

Mining Bureaucracy Debt: Process Intelligence for Evidence-Based Deregulation

Wil M.P. van der Aalst
0000-0002-0955-6940

Process and Data Science (PADS)
RWTH Aachen University, D-52074 Aachen, Germany
Email: wvdaalst@pads.rwth-aachen.de

Abstract—Administrative processes accumulate complexity over time. New forms, controls, approvals, reporting duties, and verification steps are introduced in response to incidents, audits, legal changes, political demands, or digitalization projects. These additions are usually individually defensible, but they are rarely reassessed or removed. Surprisingly, digitalization often accelerates this accumulation. Instead of eliminating obsolete procedures, organizations frequently reproduce them in software, add new digital controls, and retain parallel manual practices. The result is bureaucracy debt: the growing burden created by administrative requirements whose original purpose may have weakened, disappeared, or become achievable in a simpler way. This paper introduces bureaucracy mining, the systematic use of process mining and process intelligence to identify, quantify, compare, and reduce bureaucracy debt. We distinguish four layers at which debt may arise (legal, organizational, digital, and operational) and show how each translation from rules to execution may introduce additional burden. We outline longitudinal and comparative forms of bureaucracy mining, emphasize the need for object-centric event data, and argue that process observability and event-data portability are prerequisites for meaningful analysis.

I. INTRODUCTION

DIGITALIZATION is often introduced to make processes more efficient and faster. However, without systematic process-based evaluation, it may preserve old administrative obligations, introduce new ones, and make bureaucracy more deeply embedded in information systems. Work may also be shifted to end users (e.g., citizens or employees). Therefore, we advocate process mining as a tool for evidence-based deregulation.

Bureaucracy debt is analogous to technical debt, but it concerns administrative rules and processes rather than software. Technical debt arises when developers choose a quick or convenient solution that works in the short term but makes the software harder to maintain and change later. Bureaucracy debt arises when organizations add forms, approvals, controls, reporting obligations, and exceptions to solve immediate problems without considering their cumulative long-term effects. Surprisingly, bureaucracy debt may grow because of digitalization and the use of generative AI. This directly relates to the notion of “Process Mindlessness” introduced in [7]. In [7], we identify three main manifestations of process mindlessness: bloating, blurring, and blasting.

- *Bloating* occurs when generative AI makes produc-

ing text, reports, messages, software, and other content almost effortless. Organizations consequently generate more information than people can meaningfully process, creating additional reading, checking, coordination, and decision-making work.

- *Blurring* occurs when precise, structured, and verifiable information is replaced by plausible but probabilistic AI-generated language. Facts, assumptions, interpretations, and fabricated details become harder to distinguish, reducing transparency and accountability.
- *Blasting* occurs when digital technologies remove the natural constraints that previously limited the scale and speed of communication and action. AI-generated messages, requests, applications, reviews, and decisions can be produced and distributed in enormous volumes, shifting costs and overload to other participants in the process.

Bloating, blurring, and blasting as symptoms of “Process Mindlessness”: AI may successfully optimize a local task while making the end-to-end process noisier, less reliable, and more burdensome.

Another factor leading to bureaucracy debt is the tendency of people to systematically overlook *subtractive* solutions and add “more stuff” instead to solve problems. Adams et al. [8] demonstrate across multiple experiments that, when asked to improve an object, idea, or situation, people are more likely to consider *adding* components than *removing* components. Subtractive alternatives become even less likely under cognitive pressure or when people are not explicitly prompted to consider removal [8]. The authors identify institutional red tape as one domain in which this additive bias may have important consequences.

This tendency helps explain why bureaucracy accumulates. When an incident occurs, an audit identifies a weakness, or a new policy objective emerges, organizations typically respond by adding a control, approval, reporting obligation, form field, committee, or documentation requirement. Removing an outdated rule or simplifying an existing procedure is considered less readily. Each addition may appear reasonable in isolation, but the cumulative result is bureaucracy debt: growing administrative effort, complexity, and rigidity.

Digitalization can reinforce this additive bias. Rather than asking which activities can be eliminated, digital-transformation projects often add portals, workflow systems,

dashboards, data fields, AI checks, and reporting mechanisms on top of existing procedures. This connects directly to “Process Mindlessness”: attention is directed toward introducing new technologies and optimizing individual tasks, while insufficient attention is paid to the end-to-end process and the possibility that entire activities, controls, or information requirements should disappear.

The problem is real. A report by the Deutsche Bundesbank published in July 2026 [21] shows that administrative burden has become a substantial and rapidly increasing cost factor for German firms. Based on a survey of approximately 7,000 companies, the Bundesbank estimates that bureaucracy costs rose from around 5% of annual turnover in 2022 to around 7% in 2024. The burden is particularly severe for small firms because many compliance and documentation activities involve fixed costs that do not decrease proportionally with firm size. The strongest increases were reported in taxation and financial reporting, labor and employment regulation, environmental protection, and data protection. Manufacturing firms experienced an increase of more than 50% relative to turnover, potentially weakening the international competitiveness of German industry. Econometric estimates further suggest that increasing bureaucracy costs may have reduced German productivity growth by approximately 0.5 percentage point between 2022 and 2024. For small firms, a one-percentage-point increase in bureaucracy costs relative to turnover was associated with an average decline of around 3% in turnover per employee. *This example illustrates that bureaucracy debt is not merely a perceived inconvenience but a measurable economic burden.* Note that this refers only to German firms. Citizens are also confronted with ever-increasing red tape. This is visible to many and resulted in a new *Federal Ministry for Digital Transformation and Government Modernization* (BMDS). The ministry has a central coordinating role in identifying and removing bureaucratic barriers across the federal government and aims to reduce bureaucracy. However, without appropriate tools and a systematic approach, these efforts are unlikely to succeed.

The experience of the *U.S. Department of Government Efficiency* (DOGE) illustrates the risks of pursuing bureaucracy reduction through rapid, centrally imposed cuts without sufficiently transparent evidence, institutional safeguards, or end-to-end process analysis. DOGE was launched in January 2025 with ambitious promises to modernize federal technology, reduce expenditure, and eliminate waste. However, a substantial share of the initially announced savings could not be independently verified, contributing to the initiative’s termination. The problem was not the objective of reducing administrative burden, but the absence of a sufficiently systematic method for distinguishing unnecessary bureaucracy from capabilities required for lawful, secure, and effective government. Evidence-based deregulation requires identifying which activities create administrative burden, why they exist, which legitimate objectives they support, and what would happen if they were removed. *Process mining* [2], [5] can provide this evidence by comparing legal requirements, organizational

interpretations, digital implementations, and actual operational practices. Rather than following a disruptive “cut first, understand later” approach, we advocate *bureaucracy-debt mining* to target demonstrably redundant activities, evaluate proposed changes before implementation, and verify afterwards whether they actually reduce cost and processing time without weakening legality, service quality, data protection, or accountability.

Process mining [2], [5] can also counter the additive bias (adding more “stuff” to make things better [8]) by making subtraction an explicit redesign operation. Instead of asking only where activities can be accelerated or automated, evidence-based deregulation should systematically ask: Which steps can be removed? Which controls no longer influence outcomes? Which information requests duplicate data already available? Which approvals merely confirm earlier decisions? Process mining can identify such candidates and evaluate whether their removal reduces burden without undermining legality, fairness, quality, or risk control. In this way, bureaucracy-debt mining operationalizes the central insight of Adams et al.: meaningful improvement often requires not adding more, but deliberately searching for what can be taken away.

The remainder is organized as follows. Section II provides some background references on reducing bureaucracy and the role of digitalization. Section III introduces process mining as one of the few evidence-based approaches to tackle process inefficiencies. Section IV introduces a conceptual model describing four layers of bureaucracy debt. We elaborate on the “digitalization paradox” in Section V followed by ideas to implement forms of “bureaucracy mining” in Section VI. Section VII concludes this invited paper.

II. RELATED WORK

Research on administrative burden distinguishes learning, compliance, and psychological costs in citizen–state interactions and shows that such burdens can shape access to public services and policy outcomes [28], [24]. The related concept of red tape refers to rules and procedures that impose costs without sufficient contribution to their intended purpose [15]. This is closely related to bureaucracy debt, especially when initially useful requirements remain in place after their justification has weakened.

Digital-government research emphasizes reintegration, citizen orientation, and end-to-end service redesign [22]. However, empirical studies show that digitalization does not automatically reduce burden. E-government may simplify interactions, but it can also reproduce existing procedures, create parallel work, or shift administrative tasks to citizens [10], [26]. Proactive services may remove entire interactions, but require extensive data integration and raise concerns about transparency, privacy, and contestability [30].

“Rules as Code” and “RegTech” aim to make regulation machine-readable and compliance more efficient [27], [13]. Yet these approaches generally assume that existing requirements remain justified. They may reduce the cost of executing bureaucracy without reducing the underlying bureaucracy debt.

Process mining has become a standard tool to improve operational processes [2], [5]. It can be used to discover the actual processes, detect bottlenecks, analyze compliance, and diagnose problems. One of the main recent developments is the transition from case-centric to object-centric process mining, which provides a richer foundation for machine learning and AI. Process mining is now seen as a separate category of software [31]. Although process mining is widely used in manufacturing, automotive, logistics, finance, insurance, telecommunications, and retail, applications in public administration are limited. The reason is that there is less or no competition and there are no financial incentives. Yet the potential is huge [29], [32]. Object-centric process mining is particularly relevant because administrative processes involve interconnected citizens, applications, documents, payments, decisions, and organizational units [20]. However, existing work mainly improves the execution of current processes and rarely asks whether specific rules, controls, or activities should be removed.

This paper addresses that gap by introducing bureaucracy debt as the accumulated burden created by administrative requirements that persist after their value has declined. Process mining provides the operational evidence needed to identify such requirements and support evidence-based deregulation.

III. PROCESS MINING AND PROCESS INTELLIGENCE

AI can only improve organizations when it is grounded in how their processes actually work. Generative AI can produce text, code, or explanations; predictive AI can forecast delays or failures; and prescriptive AI can recommend actions. However, these capabilities remain isolated unless they are connected to operational data, process models, constraints, responsibilities, and organizational goals. This grounding is provided by *Process Intelligence* (PI), built around *Object-Centric Process Mining* (OCPM).

Traditional process mining [2], [5] follows one case type, such as one application, order, claim, or patient. Real processes are more interconnected: one administrative request may involve a citizen, several documents, departments, payments, decisions, and legal deadlines. OCPM [3], [4] represents these different objects and their relationships in one integrated view. This avoids oversimplifying reality and makes it possible to analyze problems that arise between organizational units rather than only within one isolated workflow.

OCPM starts from Object-Centric Event Data (OCED) [23], for example using the OCEL 2.0 standard [25]. Unlike a traditional event log, where every event must belong to exactly one case, an OCEL allows an event to be related to several objects of different types. Technically, an OCED describes:

- *events*, each with an activity (*event type*), timestamp, and attributes,
- *objects*, each belonging to an *object type* and carrying object attributes,
- *event-to-object relations* (E2O), indicating which objects participate in an event, and

- *object-to-object relations* (O2O), capturing structural or semantic relationships between objects.

Starting from event data – preferably OCED with E2O and O2O relations – process mining supports four main analytical capabilities:

- *Process discovery* automatically constructs a process model from recorded event data. Rather than starting from regulations, interviews, workshops, or manually created flowcharts, a discovery algorithm derives the observed behavior directly from the event log. The resulting model may take the form of a Directly-Follows Graph, Petri net, BPMN model, or object-centric process model.
- *Conformance checking* compares observed event data with a normative or expected process model. The reference model may represent legislation, internal regulations, service-level agreements, organizational policies, or a formally approved administrative procedure. The purpose is not limited to finding non-compliance by employees. Deviations may also indicate that the normative model is unrealistic, outdated, internally inconsistent, or unnecessarily bureaucratic. For example, regulations may require one eligibility check before granting a subsidy, while the event data show that essentially the same information is checked independently by three organizational units.
- *Performance analysis* enriches the discovered or normative process model with quantitative measurements using event data as *evidence*. Typical indicators include end-to-end throughput time, activity and handover durations, waiting and service times, queue lengths and work-in-progress, rework frequency, resource utilization, deadline violations, and processing costs. Object-centric performance analysis can localize delays at the intersections of objects and organizational units. For example, it may show that building-permit applications do not wait within the planning department itself, but become delayed when inspections, documents, payments, and external consultations must be synchronized.
- *Predictive analytics* uses historical process behavior to estimate future process states or outcomes of running process instances. Process mining supplies the process context needed to formulate meaningful machine-learning problems. Features may be derived from completed event sequences, participating objects, E2O and O2O relations, current workload, organizational unit, previous rework, and comparable historical executions. For example, every time an application reaches a particular administrative stage, a predictive model can estimate the expected delay using the application type, number of documents, related inspections, responsible unit, current queue, and history of comparable applications.

For a more comprehensive introduction to process mining, we refer to [2], [3], [4], [5]. It is important to note that the Pareto principle often applies to event data, i.e., a relatively small number of process variants often explain most of the




Fig. 1. Process model based on the top 3 variants covering 81% of the road traffic fines.




Fig. 2. Process model based on the top 10 variants covering 97% of the road traffic fines.

observed behavior. For example, 20% of the process variants may cover 80% of all cases. Figures 1, 2, and 3 illustrate this using the publicly available event data extracted from an information system supporting the management and handling of road traffic fines by a local police force in Italy [19]. The data set contains information about the handling of 145,800 road traffic fines.

Figure 1 is discovered by filtering fines based on the top 3 variants. The resulting Directly-Follows Graph (DFG) is simple but already covers 81% of all road traffic fines. Figure 2 is based on the top 10 variants, making the process model more complex, but still readable. Figure 3 is based on the full

data set. This example shows that the complexity is in the remaining 19% of cases not covered by Figure 1. Initiatives to tackle performance and compliance problems should not focus on Figure 1, but diagnose the full behavior represented by Figure 3. Starting from Figure 1 or Figure 2, one can create a normative process model describing the desired process. This normative process is then used to diagnose the deviations and their effect on throughput time, resource utilization, and outcomes.

The road traffic fine example is case-centric, i.e., we consider only one object type. In more complex scenarios, we need to consider multiple object types. OCPM is able to do this, thus providing a more faithful view of real processes by representing events in relation to multiple interacting objects, such as orders, items, deliveries, invoices, and customers, rather than forcing all behavior into a single case notion. This avoids distortions caused by traditional event-log flattening, supports multiple analytical perspectives from the same data, and enables better root-cause, conformance, and performance analysis across processes and organizational boundaries. It also reduces repeated data preparation by offering a reusable, system-independent event-data foundation. Overall, OCPM reveals the interconnected fabric of operational processes and therefore supports more accurate analysis and more effective process improvement.

Process mining has developed from an academic research field into a mature software category [31] used for evidence-based process improvement across manufacturing, finance, healthcare, logistics, and the public sector. More than 50 vendors now offer process-mining software, with platforms such as Celonis supporting process discovery, conformance checking, performance analysis, predictive monitoring, and increasingly object-centric process mining [3], [4]. The field is now moving beyond retrospective dashboards toward continuously updated, actionable process intelligence that connects process mining with automation, optimization, and generative, predictive, and prescriptive AI.

IV. FOUR LAYERS OF BUREAUCRACY DEBT

Before discussing the way in which we can use process mining and process intelligence to tackle bureaucracy in public administration, we first analyze the typical “fabric” of bureaucracy using the conceptual model shown in Figure 4. Bureaucracy should not be viewed as inherently negative. At its best, bureaucracy provides stability, predictability, accountability, and equal treatment by ensuring that decisions follow transparent rules rather than personal discretion. Standardized procedures can protect citizens, preserve institutional knowledge, and enable large and complex organizations to function reliably. However, bureaucracy gets a negative connotation when these structures become an end in themselves. Excessive formalism, redundant controls, fragmented responsibilities, and outdated rules can slow down decision-making, create unnecessary administrative burdens, and shift attention away from the actual needs of citizens and organizations.




Fig. 3. Process model based on all 145,800 road traffic fines. Although 81% of the cases can be described by a simple process model (see Figure 1), the remaining 19% add complexity. These remaining 19% of the cases cannot be ignored because most of the performance and compliance problems reside there.

Bureaucracy debt is the accumulated gap between the administrative burden generated by a process and the minimum burden necessary to achieve its legitimate legal and public-policy objectives. Like technical debt, it creates recurring “interest” in the form of additional processing, coordination, maintenance, explanation, and workaround costs.

Laws and regulations generally specify objectives, rights, obligations, constraints, and safeguards. However, these legal requirements are translated several times before they become an operational process. Each translation may introduce additional forms, approvals, controls, data fields, documentation duties, and exceptions. As shown in Figure 4, we therefore distinguish four layers of bureaucracy: the legal layer, the organizational layer, the digital layer, and the operational layer.

A. Legal Layer

The *legal layer* consists of requirements that follow directly from legislation, regulations, court decisions, and other binding rules. It defines what an administrative organization must achieve or prevent, for example, verifying eligibility, protecting personal data, ensuring equal treatment, preventing fraud, or providing opportunities for appeal.

Not every legally required safeguard constitutes bureaucracy debt. Many legal obligations serve legitimate public values such as legality, fairness, accountability, transparency, and security. Bureaucracy debt arises when legal requirements remain in force although their original justification has weakened, when overlapping regulations impose equivalent obligations, or when rules prescribe unnecessarily detailed procedures instead of defining the outcome to be achieved.

Legal bureaucracy is typically the most difficult layer to change because removing it may require legislative or regulatory action. Nevertheless, process evidence can support

such changes by showing that particular requirements consume substantial resources while contributing little to their intended objectives.

As mentioned in Figure 4, laws may be outdated, overlapping, or overly prescriptive. In fact, there are many examples where laws are even conflicting. For example, privacy laws ensure the right to erase personal data, while tax and commercial law require companies and public authorities to retain particular records for specified periods. Laboratories may restrict access to prevent infections. Fire-safety rules, however, require clearly usable escape routes and emergency exits.

B. Organizational Layer

The *organizational layer* is created when public organizations interpret and operationalize legal requirements. Laws often leave considerable discretion concerning how objectives should be achieved. Ministries, municipalities, universities, hospitals, and other organizations translate them into internal policies, responsibilities, approval hierarchies, forms, checklists, committees, and control procedures.

This translation frequently adds requirements that are not explicitly demanded by law. Organizations may introduce extra approvals to distribute responsibility, duplicate checks because departments do not trust one another, request additional documents to reduce perceived risk, or process work sequentially because decision rights are fragmented. Such additions may be individually understandable but collectively create substantial administrative burden.

Organizational bureaucracy debt therefore reflects the difference between what the law requires and what the organization requires. It is often reinforced by unclear process ownership: each department optimizes its own responsibilities, while




Fig. 4. Four layers introducing bureaucracy debt. The goal is to translate laws and regulations into operational processes that are efficient and effective. However, each layer may introduce unintentional bureaucracy debt.

no actor evaluates the end-to-end consequences for citizens, employees, or other organizations. Public administrations tend to be risk-averse. This is illustrated by the common phrase “it is better to ask for forgiveness than permission” used by employees dealing with administrations that tend to focus on legal instead of management aspects.

C. Digital Layer

The *digital layer* consists of requirements embedded in information systems, workflow configurations, portals, forms, validation rules, access controls, interfaces, and automated decision mechanisms. Once organizational interpretations are encoded in software, they become more rigid and less visible. A mandatory field, system-enforced approval, or hard-coded sequence may appear technically necessary even when it has no direct legal basis.

Digitalization may also preserve obsolete requirements. New systems are frequently designed by reproducing existing forms and procedures rather than reconsidering their purpose. Moreover, digital controls may be added while corresponding manual controls remain in place. This produces parallel paper and digital processes, duplicate data entry, repeated verification, and manual reconciliation between incompatible systems.

Digital bureaucracy debt is particularly persistent because users may no longer know why a system requires a particular field or approval. The original policy decision may be undocumented, the responsible persons may have left, and changing the system may be considered expensive or risky. Consequently, *the software itself becomes a de facto source of bureaucracy debt*.

The digital layer may cause *blasting* as described in [7]. Digital technologies remove the natural limits that once constrained the number, speed, and reach of administrative actions. Once a requirement has been encoded in software, it can automatically generate requests, reminders, validations, notifications, escalations, reports, and control activities for thousands of process instances affecting many employees or citizens. *Even a minor and weakly justified requirement may therefore create substantial system-wide burden*. Automation does not merely preserve bureaucracy; it can amplify and distribute it at negligible marginal cost, while shifting the resulting effort to citizens, employees, suppliers, and other administrative units.

D. Operational Layer

The *operational layer* comprises the processes and practices that are actually executed. It includes formal activities as well as local habits, informal consultations, private spreadsheets, email exchanges, phone calls, copy-and-paste work, shadow databases, manual tracking lists, and workarounds. These practices often remain invisible in official process descriptions and information systems.

Operational bureaucracy may arise when employees compensate for unclear rules, fragmented responsibilities, missing system functionality, poor data quality, or mistrust between organizational units. For example, employees may maintain an additional spreadsheet because the official system does not provide sufficient overview, request documents by email because systems are not connected, or conduct informal pre-

approvals before initiating a formal workflow.

Another common manifestation is administrative *ping-pong*: applications, documents, or requests are repeatedly sent back and forth between employees, departments, or organizations because responsibilities are unclear, information is incomplete, or each unit performs its own checks. Such loops create additional handovers, waiting time, repeated explanations, and rework without necessarily improving the final outcome.

Some operational practices are sensible adaptations that allow organizations to function despite inadequate rules or systems. Others have simply become habits whose original purpose has disappeared. Bureaucracy debt at this layer is therefore not identical to non-compliance. A deviation from the official process may indicate unnecessary work, but it may also reveal that the official procedure is unrealistic or that the digital system does not support the actual task. Recurrent ping-pong behavior is particularly revealing because it often indicates that no organizational unit has end-to-end responsibility and that local problem solving merely shifts work to another participant in the process.

Figure 4 illustrates how the four layers connect. Each layer may introduce additional bureaucracy debt. Using process mining this can be made visible.

V. THE DIGITALIZATION PARADOX

Digitalization is generally expected to increase efficiency and reduce bureaucracy by automating routine activities, accelerating communication, reusing data, and replacing paper-based procedures. Surprisingly, however, digitalization (and increasingly AI) often creates more rather than less bureaucracy. Existing procedures are frequently transferred into software without first questioning whether the underlying forms, approvals, controls, and reporting duties are still necessary. New validation rules, dashboards, notifications, audit trails, security checks, and AI-based controls are then added on top. Meanwhile, the corresponding manual procedures are rarely removed. The result is not simplification, but a growing combination of digital and analogue bureaucracy, including duplicate data entry, parallel documentation, and repeated verification. This reinforces the bureaucracy debt already embedded in the legal, organizational, digital, and operational layers.

A central element of this paradox is the shifting of work. A process may appear more efficient from the perspective of the administrative organization because activities previously performed by trained employees are transferred to citizens, companies, patients, students, or other employees. Digital self-service requires users to interpret instructions, create accounts, enter data, upload and convert documents, select classifications, resolve validation errors, remember credentials, and monitor the status of their requests. The processing time by the people having administrative roles may decrease, while the total end-to-end effort increases. What is reported as automation is therefore often merely a redistribution of administrative work to actors who are less familiar with the process and receive no compensation for performing it.

Moreover, once procedures are hard-coded, flexibility decreases: exceptional cases become more difficult to handle. Consider, for example, the author's official full name "Willibrordus Martinus Pancratius van der Aalst". Information systems truncate the full name in different ways leading to problems when the process is digitized.

To illustrate the digitalization paradox, consider again Figures 1, 2, and 3 based on event data extracted from an information system supporting the handling of road traffic fines [19]. When digitizing processes the goal is often to avoid the complexity of Figure 3 and squeeze cases into a simpler workflow like the one shown in Figure 1 or Figure 2. This may work well for some of the process variants. However, there are often good reasons for variability. By not supporting exceptional cases in the digital system, flexibility is reduced leading to additional efforts to work around the system. It may be the case that 80% of the cases benefit from a fully digitized process, but that the remaining 20% of the cases generate additional work and frustration.

VI. BUREAUCRACY MINING

We propose to systematically use process mining to identify, quantify, and reduce bureaucracy debt. We refer to this as *bureaucracy mining*. From a technological perspective, bureaucracy mining does not require an entirely new class of tools. Process discovery, conformance checking, performance analysis, comparative process mining, and predictive analytics can all be used. What changes is the analytical objective. Rather than merely asking how an existing process can be executed faster, bureaucracy mining asks why activities, controls, information requests, and handovers exist, which objectives they support, and whether they can be removed or simplified.

Bureaucracy mining starts from the operational layer in Figure 4, i.e., the actual process executions. However, to diagnose where bureaucracy debt is added, the analyst should understand the other three layers. A control observed in the operational process may follow directly from legislation, but it may also result from a conservative organizational interpretation, a hard-coded system constraint, or a local habit without any current formal basis.

Administrative activities are rarely performed for only one isolated case. They connect citizens, applications, documents, permits, payments, decisions, organizational units, legal provisions, and information systems. These relationships are needed to detect repeated information requests, duplicated controls, administrative ping-pong, and work shifted between actors. Bureaucracy mining should therefore measure not only the internal cost of an administrative organization, but also the burden imposed on citizens, companies, employees, and other public organizations.

A. Longitudinal Bureaucracy Mining

Bureaucracy debt accumulates over time. New requirements are often introduced after incidents, audits, political interventions, court decisions, system migrations, or organizational reorganizations. However, the resulting controls are rarely

reconsidered after the original circumstances have changed. *Longitudinal bureaucracy mining* compares process behavior across time periods to reveal this accumulation.

Event data can be divided into suitable periods, such as months, years, legislative periods, or phases before and after a system migration. Process-mining techniques can then identify changes in process structure and performance, including:

- newly introduced activities, controls, approvals, and data requests;
- increases in the number of handovers and organizational units involved;
- growth in approval depth, rework, and administrative ping-pong;
- increasing throughput times despite stable service demand;
- parallel manual and digital activities introduced during migration;
- temporary controls that have become permanent; and
- compensating workarounds created in response to inadequate systems.

This is related to the work on *concept-drift detection* in process mining [9], [14], [16].

A *bureaucracy debt index* could summarize the development of administrative burden over time. Such an index may combine the number of activities and handovers, approval depth, repeated data requests, rework frequency, waiting time, external user effort, and the proportion of controls that influence an outcome. The purpose would not be to reduce bureaucracy to one universal number. Instead, the index would provide a transparent and reproducible means of detecting whether a process is becoming more complicated without corresponding improvements in legality, quality, fairness, security, or accountability.

Longitudinal bureaucracy mining also supports the evaluation of deregulation initiatives. Measurements before and after a change can determine whether an abolished form, simplified approval, or integrated information system actually reduces total effort. This is important because removing one formal activity may merely cause new informal work elsewhere in the process.

B. Comparative Bureaucracy Mining

The same legal requirement is frequently implemented in different ways by different organizations. Municipalities, federal states, universities, hospitals, and administrative units may operate under the same legislation while requiring different forms, documents, controls, approvals, and processing steps. This variation creates a natural basis for *comparative bureaucracy mining*.

Comparisons can be made both across the four layers (cf. Figure 4) and across different organizations. Within one organization, process mining can compare: the activities directly required by legislation, the activities prescribed by internal policies, the activities enforced by information systems, and the activities actually performed in practice.

The differences between these representations reveal where bureaucracy debt may have been introduced. A legally required eligibility decision may, for example, have been translated into three organizational approvals, five mandatory system fields, and several informal pre-checks. Conversely, deviations in the operational process may reveal that an official requirement is routinely bypassed because it is impractical or adds no recognizable value.

Across organizations, comparative process mining can identify structurally different implementations of the same public service. Suppose several municipalities process the same type of building permit under the same legal framework. One municipality may use two approvals and reuse information from existing registers, whereas another may request the same information repeatedly and involve six organizational units. Differences in throughput time, workload, error rates, appeals, and outcomes can then be related to differences in process design.

An early example is the *Configurable Services for Local Governments* (CoSeLoG) project [17], [18], running from 2010 until 2014. The project involved ten Dutch municipalities and investigated how municipalities could share configurable process-support infrastructure while retaining necessary local differences. The underlying observation was that municipalities provide many comparable public services and operate under the same national legal framework, but nevertheless implement these services differently. CoSeLoG combined configurable process models, cloud-based process support, and cross-organizational process mining to identify commonalities and differences and to enable municipalities to learn from one another.

Various techniques for comparative process mining have been developed [1], [6], [11], [12]. Using dimensions such as municipality, organizational unit, case type, or information system, event data can be sliced and aggregated to compare discovered behavior, performance, and conformance. This is particularly relevant for bureaucracy mining: organizations operating under the same legislation may implement substantially different procedures, as demonstrated by the CoSeLoG project involving comparable municipal services. Such comparisons can reveal positive deviants and identify local approvals, controls, information requests, and handovers that create additional burden without measurably improving outcomes.

C. Free the Process: Observability and Portability

Longitudinal bureaucracy mining, comparative bureaucracy mining, and other forms of bureaucracy mining require high-quality event data with well-defined semantics. Activities, objects, organizational units, and outcomes need to be represented consistently so that processes can be compared across time, systems, and organizations. Without such semantic alignment, observed differences may reflect incompatible data definitions rather than genuine differences in bureaucracy. However, obtaining suitable event data is often difficult for three reasons.

First, enterprise systems may technically store the required data while *making extraction unnecessarily slow, complex, or costly*. Large software providers such as SAP offer extraction mechanisms, but these are often tightly coupled to proprietary data models, interfaces, and licensing arrangements. As a result, organizations may remain dependent on specific vendors and tools, even though the operational data concern their own processes. Bureaucracy mining therefore requires more than nominal data access. Organizations need timely, scalable, and system-independent access to event data, including stable identifiers, timestamps, activity semantics, and relations between relevant objects. Process observability and event-data portability should become explicit requirements in public procurement and system design.

Second, a substantial part of administrative work takes place *outside the official enterprise systems*. Employees use spreadsheets, emails, phone calls, informal consultations, copy-and-paste operations, shared folders, and manually maintained tracking lists. These activities often compensate for fragmented systems, missing functionality, unclear responsibilities, or poor data quality. Consequently, an analysis based only on the primary workflow system may present an incomplete and overly optimistic view of the process. Bureaucracy mining must therefore be able to incorporate data from multiple systems and, where appropriate, reconstruct recurring patterns of shadow work. The objective is not to monitor individual employees, but to reveal structural dependencies, duplicated effort, administrative ping-pong, and workarounds that remain invisible in official process descriptions.

Third, *privacy considerations are frequently invoked to restrict access to process data and thereby prevent transparency*. Privacy protection is essential, particularly when event data concern citizens, patients, employees, or sensitive administrative decisions. However, privacy should not become a blanket argument against evidence-based process analysis. Many relevant questions can be answered using pseudonymization, aggregation, access controls, data minimization, secure processing environments, and purpose-specific governance. Moreover, process mining often focuses on structural patterns, delays, handovers, repeated checks, and duplicated work rather than on the identity of individual persons. The challenge is therefore to design privacy-preserving forms of transparency instead of treating privacy and process improvement as mutually exclusive.

Freeing the process means ensuring that operational behavior can be observed independently of a particular software vendor, organizational silo, or informal communication channel. This requires open interfaces, semantically meaningful event data, traceable links between the legal, organizational, digital, and operational layers, and governance arrangements that enable responsible analysis. Without such access, bureaucracy remains hidden in proprietary systems, fragmented work practices, and overly restrictive interpretations of privacy. Without transparency, there can be no systematic identification or reduction of bureaucracy debt.

D. Evidence-Based Deregulation

The ultimate objective of bureaucracy mining is not process analysis, but evidence-based deregulation. Administrative activities can be divided conceptually into four categories:

- *value-creating work*, such as assessing an application and producing a correct decision;
- *necessary safeguards*, such as fraud prevention, data-protection checks, legal review, and opportunities for appeal;
- *administrative friction*, such as repeated data entry, avoidable handovers, and waiting caused by fragmented responsibilities; and
- *pure bureaucratic waste*, such as approvals and controls that consume resources but have no measurable influence on outcomes.

The boundaries between these categories are context dependent. A control may appear wasteful because it rarely changes a decision, while its presence may have a preventive effect. Likewise, an activity that adds no value for the administrative organization may provide transparency or procedural protection to the citizen. Bureaucracy mining should therefore not equate frequency or duration with value. It should relate activities to the legitimate objectives they are intended to support.

A central research challenge is to formalize administrative value added. An activity provides administrative value when it measurably contributes to outcomes such as correct decisions, compliance, equal treatment, fraud prevention, security, accountability, transparency, or successful service delivery. Process mining can then identify activities that increase effort and throughput time without observable contributions to these objectives.

Potential candidates for removal or redesign include:

- controls that almost never detect a problem or alter a decision;
- approvals that nearly always confirm an earlier recommendation;
- ping-pong loops involving rejected requests that are eventually accepted;
- information repeatedly requested although it is already available;
- equivalent checks performed independently by different units;
- digital and paper activities executed in parallel;
- manual workarounds surrounding supposedly digital processes;
- activities introduced during a system migration but never removed;
- controls introduced after an incident but never reassessed;
- increasing process complexity despite stable legal requirements; and
- system-enforced requirements that cannot be traced to legislation or an explicit organizational decision.

The previously introduced bureaucracy-debt index could be used as a continuous reminder to search for improvement

opportunities.

VII. RED TAPE NEVER DIES

Bureaucracy debt accumulates as legal requirements are translated into organizational rules, digital constraints, and operational workarounds. Digitalization and AI do not automatically reduce this debt. They may preserve obsolete procedures, introduce additional controls, shift work to citizens and employees, and amplify administrative activity at scale. Consequently, the gap between the promise of AI and administrative reality is widening: appointments may still take weeks, and simple procedures such as travel reimbursement may take months. This contrasts sharply with many commercial services, where products are delivered within days and responses are nearly immediate.

This paper introduced *bureaucracy mining*: the systematic use of process mining to identify, quantify, compare, and reduce unnecessary administrative burden. Longitudinal analysis reveals how bureaucracy debt accumulates over time, comparative analysis exposes avoidable differences between organizations, and object-centric process mining captures the interactions among applications, documents, decisions, payments, systems, and organizational units. Together, these capabilities provide an operational foundation for evidence-based deregulation rather than indiscriminate cuts.

The central challenge is not to automate bureaucracy, but to distinguish necessary safeguards from administrative friction and pure waste. This requires process observability, portable event data, measurable outcomes, and explicit links between the legal, organizational, digital, and operational layers. AI can accelerate existing procedures; process intelligence is needed to determine which procedures should exist at all. Red tape never dies by itself: process mining must first make it visible, measurable, and contestable.

GENERATIVE AI DISCLOSURE STATEMENT

During the preparation of this work the author used ChatGPT in order to improve the language and clarity of the manuscript. After using this tool/service, the author reviewed and edited the content as needed and takes full responsibility for the content of the publication.

REFERENCES

- [1] W.M.P. van der Aalst. How People Really (Like To) Work: Comparative Process Mining To Unravel Human Behavior. In S. Sauer, C. Bogdan, P. Forbrig, R. Bernhaupt, and M. Winckler, editors, *International Conference on Human-Centered Software Engineering (HCSE 2014)*, volume 8742 of *Lecture Notes in Computer Science*, pages 317–321. Springer-Verlag, Berlin, 2014.
- [2] W.M.P. van der Aalst. *Process Mining: Data Science in Action*. Springer-Verlag, Berlin, 2016.
- [3] W.M.P. van der Aalst. Object-Centric Process Mining: An Introduction. In A. Cerone, editor, *Formal Methods for an Informal World, ICTAC 2021 Summer School*, volume 13490 of *Lecture Notes in Computer Science*, pages 73–105. Springer-Verlag, Berlin, 2023.
- [4] W.M.P. van der Aalst. Object-Centric Process Mining: Unraveling the Fabric of Real Processes. *Mathematics*, 11(12):2691, 2023.
- [5] W.M.P. van der Aalst and J. Carmona, editors. *Process Mining Handbook*, volume 448 of *Lecture Notes in Business Information Processing*. Springer-Verlag, Berlin, 2022.
- [6] W.M.P. van der Aalst, S. Guo, and P. Gorissen. Comparative Process Mining in Education: An Approach Based on Process Cubes. In P. Ceravolo, R. Accorsi, and P. Cudre-Mauroux, editors, *IFIP International Symposium on Data-Driven Process Discovery and Analysis (SIMPDA 2013)*, volume 203 of *Lecture Notes in Business Information Processing*, pages 110–135. Springer-Verlag, Berlin, 2015.
- [7] W.M.P. van der Aalst, H.A. Reijers, O. Hinz, and J. Pfeiffer. Process Mindlessness: When we Lose Sight of What AI is Supposed to Improve. *Business and Information Systems Engineering*, 67(6):771–775, 2025.
- [8] G.S. Adams, B.A. Converse, A.H. Hales, and L.E. Klotz. People Systematically Overlook Subtractive Changes. *Nature*, 592:258–261, 2021.
- [9] J.N. Adams, S.J. van Zelst, T. Rose, and W.M.P. van der Aalst. Explainable Concept Drift in Process Mining. *Information Systems*, 114:102177, 2023.
- [10] R. Arendsen, O. Peters, M. ter Hedde, and J. van Dijk. Does E-Government Reduce the Administrative Burden of Businesses? An Assessment of Business-to-Government Systems Usage in the Netherlands. *Government Information Quarterly*, 31(1):160–169, 2014.
- [11] A. Bolt, M. de Leoni, and W.M.P. van der Aalst. A Visual Approach to Spot Statistically-Significant Differences in Event Logs Based on Process Metrics. In S. Nurcan, P. Soffer, M. Bajec, and J. Eder, editors, *International Conference on Advanced Information Systems Engineering (Caise 2016)*, volume 9694 of *Lecture Notes in Computer Science*, pages 151–166. Springer-Verlag, Berlin, 2016.
- [12] A. Bolt, M. de Leoni, and W.M.P. van der Aalst. Process Variant Comparison: Using Event Logs to Detect Differences in Behavior and Business Rules. *Information Systems*, 74(1):53–66, 2018.
- [13] M. Bolton and M. Mintrom. RegTech and Creating Public Value: Opportunities and Challenges. *Policy Design and Practice*, 6(3):266–282, 2023.
- [14] R.P. Jagadeesh Chandra Bose, W.M.P. van der Aalst, I. Zliobaite, and M. Pechenizkiy. Dealing With Concept Drifts in Process Mining. *IEEE Transactions on Neural Networks and Learning Systems*, 25(1):154–171, 2014.
- [15] B. Bozeman. A Theory of Government Red Tape. *Journal of Public Administration Research and Theory*, 3(3):273–304, 1993.
- [16] T. Brockhoff, M.S. Uysal, and W.M.P. van der Aalst. Time-aware Concept Drift Detection Using the Earth Mover's Distance. In B.F. van Dongen, M. Montali, and M. Wynn, editors, *International Conference on Process Mining (ICPM 2020)*, pages 33–40. IEEE Computer Society, 2020.
- [17] J.C.A.M. Buijs, B.F. van Dongen, and W.M.P. van der Aalst. Towards Cross-Organizational Process Mining in Collections of Process Models and their Executions. In F. Daniel, K. Barkaoui, and S. Dustdar, editors, *Business Process Management Workshops, International Workshop on Process Model Collections (PMC 2011)*, volume 100 of *Lecture Notes in Business Information Processing*, pages 2–13. Springer-Verlag, Berlin, 2012.
- [18] CoSeLoG. Configurable Services for Local Governments (CoSeLoG) Project Home Page. www.win.tue.nl/coselog.
- [19] M. de Leoni and F. Mannhardt. Road Traffic Fine Management Process (4TU.ResearchData). <https://doi.org/10.4121/uuid:270fd440-1057-4fb9-89a9-b699b47990f5>, 2015.
- [20] E. de Moura Figueiredo and A. Jalali. Object-Centric Process Mining for Public Sector Transformation. In *Joint Proceedings of the BIR 2025 Workshops and Doctoral Consortium, Co-Located with the 24th International Conference on Perspectives in Business Informatics Research*, volume 4034 of *CEUR Workshop Proceedings*, pages 184–197. CEUR-WS.org, 2025.
- [21] Deutsche Bundesbank. Monatsbericht Juli 2026, 78 Jahrgang, Nr. 7, Deutsche Bundesbank. <https://publikationen.bundesbank.de/publikationen-de/berichte-studien/monatsberichte/monatsbericht-juli-2026-1000490>, 2026.
- [22] P. Dunleavy, H. Margetts, S. Bastow, and J. Tinkler. New Public Management Is Dead: Long Live Digital-Era Governance. *Journal of Public Administration Research and Theory*, 16(3):467–494, 2006.
- [23] D. Fahland, M. Montali, J. Leberher, W.M.P. van der Aalst, M. van Asseldonk, P. Blank, L. Bosmans, M. Brenscheidt, C. Di Ciccio, A. Delgado, D. Calegari, J. Peeperkorn, E. Verbeek, L. Vugs, and M.T. Wynn. Towards a Simple and Extensible Standard for Object-Centric Event Data (OCED): Core Model, Design Space, and Lessons Learned. *Computing Research Repository (CoRR) in arXiv*, abs/2410.14495, 2024.

- [24] P. Herd and D.P. Moynihan. *Administrative Burden: Policymaking by Other Means*. Russell Sage Foundation, New York, 2018.
- [25] I. Koren, J.N. Adams, A. Berti, and W.M.P. van der Aalst. OCEL 2.0 Resources - www.ocel-standard.org. In J.M. van der Werf, C. Cabanillas, F. Leotta, and L. Genga, editors, *Doctoral Consortium and Demo Track 2023 at the International Conference on Process Mining 2023*, volume 3648 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2023.
- [26] C.Ø. Madsen, I. Lindgren, and U. Melin. The Accidental Caseworker—How Digital Self-Service Influences Citizens’ Administrative Burden. *Government Information Quarterly*, 39(1):101653, 2022.
- [27] J. Mohun and A. Roberts. Cracking the Code: Rulemaking for Humans and Machines. Technical Report 42, OECD Publishing, 2020.
- [28] D. Moynihan, P. Herd, and H. Harvey. Administrative Burden: Learning, Psychological, and Compliance Costs in Citizen–State Interactions. *Journal of Public Administration Research and Theory*, 25(1):43–69, 2015.
- [29] D. Rawiro, F.L. Gaol, S.H. Supangkat, and B. Ranti. Process Mining Applications in Government Sector: A Systematic Literature Review. In *Proceedings of the 5th European International Conference on Industrial Engineering and Operations Management*, pages 3110–3118, Rome, Italy, 2022. IEOM Society International.
- [30] H. Scholta and I. Lindgren. Proactivity in Digital Public Services: A Conceptual Analysis. *Government Information Quarterly*, 40(3):101832, 2023.
- [31] T. Srivastava, D. Sugden, and M. Kerremans. Magic Quadrant for Process Intelligence Platforms, Gartner Research Note G00840852. www.gartner.com, 2026.
- [32] J. Trottier, W. van Woensel, X. Wang, K. Mallur, N. El-Gharib, and D. Amyot. Using Process Mining to Improve Digital Service Delivery. *Computing Research Repository (CoRR) in arXiv*, abs/2409.05869, 2024.