in-the-loop design without institutional friction risks becoming a ceremonial rubber stamp.

The post-untangling hybrid architecture therefore needs specific backstops: mandatory reason-giving when officials accept or reject high-stakes algorithmic recommendations; independent second review for adverse classifications; randomized audits of accepted recommendations; adversarial review panels for contested categories; logged dissent from professional staff; confidence displays that reveal uncertainty rather than hide it; workload rules that prevent automated triage from becoming irresistible; and appeal procedures

that reach the categorization process, not merely the final output. These friction points do not slow the system for their own sake. They preserve the feedback loops and human judgment that make hybrid intelligence distinct from automated administration with a human signature at the end.

The case for hybrid intelligence as the post-untangling architecture is grounded in comparative evidence from public sector AI deployments across more than fifty countries [1]. Estonia's e-government infrastructure, including the recent Bürokratt initiative analyzed by Dreyling and colleagues [29], illustrates how hybrid configuration can be operationalized in practice while retaining administrative-law constraints. The Algorithmic State Architecture comparative framework developed by recent scholarship [30] specifies how Estonia, Singapore, India, and the United Kingdom have organized different layers of digital public infrastructure, data integration, algorithmic decision-making, and user-facing services—and how the configurations vary in the institutional conditions they preserve. AI deployments in Sweden, Norway, Denmark, South Korea, and Japan have produced configurations that maintain accountability and ethical safeguards while capturing efficiency and service-quality benefits [1,30,58]. None is a model to be replicated wholesale. They share, however, the structural feature relevant to the present argument: the institutional conditions that make hybrid configuration operative—public-sector technical capacity, contestability, oversight integrated into deployment, and human authority over categorical decisions—are present at significant levels.

Hybrid intelligence supplies the constructive answer to the question that untangling necessarily raises. Disentanglement from Palantir-class infrastructure is not the rejection of AI in public administration. It is the precondition for the responsible integration of AI under institutional conditions that public administration is competent to maintain.

8. Conclusions

The post-factual polity has acquired an infrastructural form. The categorization layer of public decision-making—the institutional apparatus by which risk, threat, eligibility, fraud, and deviance are rendered legible to state action—has migrated, across multiple high-stakes domains in the United States and allied jurisdictions, into proprietary analytical platforms operating under publicly stated normative commitments that are incompatible with the procedural constraints of liberal-democratic public administration. The current operative form of this configuration is the operational footprint of Palantir Technologies. The April 2026 publication by the company of the 22-point manifesto condensing *The Technological Republic* did not produce the configuration; it ended plausible deniability about the configuration's normative orientation.

Public administration scholarship now has both the empirical record and the ideological self-disclosure required to take the problem seriously as a problem of AI governance and ethics. The diagnostic markers—categorical opacity, contestation displacement, substitutive dependency—supply the analytical vocabulary. The existing governance instruments—administrative law, procurement reform, algorithmic impact assessment—supply necessary but individually and collectively insufficient resources. The institutional task is *untangling*: the deliberate, institutionally specified separation of proprietary categorization infrastructure from the categorization functions of the state, accomplished through capacity reconstruction, categorical repatriation, contractual restructuring, and procurement reorientation. The post-untangling architecture is *hybrid intelligence*: a configuration in which human judgment and machine processing are structured as complementary forces under institutional conditions that public administration is competent to maintain.

None of this is achievable in a single budget cycle, by a single agency, or under any of the existing political-economic configurations in which the entanglement was produced.

The paper does not minimize the difficulty. It argues that the difficulty is the actual shape of the problem, and that public administration as a discipline has an obligation to specify the shape with precision before the political conditions for addressing it close further. The Special Issue convened around the ethics and governance of AI systems is the appropriate venue for that specification. The argument advanced here is offered as a contribution to it.

For systems theory, the article specifies AI governance as a problem of boundary control, feedback integrity, and system viability. The central question is not whether a tool is accurate in isolation, but whether the administrative system retains the capacity to define its categories, observe its own transformations, correct errors through accessible feedback, and exit a vendor relationship without losing the function being governed. Those are systems properties, not merely legal or technical features.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: No new data were created or analyzed in this study. Data sharing is not applicable to this article.

Conflicts of Interest: The author serves as Guest Editor of the Special Issue to which this manuscript is submitted. In accordance with MDPI policy, the author recused from all editorial decisions concerning this submission, which are handled independently by the Editorial Office and other members of the editorial board. The author declares no other conflicts of interest.

References

1. Alibašić, H. *Hybrid Intelligence for Effective Digital Governance: AI in Administration*; Public Administration and Information Technology; Springer: Cham, Switzerland, 2026; Volume 42.
2. Wong, P. Palantir's Mini Manifesto Claims Some Cultures Are 'Harmful' and 'Middling'. *Fortune*, 22 April 2026. Available online: <https://fortune.com/2026/04/22/palantir-alex-karp-mini-manifesto-national-security-defense-tech-ai/> (accessed on 30 April 2026).
3. Hayes, C. Reading Palantir: Why the Defense Tech Giant's Manifesto May Signal Panic Inside the Company. *Gazetteer SF*, 23 April 2026. Available online: <https://sf.gazetteer.co/reading-palantir-why-the-defense-tech-giants-manifesto-may-signal-panic-inside-the-company> (accessed on 30 April 2026).
4. Amnesty International. "They Have All the Tools at Their Disposal": Tech Made by Palantir and Babel Street Pose Surveillance Threats to Pro-Palestine Student Protestors and Migrants; Index Number AMR 51/0211/2025; Amnesty International: London, UK, 2025. Available online: <https://www.amnesty.org/en/wp-content/uploads/2025/08/AMR5102112025ENGLISH.pdf> (accessed on 4 May 2026).
5. Karp, A.C.; Zamiska, N.W. *The Technological Republic: Hard Power, Soft Belief, and the Future of the West*; Crown: New York, NY, USA, 2025.
6. Reed, B. One of The Scariest Things I Have Seen': Alarms Sound Over 'Technofascist' Palantir Manifesto. *Common Dreams*, 21 April 2026. Available online: <https://www.commondreams.org/news/palantir-technofascist-manifesto-criticism> (accessed on 30 April 2026).
7. Mudde, C. Technofascism pure! Clearly, the manifesto disqualifies Palantir as a business partner for any other country than the U.S. LinkedIn, 20 April 2026. Available online: https://www.linkedin.com/posts/cas-mudde-543a642a5_because-we-get-asked-a-lot-by-palantirtech-activity-7451991415800307712-AoHg (accessed on 30 April 2026).
8. Moynihan, D. Power Without Accountability: The Palantir Manifesto. *Can We Still Govern? (Substack)*, 21 April 2026. Available online: <https://donmoynihan.substack.com/p/palantir-wants-power-without-accountability> (accessed on 30 April 2026).
9. Coeckelbergh, M. Palantir's Manifesto: Technofascism in Plain Sight. *Medium*, 21 April 2026. Available online: <https://coeckelbergh.medium.com/palantirs-manifesto-technofascism-in-plain-sight-c160ca377e9a> (accessed on 30 April 2026).
10. Alibašić, H. *The Post-Factual Polity: Ethical, Governance, Administrative, and Policy Crises in the Disinformation Era*; Information Age Publishing/Emerald Publishing Limited: Charlotte, NC, USA, 2024.
11. Alibašić, H. (Ed.) *Ethics and Governance of Artificial Intelligence (AI) Systems*. Special Issue, *Systems*. MDPI: Basel, Switzerland, 2025–2026; Available online: https://www.mdpi.com/journal/systems/special_issues/DK47159FI6 (accessed on 30 April 2026).
12. Yin, R.K. *Case Study Research and Applications: Design and Methods*, 6th ed.; SAGE Publications: Thousand Oaks, CA, USA, 2018.

13. Eisenhardt, K.M. Building theories from case study research. *Acad. Manag. Rev.* **1989**, *14*, 532–550. [CrossRef]
14. Torraco, R.J. Writing integrative literature reviews: Guidelines and examples. *Hum. Resour. Dev. Rev.* **2005**, *4*, 356–367. [CrossRef]
15. Webster, J.; Watson, R.T. Analyzing the past to prepare for the future: Writing a literature review. *MIS Q.* **2002**, *26*, xiii–xxiii. [CrossRef]
16. Brayne, S. Big data surveillance: The case of policing. *Am. Sociol. Rev.* **2017**, *82*, 977–1008. [CrossRef] [PubMed]
17. Ulbricht, L.; Egbert, S. In Palantir we trust? Regulation of data analysis platforms in public security. *Big Data Soc.* **2024**, *11*, 1–14. [CrossRef]
18. Brayne, S. *Predict and Surveil: Data, Discretion, and the Future of Policing*; Oxford University Press: New York, NY, USA, 2020.
19. Iliadis, A.; Acker, A. The seer and the seen: Surveying Palantir’s surveillance platform. *Inf. Soc.* **2022**, *38*, 334–363. [CrossRef]
20. Munn, L. *Logic of Feeling: Technology’s Quest to Capitalize Emotion*; Rowman & Littlefield: Lanham, MD, USA, 2018.
21. Engstrom, D.F.; Ho, D.E.; Sharkey, C.M.; Cuéllar, M.-F. *Government by Algorithm: Artificial Intelligence in Federal Administrative Agencies*; Report for the Administrative Conference of the United States: Washington, DC, USA, 2020.
22. Citron, D.K.; Calo, R. The automated administrative state: A crisis of legitimacy. *Emory Law J.* **2021**, *70*, 797–845.
23. Coglianese, C.; Lampmann, E. Contracting for algorithmic accountability. *Adm. Law Rev. Accord* **2021**, *6*, 175–199.
24. Brauneis, R.; Goodman, E.P. Algorithmic transparency for the smart city. *Yale J. Law Technol.* **2018**, *20*, 103–176.
25. Jobin, A.; Ienca, M.; Vayena, E. The global landscape of AI ethics guidelines. *Nat. Mach. Intell.* **2019**, *1*, 389–399. [CrossRef]
26. Mittelstadt, B.D.; Allo, P.; Taddeo, M.; Wachter, S.; Floridi, L. The ethics of algorithms: Mapping the debate. *Big Data Soc.* **2016**, *3*, 1–21. [CrossRef]
27. Floridi, L.; Cows, J. A unified framework of five principles for AI in society. *Harv. Data Sci. Rev.* **2019**, *1*. [CrossRef]
28. Margetts, H.; Dorobantu, C. Rethink government with AI. *Nature* **2019**, *568*, 163–165. [CrossRef] [PubMed]
29. Dreyling, R.; Tammet, T.; Pappel, I.; McBride, K. Navigating the AI maze: Lessons from Estonia’s Bürokratt on public sector AI digital transformation. In *Proceedings of the 22nd Annual International Conference on Digital Government Research (dg.o 2021)*; ACM: New York, NY, USA, 2021; pp. 67–75.
30. Misuraca, G.; van Noordt, C. *AI Watch—Artificial Intelligence in Public Services: Overview of the Use and Impact of AI in Public Services in the EU*; EUR 30255 EN; Publications Office of the European Union: Luxembourg, 2020. [CrossRef]
31. McIntyre, L. *Post-Truth*; MIT Press: Cambridge, MA, USA, 2018.
32. Zuboff, S. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*; PublicAffairs: New York, NY, USA, 2019.
33. Crawford, K. *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*; Yale University Press: New Haven, CT, USA, 2021.
34. Pasquale, F. *The Black Box Society: The Secret Algorithms That Control Money and Information*; Harvard University Press: Cambridge, MA, USA, 2015.
35. Checkland, P. *Systems Thinking, Systems Practice*; Wiley: Chichester, UK, 1981.
36. Churchman, C.W. *The Systems Approach*; Dell: New York, NY, USA, 1968.
37. Meadows, D.H. *Thinking in Systems: A Primer*; Chelsea Green Publishing: White River Junction, VT, USA, 2008.
38. Midgley, G. *Systemic Intervention: Philosophy, Methodology, and Practice*; Kluwer Academic/Plenum: New York, NY, USA, 2000.
39. Speed, R. UK Weighs Break Clause in Palantir NHS Deal. *The Register*, 20 April 2026. Available online: https://www.theregister.com/2026/04/20/palantir_nhs_break_clause/ (accessed on 4 May 2026).
40. House of Commons. NHS Federated Data Platform. Hansard, Volume 765, Westminster Hall Debate, 16 April 2026. Available online: <https://hansard.parliament.uk/commons/2026-04-16/debates/2FDC71C-D0C1-4738-BEE8-A4BDA311DB99/NHSFederatedDataPlatform> (accessed on 4 May 2026).
41. U.S. Congress. Clarifying Lawful Overseas Use of Data Act (CLOUD Act), Public Law No. 115–141, Division V, 132 Stat. 1213; 2018; codified at 18 U.S.C. Section 2713. Available online: <https://www.govinfo.gov/link/plaw/115/public/141> (accessed on 8 July 2026).
42. Stewart, R.; Campbell, A. Are American Tech Billionaires Threatening British Democracy? The Rest Is Politics, Episode 527, Goalhanger Podcasts, 29 April 2026. Available online: <https://podcasts.apple.com/gb/podcast/are-american-tech-billionaires-threatening-british/id1611374685?i=1000764149437> (accessed on 4 May 2026).
43. House of Commons Science, Innovation and Technology Committee. Oral Evidence: Digital Centre of Government Inquiry, HC 16290. Witness: Louis Mosley, Executive Vice-President, Palantir Technologies; 8 July 2025. Available online: <https://committees.parliament.uk/oralevidence/16290/html/> (accessed on 4 May 2026).
44. Medact. *Briefing: Concerns Regarding Palantir Technologies and NHS Data Systems*; Medact: London, UK, 2026. Available online: <https://www.medact.org/briefing-palantir-fdp> (accessed on 4 May 2026).
45. Roth, E. Palantir posts mini-manifesto denouncing inclusivity and ‘regressive’ cultures. *TechCrunch*, 19 April 2026. Available online: <https://techcrunch.com/2026/04/19/palantir-posts-mini-manifesto-denouncing-regressive-and-harmful-cultures/> (accessed on 30 April 2026).

46. Palantir Technologies. The Technological Republic, in Brief [22-Point Thread]. X, 19 April 2026. Available online: <https://x.com/PalantirTech/status/2045574398573453312> (accessed on 30 April 2026).
47. Karpf, D. Palantir's Manifesto Is as Subtle as a MAGA Hat. *Tech Policy Press*, 28 April 2026. Available online: <https://www.techpolicy.press/palantirs-manifesto-is-as-subtle-as-a-maga-hat/> (accessed on 30 April 2026).
48. Jones, J. Palantir's dystopian manifesto sparks bipartisan blowback. *MS NOW Opinion*, 24 April 2026. Available online: <https://www.ms.now/opinion/palantir-manifesto-karp-lonsdale-trump-surveillance-ai> (accessed on 30 April 2026).
49. Al Jazeera Staff. 'Technofascism': Critics Accuse Palantir of Pushing AI War Doctrine. Al Jazeera, 20 April 2026. Available online: <https://www.aljazeera.com/news/2026/4/20/technofascism-critics-accuse-palantir-of-pushing-ai-war-doctrine> (accessed on 30 April 2026).
50. Reisman, D.; Schultz, J.; Crawford, K.; Whittaker, M. *Algorithmic Impact Assessments: A Practical Framework for Public Agency Accountability*; AI Now Institute: New York, NY, USA, 2018.
51. Gwallter. Who is Louis Mosley? Gwallter, 13 February 2026. Available online: <https://gwallter.com/politics/who-is-louis-mosley.html> (accessed on 4 May 2026).
52. Trump, D.J. Establishing and Implementing the President's "Department of Government Efficiency." Executive Order 14158, 20 January 2025. *Fed. Regist.* **2025**, *90*, 8441–8443. Available online: <https://www.federalregister.gov/documents/2025/01/29/2025-02005/establishing-and-implementing-the-presidents-department-of-government-efficiency> (accessed on 4 May 2026).
53. Congressional Research Service. *Department of Government Efficiency (DOGE) Executive Order: Early Implementation; Insight* IN12493; Library of Congress: Washington, DC, USA, 2025. Available online: <https://crsreports.congress.gov/product/pdf/IN/IN12493> (accessed on 4 May 2026).
54. Davenport, C.; Rein, L. DOGE's efforts to make government more efficient are doing the opposite. *Washington Post*, 2 June 2025. Available online: <https://www.washingtonpost.com/business/2025/06/02/doge-vowed-make-government-more-efficient-its-doing-opposite/> (accessed on 4 May 2026).
55. Alon-Barkat, S.; Busuioc, M. Human-AI interactions in public sector decision making: "Automation bias" and "selective adherence" to algorithmic advice. *J. Public Adm. Res. Theory* **2023**, *33*, 153–169. [CrossRef]
56. Skitka, L.J.; Mosier, K.L.; Burdick, M. Does automation bias decision-making? *Int. J. Hum.-Comput. Stud.* **1999**, *51*, 991–1006. [CrossRef]
57. Goddard, K.; Roudsari, A.; Wyatt, J.C. Automation bias: A systematic review of frequency, effect mediators, and mitigators. *J. Am. Med. Inform. Assoc.* **2012**, *19*, 121–127. [CrossRef] [PubMed]
58. Wirtz, B.W.; Weyerer, J.C.; Geyer, C. Artificial intelligence and the public sector—Applications and challenges. *Int. J. Public Adm.* **2019**, *42*, 596–615. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.