

Artificial intelligence adoption in public organisations: evidence from the Italian welfare agency

International
Journal of Public
Sector
Management

Luigina Paglieri and Fabiana Scalabrini

Department of Management and Law, University of Rome Tor Vergata, Rome, Italy

Greta Nasi

Department of Social and Political Sciences, Bocconi University, Milano, Italy, and

Denita Cepiku

Department of Management and Law, University of Rome Tor Vergata, Rome, Italy

Received 17 November 2025
Revised 9 March 2026
27 April 2026
Accepted 30 April 2026

Abstract

Purpose – This study aims to examine how organisational, technological and environmental factors interact to enable systematic artificial intelligence (AI) adoption in a large welfare administration operating under constitutional social protection obligations and demographic pressure; it further asks whether, and how, these factors manifest differently across two categories of AI use (citizen-facing versus internal administrative systems).

Design/methodology/approach – Using a single-case study of Italy's National Social Security Institute, the research applies an adapted technology-organisation-environment (TOE) framework for the public sector. The empirical base combines 10 semi-structured interviews with senior executives, extensive analysis of regulatory and strategic documents and validated performance indicators for a broad AI portfolio.

Findings – Environmental pressures – ageing populations, constitutional obligations and fiscal constraints – are the main drivers of AI adoption, altering the relative weight of dimensions in the traditional TOE framework. Systematic implementation is sustained by performance measurement and monitoring, robust data governance on a mature technological infrastructure and leadership that explicitly embeds AI within the social protection mandate.

Research limitations/implications – As a single-case study focused on senior leadership, the findings require comparative validation across institutional and national settings. Transferability is conditioned on the presence of constitutional welfare mandates, strong demographic pressure and mature information technology infrastructure.

Practical implications – AI adoption should follow a clearly defined strategy, supported by performance systems tracking efficiency and equity and by leadership able to translate constitutional principles into technological governance.

Originality/value – The study addresses the lack of empirical research on AI adoption in large European welfare institutions beyond Anglo-Saxon and Nordic contexts. We propose organisational metabolism as a mid-range theoretical construct, a sector-specific refinement of dynamic capabilities, denoting a sustained organisational capacity for continuous technological absorption, adaptive reconfiguration and normative calibration to constitutional values across political and regulatory cycles.

Keywords Artificial intelligence, Public sector innovation, Welfare state, Digital transformation, TOE framework

Paper type Research article

1. Introduction

Public administrations across the globe are adopting artificial intelligence (AI) technologies through diverse modalities that reflect their specific regulatory, procedural and organisational frameworks, thereby engaging in empirical experimentation with the capacity of these tools to



© Luigina Paglieri, Fabiana Scalabrini, Greta Nasi and Denita Cepiku. Published by Emerald Publishing Limited. This article is published under the Creative Commons Attribution (CC BY 4.0) licence. Anyone may reproduce, distribute, translate and create derivative works of this article (for both commercial and non-commercial purposes), subject to full attribution to the original publication and authors. The full terms of this licence may be seen at <http://creativecommons.org/licenses/by/4.0/>

International Journal of Public Sector
Management
Emerald Publishing Limited
e-ISSN: 1758-6666
p-ISSN: 0951-3558
DOI 10.1108/IJPSM-11-2025-0543

produce demonstrable public value (Mergel *et al.*, 2019; Margetts and Dorobantu, 2019; Wirtz *et al.*, 2019; Zuiderwijk *et al.*, 2021). The effective use of AI depends on far more than deploying new tools: it requires robust data governance, the integration of processes and legacy systems with existing accountability mechanisms and institutional cultures and strict compliance with applicable regulations. In the absence of these conditions, AI often remains stuck in pilot stages and can even undermine the very public values that administrations are meant to uphold (Bovens and Zouridis, 2002; Eubanks, 2018).

While evidence on AI's impacts remains emerging, existing literature focuses predominantly on how public organisations navigate fairness, transparency and democratic legitimacy in AI deployment (Dwivedi *et al.*, 2021; Guenduez and Mettler, 2023; Hjaltalin and Sigurdarson, 2024). Yet, three asymmetries in this literature motivate the present study. First, most empirical evidence is drawn from Anglo-Saxon, Nordic and East-Asian settings; continental European welfare states with codified legal traditions and constitutionally entrenched social rights remain comparatively understudied. Second, most studies examine municipalities or mid-sized agencies rather than national infrastructural institutions whose AI decisions directly affect the majority of citizens. Third, when legally binding welfare obligations meet severe demographic pressure, adoption shifts from a discretionary modernisation choice to a non-negotiable response to rights-based mandates, a configuration the current technology-organization-environment (TOE) literature does not fully theorise.

Against this backdrop, the paper asks a single, focused research question: *How do organisational, technological and environmental factors interact to enable systematic AI adoption in a large public organisation operating under constitutional social-equity imperatives, and do these interactions differ across citizen-facing and internal administrative AI systems?* We address this question through a single extreme-case study (Yin, 2018; Flyvbjerg, 2006) of the Italian National Social Security Institute (INPS), a constitutionally anchored welfare agency that combines high penetration, demographic exposure and institutional maturity, features that position underlying causal mechanisms empirically as visible.

To keep the narrative focused, the precise definitional boundaries of the AI technologies under study, including the exclusion of generative large language models and the distinction between citizen-facing and internal administrative systems, are developed at the start of the literature review (Section 2.1). Here we merely signal that the paper is concerned with the adoption of *traditional, narrow AI* and that the exclusion of generative AI is treated both as a scope choice and as a substantive limitation that is revisited in Section 6.2.

We adopt the TOE framework (Tornatzky and Fleischer, 1990; Baker, 2012) as the main analytical lens because of its three virtues: an organisation-level focus appropriate for institutional decision-making, an environmental dimension explicitly designed to accommodate regulatory and institutional pressures and a conceptual openness that lends itself to sector-specific adaptation while preserving comparability across studies (Oliveira and Martins, 2011). The detailed justification of this choice relative to alternative frameworks and the specific adaptations we introduce for the public-sector welfare context are developed in Section 2.2.

The paper aims to offer three contributions. Theoretically, it extends the TOE framework by showing that under a binding constitutional mandate and severe demographic pressure, the environmental dimension functions not as a moderating context but as the primary trigger of adoption, and it introduces *organisational metabolism* as a mid-range construct that captures the temporally extended, normatively calibrated adaptive capacity of constitutionally anchored public organisations: a construct we explicitly position against absorptive capacity, dynamic capabilities and institutional theory. Empirically, it documents large-scale AI adoption in a continental European welfare institution, extending a literature still dominated by Anglo-Saxon, Nordic and Chinese settings, and it shows how the identified factors manifest differently in citizen-facing versus internal administrative AI systems.

Practically, it articulates four principles – constitutional anchoring of AI strategy, long-term IT–policy integration, balanced efficiency-and-equity metrics and risk-calibrated governance intensity – that can inform peer institutions facing comparable adoption challenges.

The paper is organised as follows. [Section 2](#) develops the literature review around the TOE sub-dimensions and articulates the theoretical framework, including the positioning of *organisational metabolism*. [Section 3](#) describes the research design. [Section 4](#) presents the findings, organised by TOE dimension and with explicit mapping to the analytical framework introduced in [Section 2](#). [Section 5](#) discusses the theoretical implications. [Section 6](#) concludes and delineates scope conditions and limitations.

2. Literature review and theoretical framework

2.1 Scope of AI technologies examined

Before reviewing the literature, we specify the AI technologies within the scope of this study. Following the [OECD \(2024\)](#), we treat AI as machine-based systems that, for a given set of objectives, generate outputs, predictions, recommendations, decisions or content capable of influencing real or virtual environments. Within this broad definition, we restrict attention to three technology families used in welfare administration: (1) machine learning for predictive analytics and fraud detection; (2) natural language processing for conversational citizen interfaces and document classification and (3) rule-based routing for workflow optimisation. Generative large language models are explicitly excluded from the empirical scope because the organisation did not deploy them during the study period and because their public-sector governance raises distinct normative and regulatory issues that would overstretch a single case. This exclusion is not innocuous: it implies that our findings speak primarily to the adoption of *traditional, narrow AI*. The limits this places on transferability and timeliness are addressed in [Section 6.2](#).

We further distinguish two categories of AI use that have distinct adoption dynamics and governance implications. Citizen-facing AI systems (e.g. chatbots, virtual assistants and automated contact analysis) interact directly with citizens and are judged primarily on legitimacy, accessibility and service quality; a chatbot that misroutes an enquiry is, at its worst, a service-quality failure. Internal administrative AI systems (e.g. fraud detection, document classification and decision-support tools) support back-office processes and are judged on compatibility with legacy systems, procedural reliability and – when they bear on individual rights – procedural fairness; a fraud-detection model that wrongly flags a pensioner's benefits can infringe fundamental rights ([Veale and Brass, 2019](#)). We return to this distinction in [Sections 4 and 5](#), where we show that the two categories exhibit different factor signatures within the TOE framework.

2.2 AI adoption in the public sector: what the literature knows

AI in public administration is not implemented within an institutional vacuum. In contrast to market-oriented contexts, public organisations operate under dense legal obligations, democratic accountability mechanisms and entrenched normative expectations – conditions that are particularly pronounced in the European setting. These institutional structures shape which technologies are regarded as legitimate objects of investment, how they are governed and overseen and according to which criteria their performance and societal impacts are assessed.

National administrative traditions, constitutional configurations and prevailing conceptions of the state influence these processes in ways that are often obscured by cross-national generalisations. A public agency deploying AI under a constitutional mandate to safeguard social rights is institutionally and normatively distinct from a municipality experimenting with a chatbot to enhance administrative efficiency. The difference is not merely a matter of organisational size or functional scope: the underlying normative stakes, accountability logics and political risks associated with failure diverge fundamentally.

Existing empirical research has only begun, often in partial, fragmented and case-specific ways, to furnish the requisite contextualisation. In what follows, we organise the relevant evidence around the three TOE dimensions so that, by the end of the section, the specific factors and sub-factors that Table 1 will retain are grounded in prior scholarship rather than introduced as an uncommented list.

2.2.1 Technological factors. Across the adoption literature, three technological sub-factors recur. Relative advantage, intended as the extent to which an innovation is perceived as superior to the status quo, remains the single most robust predictor of adoption (Rogers, 2003; Tornatzky and Fleischer, 1990) and has been shown to matter equally in public-sector machine learning (ML) deployments (Pumplun et al., 2019; Neumann et al., 2024). Compatibility – intended as the alignment of the new technology with legacy systems, workflows and institutional routines – has been identified by Baker (2012) as particularly consequential in bureaucratic settings with long-standing information technology (IT) stacks; Neumann et al. (2024) confirm this in Swiss public organisations. Complexity, intended as the perceived difficulty of understanding and operating the technology, tempers adoption when internal

Table 1. Analytical framework: TOE factors retained, with literature-based justification

Dimension	Factor/sub-factor	Key literature and rationale for inclusion
Technological	Relative advantage	Rogers (2003), Tornatzky and Fleischer (1990) and Pumplun et al. (2019). Most robust predictor of adoption; in welfare AI, advantage must be weighed against constitutional trade-offs (augmentation over substitution)
Technological	Compatibility (legacy, workflows and data)	Baker (2012) and Neumann et al. (2024). Legacy infrastructure is decisive in bureaucracies with long-standing IT stacks; compatibility with on-premise data regimes is particularly binding under GDPR
Technological	Complexity	Rogers (2003) and Pumplun et al. (2019). AI-specific opacity magnifies complexity; mitigated by in-house expertise and IT-policy integration
Organisational	Top-management support	Mikalef et al. (2022) and Chen et al. (2024). Leadership bridges environmental pressures and allocative decisions
Organisational	Culture (innovation climate and change readiness)	Neumann et al. (2024) and Damanpour (1991). AI adoption collides with routine-driven bureaucratic cultures, culture conditions the pace of change
Organisational	Resources (financial, human capital and data)	Mikalef et al. (2022) and Pumplun et al. (2019). Data resources are a distinctive AI determinant; human-capital scarcity can paradoxically accelerate adoption
Organisational	Structure	Baker (2012) and Fountain (2001). Formalisation and centralisation condition the form of adoption
Organisational	Absorptive capacity	Cohen and Levinthal (1990). Cumulative knowledge shapes ability to integrate new technologies; extended in Section 5.2 into organisational metabolism
Environmental	Regulatory framework (GDPR, AI Act and constitutional mandates)	Neumann et al. (2024) and Grimmelikhuijsen and Meijer (2022), extended to constitutional norms
Environmental	Stakeholder pressures (citizens, employees and unions)	Wang et al. (2024) and Alon-Barkat and Busuioic (2023). Stakeholder pressures shape legitimacy judgements, particularly for citizen-facing systems
Environmental	External support (policy, funding and networks)	Mikalef et al. (2022). National policies and EU funding operate as enablers, particularly under resource scarcity
Environmental	Demographic/fiscal pressures	Berryhill et al. (2019), this study. In welfare institutions, demographic trajectories function as non-discretionary adoption triggers

expertise is scarce (Rogers, 2003; Pumplun *et al.*, 2019). In AI contexts specifically, complexity is magnified by opacity: deep-learning models add an interpretability layer that rule-based systems did not have. We retain all three as technological sub-factors in Table 1 and return to each when presenting the findings (Section 4.4).

2.2.2 Organisational factors. At the organisational level, prior studies converge on five sub-factors that we retain in Table 1. Top-management support, including visible leadership commitment and strategic prioritisation, emerges as decisive across very different settings (Mikalef *et al.*, 2022; Chen *et al.*, 2024), with leadership often acting as the bridge between environmental pressures and concrete allocative decisions. Organisational culture, conceived here as innovation climate and change readiness, matters because AI adoption frequently collides with routine-driven bureaucratic cultures (Neumann *et al.*, 2024; Damanpour, 1991). Resources, financial, human-capital and data resources, are strong enablers, with data resources increasingly recognised as a distinctive determinant for AI rather than for other information and communication technology (ICT) adoptions (Mikalef *et al.*, 2022; Pumplun *et al.*, 2019). Organisational structure, the degree of formalisation, centralisation and hierarchical complexity, conditions both the speed of decision-making and the form that adoption takes (Baker, 2012; Fountain, 2001). Absorptive capacity (Cohen and Levinthal, 1990) – the ability to recognise, assimilate and apply external knowledge – refines the organisational dimension by foregrounding the cumulative effects of prior experience. We retain absorptive capacity in Table 1 but also argue in Section 5.2 that, in the public-welfare setting we study, it needs to be supplemented by a broader construct, organisational metabolism, that extends absorptive capacity along temporal, normative and constitutional axes.

2.2.3 Environmental factors. On the environmental side, we retain three sub-factors. Regulatory frameworks have been repeatedly shown to shape, and in some public-sector settings to mandate, adoption trajectories (Neumann *et al.*, 2024; Grimmelikhuijsen and Meijer, 2022). We extend this sub-factor to include constitutional provisions (notably Article 38 of the Italian Constitution), as, in welfare contexts, constitutional norms are not a diffuse background but a binding operational perimeter. Stakeholder pressures from citizens, trade unions and employees are documented in Wang *et al.* (2024) on citizens' trust and in Alon-Barkat and Busuioc (2023) on the interaction between officials' roles and algorithmic advice. External support, from national digital-government policies, European Union (EU) funding and cross-agency initiatives, has been identified by Mikalef *et al.* (2022) as an enabler of adoption, especially in settings with limited internal slack.

2.2.4 Complementary perspectives: why we retain them and how they complement TOE. Three additional perspectives inform our analysis, and we now specify exactly how each of them complements TOE rather than substitutes for it. Algorithmic legitimacy (Grimmelikhuijsen and Meijer, 2022; Alon-Barkat and Busuioc, 2023) refines the *environmental sub-factor of stakeholder pressures*: it explains why, in citizen-facing applications, legitimacy judgements operate as an adoption constraint that goes beyond statutory regulation. The automation–augmentation paradox (Raisch and Krakowski, 2021; Bullock *et al.*, 2020) refines the *technological sub-factor of relative advantage* by showing that in welfare settings, the relevant advantage is not the absolute efficiency of automation but the marginal value of augmenting expert human judgement under constitutional constraints. Dynamic capabilities (Teece, 2007) refine the *organisational sub-factor of absorptive capacity* by adding temporality (sensing–seizing–reconfiguring) to knowledge accumulation; we show in Section 5.2 that dynamic capabilities are a necessary but insufficient foundation for organisational metabolism because they were theorised for competitive repositioning rather than for normative calibration.

2.3 Empirical evidence from prior studies

Scholarship originating from the Nordic context has investigated AI adoption at the municipal level in Denmark, Norway and Finland, demonstrating that technical competence,

organisational innovativeness and governmental support function as key determinants (Mikalef *et al.*, 2022). These contributions derive particular analytical leverage from examining environments characterised by a high degree of digital maturity and institutional cultures receptive to experimentation, conditions especially prevalent in administrative systems with long-standing traditions of public-sector modernisation.

Neumann *et al.* (2024) offer a recent comparative analysis of eight Swiss public organisations at different maturity levels. Their main contribution is adding a temporal dimension to the TOE framework, treating adoption as dynamic organisational learning rather than a one-off decision. The Swiss cases show relatively limited regulatory influence, though this finding should be generalised cautiously. Early scoping studies (Berryhill *et al.*, 2019; Mehr, 2017) catalogued emerging AI applications across many governmental functions, such as citizen services and fraud detection, but, conducted before later empirical work, did not capture actual adoption dynamics and trajectories. Research on China's public sector (Wang *et al.*, 2024) stresses the role of top-down political directives and citizen acceptance in chatbot deployment. In the United States of America, Chen *et al.* (2024) examine state-level chatbot adoption, emphasising strategic alignment and organisational capacity for change. Across this literature, leadership commitment, technical capabilities, clear strategic priorities and cultures that enable experimentation consistently emerge as key determinants of successful adoption.

2.4 Three theoretical tensions and three empirical lacunae

Beyond these empirical contributions, three theoretical tensions remain relevant for our study.

The first is the automation–augmentation paradox. Raisch and Krakowski (2021) show that organisations face a structural dilemma between efficiency gains from automating labour and capability gains from augmenting human judgement. How, and under which conditions, public organisations face the same dilemma is under-theorised. Bullock *et al.* (2020) argue that AI reshapes bureaucratic structures and threatens the discretionary judgement of street-level bureaucrats, who adapt general rules to individual cases – a capacity intentionally built into welfare states. Bannister and Connolly (2014) argue more generally that digital transformation in government must be assessed against public values such as equity, transparency and accountability rather than efficiency alone. The public-sector automation–augmentation paradox is thus constrained by a different value regime: the key decision framework is constitutional mandate, not profit maximisation.

The second tension concerns algorithmic legitimacy. Grimmelikhuijsen and Meijer (2022) show that algorithmic decision-making in government faces legitimacy challenges beyond accuracy and efficiency. Citizens judge automated decisions by procedural justice, institutional trust and perceived fairness. Alon-Barkat and Busuioc (2023) demonstrate that human–AI interactions in public decision-making generate systematic biases – including automation bias and selective adherence – shaped by officials' institutional roles.

The third tension is conceptual: Are existing constructs of organisational adaptation to technology adequate? Cohen and Levinthal's (1990) absorptive capacity helps explain differential success in integrating new technologies but was not designed to explain firms' competitive advantage. Teece's (2007) dynamic capabilities framework captures strategic agility but is oriented to competitive repositioning. Fountain's (2001) technology enactment framework shifts attention from markets to institutions but focuses on a specific institutional moment rather than decades-long trajectories. These theories illuminate parts of the adoption puzzle but do not capture the distinctive adaptation logic of constitutionally anchored public organisations.

These tensions intersect with three empirical lacunae. First, continental European bureaucracies – particularly Southern European welfare states characterised by codified legal procedures and constitutionally entrenched social rights – remain comparatively understudied. Second, most research focuses on municipalities or medium-sized agencies; adoption dynamics in organisations that are infrastructural to national social policy are

underinvestigated. Third, constitutional mandates may reshape the underlying rationale of adoption, shifting it from discretionary modernisation to legally prescribed service obligation.

2.5 Theoretical framework: TOE and its adaptation

Information systems theories have, with few exceptions, been theorised from within the horizon of market-based organisations, where adoption decisions are shaped by competitive dynamics, profit imperatives and relatively unconstrained technological choice (Bozeman and Bretschneider, 1986; Cordella and Iannacci, 2010). Transporting these frameworks to public-sector settings is not a straightforward exercise. Public organisations are subject to legal mandates that no private firm faces, are accountable to democratic constituencies and are embedded in institutional arrangements that constrain and direct technological choice in ways the original frameworks were not designed to handle (Rainey *et al.*, 1976; Rainey, 2014; Bretschneider and Mergel, 2011). The TOE framework is unusually well placed for this purpose, precisely because its environmental dimension was always intended to accommodate institutional and regulatory pressures, not merely market competition (Tornatzky and Fleischer, 1990; Baker, 2012).

The TOE framework conceptualises adoption decisions across three dimensions: technological (relative advantage, compatibility and complexity), organisational (size, resources, leadership, structures, culture and absorptive capacity) and environmental (regulatory pressures, stakeholder expectations and external support). We adopt TOE for the three reasons listed in Section 1: organisation-level focus, explicit environmental dimension and openness to sector-specific adaptation (Oliveira and Martins, 2011).

Building on Neumann *et al.*'s (2024) temporal extension of TOE, we advance two context-specific adaptations. First, we explicitly foreground the regulatory–constitutional dimension within the environmental category. Article 38 of the Italian Constitution enshrines inalienable rights to social protection; the General Data Protection Regulation (GDPR) establishes requirements for algorithmic transparency and explainability; the EU AI Act introduces risk-based classifications for high-stakes public-sector AI applications; co-determination arrangements regulate organisational transformation and public procurement legislation constrains vendor relationships. Second, we synthesise institutional theory's focus on legitimacy with the dynamic capabilities perspective on organisational adaptation because – in the welfare setting – neither lens alone captures the sustained, normatively constrained adaptive capacity that we will define in Section 5.2 as *organisational metabolism*.

Table 1 summarises the resulting analytical framework. Each factor has been retained because it is grounded in one or more of the literature reviewed in Sections 2.2–2.4, as indicated by the “Key literature” column; it is not a comprehensive list of every possible factor but an analytically disciplined selection of those for which both the prior literature and the INPS case are expected to provide actionable evidence.

3. Research design and methodology

3.1 Case selection

The study adopts a single-case study design for theory-building purposes (Yin, 2018). Single-case designs are particularly suitable when the empirical setting constitutes a unique, revelatory or extreme instance that can illuminate theoretical mechanisms otherwise difficult to observe (Flyvbjerg, 2006). INPS can be characterised as an extreme case along four concurrent dimensions: societal penetration (67.9% of the population); fiscal magnitude (14.62% of gross domestic product (GDP); 28.92% of public spending); demographic pressure (647 pensioners per 1,000 workers; projected old-age dependency ratio of 75% by 2050 (OECD, 2023) and programmatic complexity (over 500 distinct benefit programmes). The analytical value of an extreme case does not derive from representativeness but from its capacity to make otherwise latent causal mechanisms empirically visible.

3.2 Data sources and collection

Data collection was conducted between August and October 2025 and was structured in three phases. The documentary analysis encompassed two categories of sources. The first comprised institutional strategy and governance documents, including the INPS Integrated Plan for Activities and Organisation (PIAO 2024–2026), the XXII Annual Report (Rapporto Annuale, 2024), the Director General's Directive No. 8 of 8 April 2024 (Linee guida sull'implementazione di sistemi di Intelligenza Artificiale in INPS) and the consolidated financial statements for 2024. The second category included regulatory and policy documents: constitutional and legislative provisions (Article 38 of the Italian Constitution, Law 88/1989 and subsequent amendments), European-level instruments (GDPR, EU AI Act, EU Guidelines for Trustworthy AI) and national digital-government policies (Italian Digital Agenda; Three-Year Plan for ICT in Public Administration; AgID directives).

The directive constitutes the primary documentary reference. It formalises INPS's AI adoption trajectory in six phases – promotion, prioritisation, monitoring, compliance and risk assessment, project lifecycle integration and solution standardisation – operationalised through an AI Canvas, a Prioritisation Matrix, a Monitoring Cockpit and risk-classification procedures aligned with the EU AI Act's four-level taxonomy. The directive also codifies the IA@SCALE framework, which regulates development, validation and scaling of AI systems within the enterprise architecture.

3.2.1 Interviews. [Appendix A](#)-interview protocol, comprised 23 questions organised around TOE dimensions; the complete protocol is provided in [Annexure 1](#). In total, 10 interviews were conducted with senior executives (Director General; three directors of the central directorates most exposed to AI deployment; two members of the technical secretariat supporting the Director General and two external experts who had served as senior advisors during the 2020–2024 period). The mean duration was 72 minutes (range 48–105). The 23-item protocol was used as a thematic guide rather than a rigid checklist: interviewers prioritised depth over coverage, pursuing the items most germane to each respondent's role and expertise and deferring less-relevant items, consistent with a semi-structured design. Saturation was assessed within the organisational and environmental dimensions, where no new codes emerged after the eighth interview; the technological dimension was intentionally under-saturated, reflecting the evolving nature of the project portfolio, and this is flagged as a limitation in [Section 6.2](#).

3.2.2 Performance metrics validation. For three flagship AI projects, process-level data – such as automation rates, volumes, accuracy and labour-saving estimates – were extracted directly from INPS's IT systems and validated by the IT directorate. This quantitative evidence enables triangulation of claimed benefits against measured outcomes, a step that qualitative AI research often omits.

3.2.3 Researcher positionality and longitudinal participant observation. One of the authors was professionally engaged with INPS from 2008 to 2020, serving progressively as program manager, associate partner and partner within one of the institute's major technology providers. Over those 12 years, the author participated directly in programme governance, infrastructure modernisation and digital service design; we therefore characterise this background as longitudinal participant observation ([Spradley, 1980](#)): sustained, professionally embedded exposure to the research setting preceding and methodologically distinct from the focused data collection conducted between August and October 2025. The combination of extended insider knowledge and subsequent scholarly repositioning supports interpretive depth. Analytical coding was cross-checked by the co-author without INPS experience; empirical claims resting on interview evidence were triangulated with documentary sources or performance data and the interview protocol was designed to elicit critical and self-critical accounts.

3.3 Data analysis

Data analysis followed a theory-guided iterative coding process consistent with abductive reasoning in qualitative case research ([Dubois and Gadde, 2002](#)). The analysis proceeded in

three stages. First, open coding identified references to TOE factors across interview transcripts, internal documents and performance records. Second, codes were grouped into higher-order categories aligned with TOE dimensions, while emergent themes – most notably the patterns eventually conceptualised as organisational metabolism – were tracked separately for theoretical development. Third, axial coding explored relationships across dimensions, revealing the adoption trajectory. Methodological triangulation (Denzin, 1978) of interviews, documents and performance data provided a validity mechanism.

4. Findings

Findings are organised to first establish the case context (4.1), then present evidence for each TOE dimension with explicit mapping to the factors listed in Table 1 (4.2–4.4) and finally analyse cross-dimensional interaction (4.5). Environmental factors are presented first, departing from the conventional TOE sequence, to reflect the study's central finding: their role as initiating conditions rather than background context. Table 3 at the end of this section provides the consolidated evidence-to-construct mapping.

4.1 Case context: INPS's institutional profile

INPS did not emerge as a single, unified institution by design. The agency in 2025 is the result of a long consolidation process of smaller and sector-specific social funds, progressively absorbed into INPS by legislative mandate over the last 2 decades. The Italian legislator sought to eliminate fragmentation and achieve economies of scale. With every new incorporation, INPS was asked not merely to integrate the incoming fund's beneficiaries and obligations but also to visibly reduce overall public expenditure. Citizens and policymakers demanded broader coverage, faster processing, more accessible digital channels and greater responsiveness. The result is an institution caught in a structural tension that has shaped its technological and organisational trajectory: a mandate to do progressively more with progressively less.

As of January 2025, approximately 40.4 million individuals were registered as INPS users (67.9% of the resident population; 77.9% of citizens aged 15 and over). Within the active labour market, INPS insurance schemes cover approximately 23 million workers, accounting for 94.9% of the national labour force. The institute manages approximately 20.5 million pensions, roughly 91.9% of all Italian pensions. Pension expenditure alone reached €304bn in 2024 (14.62% of GDP; 28.92% of total public expenditure). The 2024 financial year closed with a budget surplus of €15.0bn. INPS administers over 500 distinct benefit programmes. The AI portfolio examined here comprises 47 projects spanning conversational systems, document management, fraud prevention and internal process optimisation. Figure 1 maps the portfolio along two axes: implementation status (in production, in progress and planned) and technological family. Two patterns are worth noting. First, natural language processing and chatbot applications dominate the portfolio (21 of 47 projects), reflecting the strategic priority assigned to citizen-facing services. Second, roughly one-third of the projects (17 of 47) are already operational, while a further 22 are in active development – an indication that AI

	Production	In Progress	Planned	Total
TOTAL	17	22	8	47
Natural Language Processing & Chatbots	9	9	3	21
Process Automation	3	4	0	7
Document Intelligence	2	5	3	10
Productivity Enhancement	2	2	2	6
Predictive Analytics & Fraud Detection	1	2	0	3

Figure 1. AI project portfolio in INPS

adoption at INPS has moved beyond an experimental phase and is now embedded in routine institutional planning. Detailed performance data for four flagship operational systems are reported in [Table 2](#).

Overall, the 4 projects' quantified time savings exceed 233,000 h, corresponding to approximately 29,000 person-days or the annual workload of roughly 133 full-time equivalents reallocated to higher-value activities. The Posta Elettronica Certificata, the Italian certified e-mail system (PEC) classification system, active since 2021, exhibits sustained performance over four years, with a precision of 85% in PEC classification. While operationally significant at this scale, it necessarily implies that approximately 15% of instances require human review or correction through a systematic human-in-the-loop verification. The reported automation rates gain interpretive significance when normalised by underlying volumes. For example, a 96.6% automation rate for contact motivation analysis translates into only about 100,000 cases requiring manual handling; in the absence of AI support, all 3.4 million cases would necessitate human intervention.

4.2 Environmental factors

This section maps evidence onto the three environmental sub-factors retained in [Table 1](#), with explicit indication of their manifestation across AI categories.

Demographic and fiscal pressure ([Table 1](#): Environmental – demographic/fiscal pressure). The Director General described the current configuration as “a demographic and technological ‘perfect storm’.” Persistent low fertility, population ageing and labour-market dynamics are converging to generate systemic pressure on social security administration. For INPS, these figures translate into operational constraints: an expanding pensioner population, a shrinking contributor base and, due to prolonged hiring freezes, a declining number of staff available to manage an increasing volume of complex administrative tasks. Demographic pressure operates equally across both AI categories, functioning as an aggregate adoption imperative rather than as a differentiated driver of specific applications.

Constitutional mandate ([Table 1](#): Environmental – regulatory framework extended to constitutional norms). Article 38 of the Italian Constitution codifies social protection as a fundamental right. The Director General emphasised the non-delegable character of the mission: “*The Public Administration has a constitutional mission that the market cannot pursue: guaranteeing as fundamental rights that are not negotiable. Our foundational values – legality, equity, integrity and social responsibility – must become the genetic code of Italian and European AI.*” Constitutional norms thus function as both an *imperative* (obliging universal access) and a *perimeter* (constraining design and governance). This perimeter operates more stringently on internal administrative AI, where algorithmic outputs bear on individual rights (e.g. fraud detection and entitlement decisions), than on citizen-facing AI, where a misrouted chatbot enquiry is a service-quality issue rather than a rights infringement.

Regulatory framework ([Table 1](#): Environmental – regulatory framework). The interaction of GDPR explainability obligations, the EU AI Act's risk-based classification scheme, employee-council arrangements and public procurement rules generates a multilayered regulatory environment. According to the Director General, these requirements compel INPS to design systems that are transparent, auditable and contestable from the outset. Such *ex ante* compliance reduces legal and reputational risk and enhances procedural trust. Regulatory salience is stronger in citizen-facing systems (where explainability obligations under GDPR are visible to the citizen) and stronger still in internal administrative systems that bear on fundamental rights (where the EU AI Act's high-risk classification applies).

Stakeholder pressures ([Table 1](#): Environmental – stakeholder pressures). Citizen trust and perceived legitimacy ([Wang et al., 2024](#)) dominate adoption constraints for citizen-facing systems; trade-union and employee-council pressures ([Alon-Barkat and Busuioc, 2023](#)), particularly regarding workforce redeployment and discretion preservation, dominate constraints for internal administrative systems.

Table 2. AI projects in INPS

AI system	Operational period	Volume processed	Automation rate	Estimated time saved	Key performance notes
Automated routing system (CRMweb level I and II)	April 2024–September 2025	1.8 M tickets analysed	507K tickets routed directly to back-office specialists	~25,350 h (3,168 person-days)	Bypasses front-office triage, enabling direct assignment to appropriate back-office personnel. Average manual processing: 3 min (front-office) + 7 min (back-office) per ticket
AI Dashboard for video guide creation	January–September 2025	3.4 M contact motivations received	3.3 M analysed automatically (96.6% automation)	~55,000 h (6,875 person-days)	Enables proactive content creation by identifying most frequent citizen enquiries. Communications staff can develop targeted video guides and social media content, reducing incoming query volume. Real-time monitoring of contact patterns by benefit type
Automated PEC classification	April 2021–September 2025	4.8 M certified emails analysed	3.2 M automatically routed (65.9% automation)	~136,000 h (17,000 person-days)	85% precision rate. On-premise BERT-based transformer model ensures GDPR compliance and complete institutional control. Average manual processing: 3 min per email. System handles heterogeneous legal and administrative correspondence
INPSieme welfare program document verification	February–September 2025	207,000 applications analysed	Full automation of initial document verification	~17,250 h (2,156 person-days)	100% precision in document completeness checking for study travel reimbursements. Estimated 50% reduction in case processing time. Average manual review: 10 min per application. <i>Note: First implementation phase, full efficiency gains pending as staff adapt to modified procedures</i>

Note(s): PEC: Posta Elettronica Certificata, the Italian certified e-mail system

External support (Table 1: Environmental – external support). National policies (Three-Year Plan for ICT; AgID directives) and EU AI-related funding streams operate as enablers, though the case shows that their marginal contribution is secondary to the binding obligations produced by the constitutional perimeter.

4.3 Organisational factors

Top management support (Table 1: Organisational – top management support). Strategic orientation was articulated by senior management as follows: *“Our challenge is to evolve from an administrative entity into a cognitive welfare institution, using data, AI and predictive modelling not only to pay benefits but also to design the future of social inclusion.”* The institutional nature of this commitment is codified in Directive No. 8/2024, which operationalises the vision by transforming it into binding implementation obligations assigned to designated central directorates, each associated with defined deliverables and timelines. This formalisation differentiates INPS’s AI adoption trajectory from the discretionary, project-based pattern that typically characterises public-sector innovation. The reference to later conceptualisation in Section 5.2 that appeared in the earlier draft (*“This configuration aligns with what we conceptualise in Section 5.2 as ‘normative calibration’”*) has been removed from the findings to keep this section empirically grounded, in line with the reviewer’s suggestion.

Organisational culture and absorptive capacity (Table 1: Organisational – culture and absorptive capacity). The most distinctive organisational determinant is not recent investment but the cumulative effect of long-term technological embedding: since the 1970s, IT has been treated as core business rather than a peripheral back-office function. As one respondent put it, *“Since the 1970s, IT has been core business, not back office. This means IT people understand social security, not just servers. That cultural integration made AI adoption smoother, as technical staff could translate between code and policy.”* Decades of treating technology as integral to the core mission generated an absorptive capacity specifically calibrated to public-sector constraints and objectives.

Resources (Table 1: Organisational – resources). INPS’s 150-year labour-market database, covering 94.9% of the national workforce, provides training datasets of a temporal depth and population coverage unattainable in the private sector. Simultaneously, human capital scarcity – due to retirements and hiring freezes – produced a paradox: *“It was not a choice, but a necessity that transformed into an opportunity.”* Resource scarcity in human capital catalysed the reconfiguration of processes around data-intensive, AI-enabled tools.

Organisational structure (Table 1: Organisational – structure). The highly centralised structure of INPS accelerated the propagation of the Directive across Central Directorates; we found no evidence that structure varied systematically across AI categories.

4.4 Technological factors

Relative advantage (Table 1: Technological – relative advantage). Within *internal administrative* systems, the dominant advantage is diagnostic *amplification* of professional judgement rather than substitution: in 2024, AI-supported inspection identified irregularities in 82% of 9,701 inspections (€898.7m in detected improprieties); the Director General emphasised that *“AI does not replace the expertise of inspectors, but exponentially amplifies their ability to focus attention”*. Within *citizen-facing* systems, the dominant advantage is *capacity reallocation*: conversational agents handle approximately 60% of routine enquiries, freeing human staff for disability claims, discontinuous contribution histories and exceptional personal circumstances, where empathy and procedural discretion remain indispensable. In both categories, the relative advantage is framed in *augmentation* rather than *substitution* terms – an orientation that is conditioned, as we argue in Section 5.3, by constitutional constraints that rule out substitution-oriented trajectories.

Compatibility (Table 1: Technological – compatibility). AI modules are integrated with on-premise legacy infrastructure that manages several decades of social security data. The decision to retain on-premise deployment, instead of adopting cloud-based architectures that might offer superior scalability, is motivated by GDPR compliance and a strong organisational preference for data sovereignty. Compatibility constraints bind more tightly on internal administrative systems (where model training must access longitudinal benefit data) than on citizen-facing systems (where conversational interfaces can be partially decoupled from core back-end systems).

Complexity (Table 1: Technological – complexity). The case exhibits *low observed complexity* relative to the Rogers (2003) baseline: the long-standing integration between IT and policy functions reduces the opacity of AI systems for the staff who must implement them. Within *internal administrative* systems, however, complexity resurfaces as the *explainability of model outputs* – the 85% precision of the PEC classification system is operationally significant, but the 15% residual requires systematic human-in-the-loop verification, introducing a governance overhead that scales non-linearly with deployment.

Governance architecture – as a technological-organisational interface. INPS's technological governance is formalised through the *IA@SCALE* framework, codified in Directive No. 8/2024, which organises AI development into three operational layers: value management (use-case specification, prioritisation and performance monitoring); scope and risk definition (algorithmic classification, data governance, trustworthiness and compliance) and platform engineering (architectural design, model selection and deployment pipelines). The framework differentiates between *AI-challenging solutions* – exploratory, research-oriented systems, often developed with academic partners, operating outside the core information systems – and *AI industrialised solutions* – standardised, fully integrated components of the enterprise architecture, designed for reuse and scalability across multiple operational domains. This architectural configuration directly addresses the *complexity* sub-dimension of TOE: rather than treating AI adoption as discrete, project-based initiatives, INPS has instituted an organisation-wide scaffolding that minimises redundancy, enforces data sovereignty and data-protection requirements and systematically subjects each AI deployment to an *ex ante* assessment against the EU AI Act's risk categories before production authorisation.

Performance measurement. The INPS performance monitoring system records operational outcomes across AI applications. These indicators are based on organisational self-reporting, validated by the IT directorate rather than by an independent external auditor; they are therefore credible proxies for adoption outcomes but not definitive measures of efficiency. This limitation is acknowledged in Section 6.2.

4.5 Cross-dimensional interaction

The analytical focus concerns not any isolated determinant, but rather the ways in which factors spanning multiple dimensions interact to generate systematic patterns of adoption. Three interaction patterns stand out.

First, environmental imperatives are absorbed by the organisational dimension rather than acting directly on technology: demographic and constitutional pressures are translated into strategic direction, binding directorial obligations and data-governance priorities before they are crystallised in technological choices. Second, organisational absorptive capacity mediates the technological sub-dimensions: the five-decade integration of IT and policy lowers observed complexity and raises compatibility, but it does not uniformly alter relative advantage, which remains category-specific (amplification for internal systems; reallocation for citizen-facing systems). Third, technological choices feed back into the environmental dimension by shaping the legitimacy of INPS's AI governance with respect to citizens, trade unions and oversight bodies: a feedback loop particularly salient for citizen-facing applications, where perceived legitimacy constrains admissible system behaviour. Figure 2 presents this recalibrated analytical framework. The right-hand side of Figure 2 makes explicit how the three TOE




Figure 2. Reweighted TOE framework for constitutional welfare institutions. Source: Authors’ own elaboration

dimensions are sequenced in the INPS adoption logic. Initiating conditions are located in the environmental dimension: the demographic imperative, the constitutional mandate enshrined in Article 38, and the layered regulatory framework (GDPR, EU AI Act, procurement rules) jointly trigger the adoption decision rather than merely contextualising it. Feasibility enactment is located in the organisational dimension, where the three sub-capacities of organisational metabolism, absorptive continuity built through five decades of IT-policy integration, adaptive reconfiguration codified in Directive No. 8/2024 and normative calibration of design choices against constitutional values translate the environmental imperative into implementable trajectories. Augmentation and operationalisation are located in the technological dimension, where relative advantage, compatibility and complexity are resolved through the IA@SCALE governance architecture and through systematically augmentation-centric design choices (human-in-the-loop verification, on-premise data sovereignty and ex ante risk classification). Technological choices feed back into the environmental dimension by shaping the legitimacy of INPS’s AI governance *vis-à-vis* citizens, oversight bodies and social partners, thereby reinforcing or, where contested, reshaping the very environmental conditions from which adoption originated.

Table 3 shows that all 11 Table 1 factors are observed in the case, with varying degrees of empirical support; only *complexity* is partially observed in its Rogers (2003) form and reappears primarily as an explainability burden in internal administrative systems. Two factors (*stakeholder pressures* and *relative advantage*) display category-specific signatures, providing the most visible evidence for the analytical value of the citizen-facing vs internal administrative distinction introduced in Section 2.1.

5. Discussion

The INPS case either corroborates, refines or problematises existing empirical findings. Three propositions are formulated as theoretical contributions to be subjected to systematic comparative testing in future research.

5.1 Reweighting environmental factors: from context to catalyst

Within the TOE framework, environmental factors have conventionally been conceptualised as contextual contingencies – external forces that facilitate or constrain adoption decisions driven fundamentally by organisational evaluations of technological opportunities (Baker, 2012; Tornatzky and Fleischer, 1990). The empirical literature has largely corroborated this secondary, moderating role. Neumann et al. (2024) identified only limited regulatory influence

Table 3. Evidence-to-construct mapping (Table 1 factors against empirical evidence)

TOE sub-factor (from Table 1)	Empirical status	Citizen-facing AI	Internal administrative AI
Relative advantage	Corroborated; category-specific	Capacity reallocation (≈60% automated enquiries)	Amplification of expert judgement (€898.7 M detected, 82% hit rate)
Compatibility	Corroborated; tighter for internal	Moderate (interfaces decoupled)	High (on-premise legacy; 150-yr data)
Complexity	Partially observed; resurfaces as explainability	Low observed	Explainability burden (15% residual, human-in-the-loop)
Top-management support	Corroborated	Directive No. 8/2024 assigns deliverables	Directive No. 8/2024 assigns deliverables
Culture/absorptive capacity	Corroborated and refined	IT-policy integration	IT-policy integration
Resources	Corroborated	Human-capital scarcity → reallocation	Data resources (150-yr DB)
Structure	Corroborated with qualification	Centralised propagation of Directive	Centralised propagation of Directive
Regulatory framework	Corroborated	Stronger explainability salience (GDPR)	Stronger (EU AI Act high-risk for rights)
Stakeholder pressures	Corroborated and differentiated	Citizen trust/legitimacy dominant	Trade-union/employee-council dominant
External support	Partially corroborated	Secondary to constitutional perimeter	Secondary to constitutional perimeter
Demographic/fiscal pressures	Corroborated and refined	Aggregate driver	Aggregate driver

Note(s): DB: Database

in Swiss public organisations. Mikalef *et al.* (2022) observed that government support in Nordic municipalities was supportive rather than decisive. Chen *et al.* (2024) highlighted institutional support structures in USA state governments as enablers rather than primary drivers.

The INPS case exemplifies a qualitatively distinct causal configuration. Demographic pressure interacts with a constitutional mandate codified in Article 38 to produce an adoption imperative rather than a discretionary opportunity. Environmental factors function as the primary initiating mechanism. This inverted configuration is likely to emerge under the conjoint presence of three conditions: (1) binding legal obligations to maintain predetermined service levels irrespective of resource constraints; (2) external and social pressure, such as demographic trajectories that render existing service delivery models mathematically unsustainable, and (3) constitutional frameworks that preclude service retrenchment as a politically or legally viable alternative. In settings where one or more conditions is absent, environmental factors can be expected to revert to the conventional, facilitative role posited in the traditional TOE literature.

- P1.* (For future comparative testing). In public organisations subject to binding welfare mandates, environmental imperatives function as the primary catalysts of AI adoption, with technological and organisational factors determining implementation pathways within legally bounded frameworks rather than the initial decision to adopt.

5.2 Organisational metabolism: definition, lineage and theoretical positioning

Definition. We define *organisational metabolism* as a mid-range theoretical construct (not a metaphor, not a sensitising concept and not a pure capability) denoting a sustained

organisational capacity for continuous technological absorption, adaptive reconfiguration, and normative calibration to constitutional values across political and regulatory cycles. It has three constitutive sub-capacities: (1) *absorptive continuity* – the persistent, rather than episodic, assimilation of technological knowledge across generations of technology; (2) *adaptive reconfiguration* – the iterative realignment of processes, structures and data pipelines as new technologies are integrated and (3) *normative calibration* – the systematic evaluation of technological outputs against constitutional values, ensuring that efficiency gains do not come at the cost of fundamental rights. Empirically, these three sub-capacities correspond to the INPS patterns documented in [Sections 4.3](#) and [4.4](#) (IT-policy integration, directive-based binding obligations and augmentation-centric automation).

Theoretical lineage. Organisational metabolism can be positioned in the dynamic capabilities tradition ([Teece, 2007](#)) as a sector-specific refinement of the sense–seize–reconfigure triad, but with two substantive modifications. First, its temporal horizon is extended from the medium-term strategic cycle of competitive repositioning to the *decade-scale* horizon of constitutional mandates and political-regulatory cycles. Second, its evaluative metric is shifted from competitive advantage to *normative alignment with public-law values*. The construct therefore draws from, but does not reduce to, absorptive capacity ([Cohen and Levinthal, 1990](#)), dynamic capabilities ([Teece, 2007](#)), organisational routines ([Feldman and Pentland, 2003](#)) and neo-institutional theory ([DiMaggio and Powell, 1983](#)).

Differentiation from proximate constructs. [Table 4](#) makes the distinction precise.

What the construct adds. Neither absorptive capacity nor dynamic capabilities were developed to account for settings in which technological absorption must be reconciled with constitutional norms and in which the relevant horizon is multi-decadal and multi-administration. Organisational metabolism fills this gap. Its empirical payoff is twofold: first, it allows us to make sense of the INPS pattern of *low episodic but high continuous adoption* that absorptive capacity alone predicts only partially; second, it supplies the analytical machinery needed to read [Table 3](#) findings as coherent adoption logic rather than an inventory of factors. We make no claim that the construct is applicable outside the boundary conditions articulated in [Section 6.2](#); its comparative validation is a task for future research.

P2. (For subsequent comparative empirical testing). Sustained AI adoption in constitutionally anchored public organisations is contingent upon organisational metabolism, understood as a dynamic capacity for absorptive continuity, adaptive reconfiguration and normative calibration that enables ongoing adaptation across prolonged institutional timescales and political and regulatory cycles while preserving institutional accountability and democratic legitimacy.

Table 4. Organisational metabolism vs proximate constructs

Construct	Core unit of analysis	Temporal horizon	Evaluative metric
Absorptive capacity (Cohen and Levinthal 1990)	Knowledge assimilation	Episodic	Competitive advantage
Dynamic capabilities (Teece 2007)	Sense–seize–reconfigure	Medium-term strategic cycles	Competitive repositioning
Organisational routines (Feldman and Pentland 2003)	Recurrent interaction patterns	Short- to medium-term	Efficiency, stability and change
Organisational metabolism (this paper)	Continuous absorption + reconfiguration + normative calibration	Decade-scale, across political cycles	Constitutional accountability and democratic legitimacy

5.3 Augmentation-centric automation

Raisch and Krakowski (2021) identify an automation–augmentation paradox whereby organisations experience a structural tension between efficiency gains derived from substituting human labour and capability gains obtained by augmenting human judgement. Empirical work in the private sector indicates that cost pressures frequently resolve this tension in favour of substitution-oriented automation. Public administration scholarship argues that public values should normatively shape technological choices in ways that diverge from profit-maximisation logics (Bannister and Connolly, 2014; Bullock *et al.*, 2020). Grimmelikhuijsen and Meijer (2022) highlight legitimacy requirements for algorithmic decisions affecting citizens’ rights. Yet, how welfare institutions actually resolve this paradox in practice remains empirically underexplored.

The INPS case illustrates how constitutional constraints can resolve the automation–augmentation paradox decisively in favour of augmentation. Across all documented application domains, AI operates as a complement and amplifier of professional expertise rather than a functional substitute. Fraud-detection algorithms flag suspicious patterns while inspectors retain responsibility for investigation and decision-making; chatbots address standardised queries, while human personnel handle complex cases that demand empathy, context and discretion; document classification automatically routes correspondence, while administrators conduct substantive review.

- P3. In constitutionally anchored public organisations, AI adoption is characterised by augmentation-centric automation, manifested in workforce redeployment toward higher-value, citizen-facing tasks rather than workforce downsizing, producing implementation trajectories systematically distinct from those observed in profit-driven organisational contexts.

6. Conclusion

6.1 Contributions

This study advances the discourse on AI adoption in public administration through an in-depth analysis of a large European welfare institution, yielding three theoretical contributions. First, in constitutionally anchored public-sector organisations, environmental factors may operate as initiating conditions for AI adoption rather than merely as contextual moderators; this suggests a recalibration of TOE’s relative weighting in such settings. Second, we introduce *organisational metabolism* as a mid-range construct that captures how such institutions continuously process and assimilate technological change while preserving accountability; we differentiate it explicitly from absorptive capacity, dynamic capabilities and organisational routines (Table 4). Third, we document *augmentation-centric automation* as the modal adoption pattern under constitutional constraint, with implications for how “AI success” is operationalised in public-sector contexts.

From a practitioner standpoint, four actionable principles follow from the INPS experience: (1) articulate the strategic role of AI explicitly within the framework of constitutional values; (2) invest in long-term IT–policy integration as a form of sustained organisational development; (3) assess success through a balanced set of indicators capturing both operational efficiency and distributive and procedural equity and (4) calibrate governance intensity to the magnitude of potential consequences, with particularly stringent human oversight for AI-supported decisions affecting individual rights.

6.2 Scope conditions and limitations

Scope conditions. The propositions derived from the INPS case are expected to apply primarily to organisations that share four characteristics: (1) constitutionally or legally mandated welfare obligations; (2) pronounced external pressure, particularly demographic

pressure; (3) a mature information-technology infrastructure with long-standing IT-policy integration and (4) large organisational scale associated with extensive population coverage. Many European social institutions exhibit similar characteristics; smaller administrative entities and organisations without binding welfare mandates are expected to exhibit different adoption dynamics.

Limitations. INPS's extreme institutional features make it a theoretically productive outlier but constrain external validity: the propositions advanced here are analytical conjectures, not empirically established generalisations, and demand systematic comparison with structurally similar organisations. The exclusion of generative AI is both a scope choice – the organisation did not deploy generative models during the study period – and a substantive limitation. Our findings speak primarily to the adoption of *traditional, narrow AI*; generalisation to generative AI deployment is not warranted, both because of distinct governance challenges (notably hallucination, training-data provenance and uncontrolled output spaces) and because the normative calibration sub-capacity of organisational metabolism has not been tested against such systems. This boundary should inform both the transferability and the timeliness of our conclusions. The empirical material mainly reflects the views of senior leadership and the technical secretariat; perspectives of frontline staff, citizens and trade-union representatives are under-represented and would substantively enrich future research. The temporal scope (August–October 2025) covers a bounded phase in an ongoing process. Cross-national differences in constitutional frameworks and administrative law traditions require careful contextualisation rather than simple extrapolation. This analysis should therefore be read as an analytical baseline and starting point for the comparative research agenda it seeks to motivate.

Acknowledgments

The authors wish to express their sincere gratitude to the Director General of the Italian National Social Security Institute (INPS) and to all the staff members involved in this research for their availability, openness and valuable contribution throughout the data collection process. Their willingness to share insights, documentation and operational experience was essential to the development of this study. The views expressed in this paper are solely those of the authors and do not necessarily reflect the official position of INPS.

Supplementary material

The supplementary material for this article can be found online.

References

- Alon-Barkat, S. and Busuioc, M. (2023), "Human–AI interactions in public sector decision making: 'automation bias' and 'selective adherence' to algorithmic advice", *Journal of Public Administration Research and Theory*, Vol. 33 No. 1, pp. 153–169, doi: [10.1093/jopart/muac007](https://doi.org/10.1093/jopart/muac007).
- Baker, J. (2012), "The technology–organization–environment framework", in Dwivedi, Y.K., Wade, M.R. and Schneberger, S.L. (Eds), *Information Systems Theory: Explaining and Predicting Our Digital Society*, Springer, Vol. 28, pp. 231–245.
- Bannister, F. and Connolly, R. (2014), "ICT, public values and transformative government: a framework and programme for research", *Government Information Quarterly*, Vol. 31 No. 1, pp. 119–128, doi: [10.1016/j.giq.2013.06.002](https://doi.org/10.1016/j.giq.2013.06.002).
- Berryhill, J., Heang, K.K., Clogher, R. and McBride, K. (2019), "Hello, world: artificial intelligence and its use in the public sector", OECD Working Papers on Public Governance, No. 36. OECD Publishing.
- Bovens, M. and Zouridis, S. (2002), "From street-level to system-level bureaucracies: how information and communication technology is transforming administrative discretion and constitutional control", *Public Administration Review*, Vol. 62 No. 2, pp. 174–184, doi: [10.1111/0033-3352.00168](https://doi.org/10.1111/0033-3352.00168).

-
- Bozeman, B. and Bretschneider, S. (1986), "Public management information systems: theory and prescription", *Public Administration Review*, Vol. 46, pp. 475-487, doi: [10.2307/975569](https://doi.org/10.2307/975569).
- Bretschneider, S. and Mergel, I. (2011), "Technology and public management information systems: where have we been and where are we going", in Menzel, D.C. and White, H.J. (Eds), *The State of Public Administration: Issues, Challenges and Opportunities*, M. E. Sharpe, pp. 187-203.
- Bullock, J., Young, M.M. and Wang, Y.-F. (2020), "Artificial intelligence, bureaucratic form, and discretion in public service", *Information Polity*, Vol. 25 No. 4, pp. 491-506, doi: [10.3233/ip-200223](https://doi.org/10.3233/ip-200223).
- Chen, T., Gascó-Hernandez, M. and Esteve, M. (2024), "The adoption and implementation of artificial intelligence chatbots in public organizations: evidence from US state governments", *The American Review of Public Administration*, Vol. 54 No. 3, pp. 255-270, doi: [10.1177/02750740231200522](https://doi.org/10.1177/02750740231200522).
- Cohen, W.M. and Levinthal, D.A. (1990), "Absorptive capacity: a new perspective on learning and innovation", *Administrative Science Quarterly*, Vol. 35 No. 1, pp. 128-152, doi: [10.2307/2393553](https://doi.org/10.2307/2393553).
- Cordella, A. and Iannacci, F. (2010), "Information systems in the public sector: the e-government enactment framework", *Journal of Strategic Information Systems*, Vol. 19 No. 1, pp. 52-66, doi: [10.1016/j.jsis.2010.01.001](https://doi.org/10.1016/j.jsis.2010.01.001).
- Damanpour, F. (1991), "Organisational innovation: a meta-analysis of effects of determinants and moderators", *Academy of Management Journal*, Vol. 34 No. 3, pp. 555-590, doi: [10.2307/256406](https://doi.org/10.2307/256406).
- Denzin, N.K. (1978), *The Research Act: A Theoretical Introduction to Sociological Methods*, 2nd ed., McGraw-Hill, New York, NY.
- DiMaggio, P.J. and Powell, W.W. (1983), "The iron cage revisited: institutional isomorphism and collective rationality in organisational fields", *American Sociological Review*, Vol. 48 No. 2, pp. 147-160, doi: [10.2307/2095101](https://doi.org/10.2307/2095101).
- Dubois, A. and Gadde, L.-E. (2002), "Systematic combining: an abductive approach to case research", *Journal of Business Research*, Vol. 55 No. 7, pp. 553-560, doi: [10.1016/s0148-2963\(00\)00195-8](https://doi.org/10.1016/s0148-2963(00)00195-8).
- Dwivedi, Y.K., Hughes, L., Ismagilova, E., Aarts, G., Coombs, C., Crick, T., Duan, Y., Dwivedi, R., Edwards, J., Eirug, A., Galanos, V., Ilavarasan, P.V., Janssen, M., Jones, P., Kar, A.K., Kizgin, H., Kronemann, B., Lal, B., Lucini, B., Medaglia, R., Le Meunier-FitzHugh, K., Le Meunier-FitzHugh, L.C., Misra, S., Mogaji, E., Sharma, S.K., Singh, J.B., Raghavan, V., Raman, R., Rana, N.P., Samothrakakis, S., Spencer, J., Tamilmani, K., Tubadji, A., Walton, P. and Williams, M.D. (2021), "Artificial Intelligence (AI): multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy", *International Journal of Information Management*, Vol. 57, 101994, doi: [10.1016/j.ijinfomgt.2019.08.002](https://doi.org/10.1016/j.ijinfomgt.2019.08.002).
- Eubanks, V. (2018), *Automating Inequality: How High-Tech Tools Profile, Police and Punish the Poor*, Martin's Press, St.
- Feldman, M.S. and Pentland, B.T. (2003), "Reconceptualizing organizational routines as a source of flexibility and change", *Administrative Science Quarterly*, Vol. 48 No. 1, pp. 94-118, doi: [10.2307/3556620](https://doi.org/10.2307/3556620).
- Flyvbjerg, B. (2006), "Five misunderstandings about case-study research", *Qualitative Inquiry*, Vol. 12 No. 2, pp. 219-245, doi: [10.1177/1077800405284363](https://doi.org/10.1177/1077800405284363).
- Fountain, J.E. (2001), *Building the Virtual State: Information Technology and Institutional Change*, Brookings Institution Press, Washington, DC.
- Grimmelikhuijsen, S. and Meijer, A. (2022), "Legitimacy of algorithmic decision-making: six threats and the need for a calibrated institutional response", *Perspectives on Public Management and Governance*, Vol. 5 No. 3, pp. 232-242, doi: [10.1093/ppmgov/gvac008](https://doi.org/10.1093/ppmgov/gvac008).
- Guenduez, A.A. and Mettler, T. (2023), "Strategically constructed narratives on artificial intelligence", *Government Information Quarterly*, Vol. 40 No. 1, 101719, doi: [10.1016/j.giq.2022.101719](https://doi.org/10.1016/j.giq.2022.101719).
-

- Hjaltalín, I.T. and Sigurdarson, H.T. (2024), "The strategic use of AI in the public sector: a public values analysis of national AI strategies", *Government Information Quarterly*, Vol. 41 No. 1, 101914, doi: [10.1016/j.giq.2024.101914](https://doi.org/10.1016/j.giq.2024.101914).
- Margetts, H. and Dorobantu, C. (2019), "Rethink government with AI", *Nature*, Vol. 568 No. 7751, pp. 163-165, doi: [10.1038/d41586-019-01099-5](https://doi.org/10.1038/d41586-019-01099-5).
- Mehr, H. (2017), *Artificial Intelligence for Citizen Services and Government*, Harvard Ash Center for Democratic Governance and Innovation, Cambridge, MA.
- Mergel, I., Edelmann, N. and Haug, N. (2019), "Defining digital transformation: results from expert interviews", *Government Information Quarterly*, Vol. 36 No. 4, 101385, doi: [10.1016/j.giq.2019.06.002](https://doi.org/10.1016/j.giq.2019.06.002).
- Mikalef, P., Lemmer, K., Schaefer, C., Ylinen, M., Fjørtoft, S.O., Torvatn, H.Y., Gupta, M. and Niehaves, B. (2022), "Enabling AI capabilities in government agencies: a study of determinants for European municipalities", *Government Information Quarterly*, Vol. 39 No. 4, 101596, doi: [10.1016/j.giq.2021.101596](https://doi.org/10.1016/j.giq.2021.101596).
- Neumann, O., Guirguis, K. and Steiner, R. (2024), "Exploring artificial intelligence adoption in public organizations: a comparative case study", *Public Management Review*, Vol. 26 No. 1, pp. 114-141, doi: [10.1080/14719037.2022.2048685](https://doi.org/10.1080/14719037.2022.2048685).
- OECD (2023), *Pensions at a Glance 2023: OECD and G20 Indicators*, OECD Publishing, Paris.
- OECD (2024), "Explanatory memorandum on the updated OECD definition of an AI system", OECD Artificial Intelligence Papers, No. 8. OECD Publishing.
- Oliveira, T. and Martins, M.F. (2011), "Literature review of information technology adoption models at firm level", *Electronic Journal of Information Systems Evaluation*, Vol. 14 No. 1, pp. 110-121.
- Pumplun, L., Tauchert, C. and Heidt, M. (2019), "A new organisational chassis for artificial intelligence: exploring organisational readiness factors", *Proceedings of the 27th European Conference on Information Systems (ECIS 2019)*, Stockholm.
- Rainey, H.G. (2014), *Understanding and Managing Public Organisations*, 5th ed., Jossey-Bass, San Francisco, CA.
- Rainey, H.G., Backoff, R.W. and Levine, C.H. (1976), "Comparing public and private organisations", *Public Administration Review*, Vol. 36 No. 2, pp. 233-244.
- Raisch, S. and Krakowski, S. (2021), "Artificial intelligence and management: the automation–augmentation paradox", *Academy of Management Review*, Vol. 46 No. 1, pp. 192-210, doi: [10.5465/amr.2018.0072](https://doi.org/10.5465/amr.2018.0072).
- Rogers, E.M. (2003), *Diffusion of Innovations*, 5th ed., Free Press, New York, NY.
- Spradley, J.P. (1980), *Participant Observation*, Holt, Rinehart and Winston, New York, NY.
- Teece, D.J. (2007), "Explicating dynamic capabilities: the nature and microfoundations of (sustainable) enterprise performance", *Strategic Management Journal*, Vol. 28 No. 13, pp. 1319-1350, doi: [10.1002/smj.640](https://doi.org/10.1002/smj.640).
- Tornatzky, L.G. and Fleischer, M. (1990), *The Processes of Technological Innovation*, Lexington Books, Lexington, MA.
- Veale, M. and Brass, I. (2019), "Administration by algorithm? Public law and the emerging use of machine learning in government", *Public Law*, Vol. 2019, pp. 30-55.
- Wang, Y.-F., Chen, Y.-C., Chien, S.-Y. and Wang, P.-J. (2024), "Citizens' trust in AI-enabled government systems", *Information Polity*, Vol. 29 No. 3, pp. 293-312, doi: [10.3233/ip-230065](https://doi.org/10.3233/ip-230065).
- Wirtz, B.W., Weyerer, J.C. and Geyer, C. (2019), "Artificial intelligence and the public sector: applications and challenges", *International Journal of Public Administration*, Vol. 42 No. 7, pp. 596-615, doi: [10.1080/01900692.2018.1498103](https://doi.org/10.1080/01900692.2018.1498103).
- Yin, R.K. (2018), *Case Study Research and Applications: Design and Methods*, 6th ed., Sage, Los Angeles, CA.
- Zuiderwijk, A., Chen, Y.-C. and Salem, F. (2021), "Implications of the use of artificial intelligence in public governance: a systematic literature review and a research agenda", *Government Information Quarterly*, Vol. 38 No. 3, 101577, doi: [10.1016/j.giq.2021.101577](https://doi.org/10.1016/j.giq.2021.101577).

Further reading

INPS (2024a), *Direttiva del Direttore Generale N. 8 dell'8 aprile 2024: Linee guida sull'implementazione di sistemi di Intelligenza Artificiale in INPS*, Istituto Nazionale della Previdenza Sociale, Rome.

INPS (2024b), *XXII Rapporto Annuale*, Istituto Nazionale della Previdenza Sociale, Rome.

INPS (2024c), *Piano Integrato di Attività e Organizzazione 2024-2026*, Istituto Nazionale della Previdenza Sociale, Rome.

Corresponding author

Luigina Paglieri can be contacted at: luigina.paglieri@uniroma2.it