

Algorithmic Incorporation: Reconciling SECP's eZfile with Constitutional Due-Process Guarantees

Muhammad Usman Subhani

usmansubhani333@gmail.com

Lecturer Department of Law, Mirpur University of Science and Technology Azad Kashmir, Pakistan.

Hazrat Usman

hazrat.usmanadvocate@gmail.com

Advocate High Court, Punjab Bar Council, Punjab, Pakistan.

Corresponding Author: *Muhammad Usman Subhani usmansubhani333@gmail.com

Received: 22-11-2025	Revised: 08-12-2025	Accepted: 22-12-2025	Published: 06-01-2026
-----------------------------	----------------------------	-----------------------------	------------------------------

ABSTRACT

Automated refusal of incorporation and reflexive freezing of bank accounts are frequently presented as unavoidable tolls of algorithmic celerity; however, our multistage experiment with the Securities and Exchange Commission of Pakistan's e-Zfile, the Competition Commission portal, and the AML engines of three tier-I banks demonstrates the contrary. It was shown that the insertion of a narrowly circumscribed notice-and-reply module consisting only of rule-trigger logging, machine-readable reasoning, and a seven-hour rebuttal window accords well with calculations based on the classical triad of audi alteram partem, reasoned decision making, and proportionality. The error rates in registration, merger control, and suspicious activity detection decreased by 26.9 %, 48.3 %, and 14.2 %, respectively, while the ratio of demographic phenotypic classes converged toward parity ($\chi^2 = 5.62$; $P > 0.05$) and the share of irreversible actions dropped below 3 %. It is necessary to explain why this constitutional graft required only 200 lines of additional code and eleven-minute deployment cycles. The answer, it appears, is that due-process norms operate as design affordances, not external shackles. The findings indicate that legality and efficiency may be mutually reinforcing and allow us to conclude that "constitutional-by-design" should migrate from pilot demonstrations to statutory mandates across diverse regulatory domains, repositories, and organizational cultures already in production environments.

Keywords: Algorithmic decision-making; Constitutional due process; Administrative law automation; SECP e-Zfile; AML scoring engines; Procedural fairness.

INTRODUCTION

"Algorithmic regulation," once a shorthand for predictive policing in Silicon Valley techno-optimism, now defines the quotidian life-cycle of corporate personality in Pakistan: a private entrepreneur enters the Securities and Exchange Commission of Pakistan (SECP) eZfile portal, uploads a digital memorandum, and, within milliseconds, "an algorithm freezes an account or rejects a registration due to a data mismatch," silently eclipsing the audi alteram partem guarantee that undergirds administrative justice (Andrews et al., 2017, p. 2). Lawyers accustomed to oral hearings confront what Williams (2022) calls "the increasing prevalence of algorithmic decision-making ... by public authorities" and the attendant need for doctrinal recalibration (p. 1). Yet while Euro-American scholarship refines proportionality tests for machine-logic (Allars, 2024) and global policy coalitions celebrate "algorithmic accountability for the public sector" (Ada Lovelace Institute et al., 2021, p. 3), Pakistan's statutory edifice still presumes a human decision-maker who writes reasons in prose rather than code.

The disjunction is sequential. The Securities and Exchange Commission of Pakistan (SECP)'s eZfile, the Public Credit Registry's online dispute console, and commercial banks' AML-CFT scoring engines jointly govern the formation, persistence, and liquidity of economic actors; each system executes "regulation of and by algorithm" (Andrews et al., 2017, p. 2) through rule sets that were neither tabled in Parliament nor gazetted in the manner contemplated by Article 89 of the Constitution. When a flagged beneficial-owner field or an anomalous CIF "the natural person on whose behalf a transaction is being conducted" (State Bank of Pakistan, 2022, Definitions section) triggers an automated freeze, the individual first learns of the adverse act not through a show-cause notice but through the sudden inability to operate a bank account. The State Bank's consumer booklet enumerates the "Rights and Responsibilities of Credit-Bureaus' Consumers," but its procedural roadmap presumes an ex-post remedy rather than an ex-ante hearing (State Bank of Pakistan, 2021, p. 3). Shah, Ali, and Tahir (2025) observe that while "AI-supported technologies can modernise legal procedures," they also risk "structural delays" structural delays by entrenching opacity (p. 88). becauseative lacuna is thus twofold: first, classic administrative-law canons notice, evidence disclosure, reasoned decision, and proportionality are frustrated by speed and secrecy; second, judicial review itself struggles because, as Green (2024) notes in the policing context, "courts cannot interrogate a decision whose logic remains undisclosed" (p. 4).

Therefore, this article advances an original jurisprudential and empirical inquiry, neither doctrinal commentary nor sectoral survey, into what we term *algorithmic incorporation*: the fusion of corporate-law registration workflows with machine-executed public powers. Our research question asks how the foundational principles of Pakistani administrative law, especially the right to be heard and the duty to give reasons, can be operationalized inside the code so that automated refusals or freezes comply ex ante, not remedially ex post, with constitutional due-process guarantees. The stakes are evident. In the fiscal years 2023–24, 86 percent of new companies filed entirely through eZfile; parallel AML engines screened 312 million transactions. Thus, in the Weberian sense, each line of code substitutes for a human bureaucrat, yet remains outside the purview of the "structured proportionality that Williams (2022) argues must now migrate from paper to algorithm (p. 7).

We proceed on the basis of two empirical premises. First, algorithmic systems already embed de facto discretion. For example, eZfile's name-clearance module implements a phonetic-similarity threshold calibrated by private vendors, whereas SBP-mandated AML filters assign weighted risk scores based on proprietary typologies. As the Ada Lovelace report reminds policymakers, "transparency of input data and model logic is a pre-condition for meaningful accountability" (Ada Lovelace Institute et al., 2021, p. 36). Second, Pakistani constitutional jurisprudence supplies doctrinal tools, Article 10A's right to fair trial, and the *Ms. Shehla Zia* proportionality trilogy that can be *translated* into computational safeguards. Ranjha (2025) frames this translation project as designing "accountability triggers" that fire whenever an algorithmic rule deprives a citizen of a statutory benefit (p. 3). Our contribution is to specify, prototype, and empirically evaluate such triggers in an SECP's live environment.

Methodologically, multilayered audits were adopted. We decapitalize the black box by replicating eZfile's decision tree, then subject the reconstructed model to the "meaningful information" standard articulated under the GDPR and recently adapted by Williams (2022, p. 9) for common-law proportionality analysis. We then embed a *reason API* as a machine-readable justification string at each rejection node and assess its sufficiency against Allars' (2024) typology of reasons in automated administration (pp. 6–7). Parallel simulations on an anonymized AML data gauge whether proportionality thresholds can be adjusted ex ante to mitigate type (over blocking) and type-II (under-detection) errors without eroding constitutional minima. Throughout, we heed Butt's (2024) warning that "natural justice must not be sacrificed on the altar of efficiency" (p. 10).

Intellectual payoff transcends local compliance. By demonstrating that *constitutional code* can be compiled into existing regulatory stacks without debilitating latency, we contest the fatalism that Nordell (2021) attributes to the “deterritorialization of law in cyberspace” (p. 2). Professionally designed algorithmic incorporation does not annul administrative law norms; it materializes them in silicon. In this sense, our project aligns with global calls for “regulation and control of algorithmic codes” as “a necessity of our times” (Rott-Pietrzyk et al., 2025, p. 101) while delivering a doctrinally grounded, jurisdiction-specific blueprint.

Thus, the pages that follow move from conceptual mapping to technical implementations. We first reconstructed the legal anatomy of eZfile, product-channel-customer (PCC) scoring, and AML engines, engineered due-process modules at decisive junctures, and empirically evaluated the constitutional performance of the augmented systems. In so doing, we respond to the lacuna that Andrews et al. (2017) diagnose the gulf between “regulation of algorithms and regulation by algorithms” (p. 2) and offer the first Pakistan-focused, code-level solution to reconcile automated governance with the right to be heard.

Foundational Principles of Pakistani Administrative Due Process

Under Pakistani public law orthodoxy, “every administrative act must speak through reasons that are intelligible to those whom it binds (Williams, 2022, p. 470). Yet the constitutional promise of *audi alteram partem* travels across a terrain shaped as much by statutory text as by the “underlying structural values of fairness, transparency and proportionality” that Lord Sales calls the “north stars of judicial review” (Sales, 2024, p. 2). In Pakistan, these values emerged piecemeal: early High-Court jurisprudence read Article 4 of the 1973 Constitution as a freestanding right to be treated “in accordance with law,” while later benches anchored procedural fairness in Article 10A’s “fair trial and due process” clause. Butt (2024) crystallises the trajectory: administrative power moved “from bureaucracy to black box,” yet the courts insisted that even silent algorithms remain answerable to the same “reason-giving discipline that legitimates discretion” (p. 10).

This discipline comprises of three interlocking duties. First, notice: an affected person must understand the case that he/she must meet. The State Bank’s *Rights and Responsibilities of Credit-Bureaus’ Consumers* booklet formalizes this by obliging credit bureaux to furnish “all data underlying an adverse score” upon request (State Bank of Pakistan, 2021, p. 3). Second, opportunity to be heard of classical Pakistani cases requires a *real* chance to rebut, not a post hoc appeal, after a freeze has already been bitten. Ajuzieogu (2025) warns that delayed hearings transform due process into “an empty promise, effectually posthumous” when the decision is executed in real time (p. 5). Third, reasons: the administrator must reveal a logical bridge between facts and results. Allars (2024) calls this the “grammar of justification” and shows that, even for machine-generated outcomes, a court will probe whether the chain of inference was “coherent, transparent and bounded by the statute” (p. 1150).

These duties become constitutionally potent, when coupled with proportionality. Pakistani benches, borrowing from the German doctrine via the Commonwealth precedent, test invasions of rights against a four-step ladder: suitability, necessity, minimal impairment, and overall balancing. Williams (2022) argues that proportionality “already contains the scaffolding for algorithmic accountability” because each rung demands a showing of why a particular threshold, weight, or risk score is justified (p. 476). Ranjha (2025) applies this insight locally, recommending “accountability triggers” that fire when an e-governance rule “materially impairs a recognised civil or commercial liberty” (p. 3). The jurisprudential core is therefore settled: notice, hearing, reason, and proportionality are not aspirational add-ons, but constitutional predicates of legality.

However, the locus of adjudicative fact-finding shifted. In SECP's eZfile, the clearance of a company name or the validation of a beneficial-ownership field is now performed by what Andrews et al. (2017) describe as an "algorithm that applies regulatory rules in milliseconds, leaving no textual trace of deliberation" (p. 2). Green (2024) further notes that such systems "invert the traditional hearing model" because the decision precedes any human interaction (p. 4). The practical consequence is that constitutional principles, although rhetorically robust, risk becoming technologically toothless. The nub of our inquiry lies here: can we encode constitutional predicates *within* the decision pipeline so that rights are preserved contemporaneously with, rather than subsequent to, an algorithmic act?

To answer, we must first accept Butt's (2024) caution that "efficiency cannot eclipse legality" (p. 12). The proportionality doctrine already anticipates trade-offs: a measure that pursues a legitimate aim but is "manifestly excessive" fails the minimal-impairment test. For example, in AML screening, the Revised Regulations of Microfinance Banks mandate the bank to freeze an account once the risk score is greater than a predetermined limit; however, they also state that such a bank must notify immediately and offer a mechanism for quick review (State Bank of Pakistan, 2022, s. 15). Thus, a nascent *audi alteram partem* layer is inserted into the regulatory text, although it continues to work after the freeze. In our research, we suggest that the forward layer issues a pre-freeze notification and enables the customer to correct mismatched data before sanction bites.

Nordell (2021) puts the problem in cognitive terms: Code DE-territorializes law and law. aims at exteriorizing the code" (p. 2). The reterritorializing movements of Pakistani law are based on this doctrine of necessity. of error of jurisdiction: If an authority does not follow a compulsory requirement such as hearing or reason, it acts unauthorized, and the order is void *ab initio*. The Supreme Court's doctrine, as applied in electronic tax appeals, implies judicial eagerness to make extensions. classic due process justifications for the digital realm. As Allars (2024) demonstrated, Australian courts face the challenge of dealing with them. Related questions require an opaque machine-learning model to exhibit a traceable logic pathway, even for similar queries (p. 1158). The courts in Pakistan have similar remedial powers, and the missing link is operational in Pakistan.

This detail emerges from the comparative design proposals. The Ada Lovelace report urges governments to mandate "reason APIs" that surface the rule fired, the data points used, and the threshold crossed (Ada Lovelace Institute et al., 2021, p. 36). Williams (2022) suggests embedding proportionality checks directly into algorithmic logic, ensuring that each decision is accompanied by a self-audit string documenting suitability and necessity. Such checks resonate with Pakistan's *Shehla Zia* proportionality doctrine, which requires a demonstrable balance of rights and objectives. If a bank's AML engine blocks a widow's remittance because its surname resembles that on a sanctions list, a built-in minimal-impairment module could offer a manual override or escalate for rapid human review. In doctrinal terms, this would satisfy the requirement that "power must be exercised in a manner that least restricts the affected right" (Williams, 2022, p. 478).

The synthesis is unmistakable: Pakistani constitutional law provides a mature matrix of notice, hearing, reason, and proportionality; global scholarship supplies blueprints for translating those norms into code; and sectoral regulations signal a readiness to embed procedural hooks. The task now is to prototype, evaluate, and judicially validate an architecture in which the due process is not merely *retrospective redress* but *prospective design*. Our research proceeds from that premise, seeking to engineer constitutional guarantees directly into the decision-making kernels of eZfile, PCC scoring engines, and AML filters, thereby converting formal rights into executable safeguards.

Comparative Insights from Global Algorithmic-Governance Scholarship

“Algorithmic regulation,” Yeung’s phrase for rule-enforcement “in milliseconds, leaving no textual trace of deliberation” (Andrews et al., 2017, p. 2), now forms the common substrate of public administration from Tallinn to Toronto. Comparative scholars have, therefore, turned from describing codes to prescribing safeguards, asking how legal systems can extract “a grammar of justification” from opaque statistical engines (Allars, 2024, p. 1150). The most far-reaching proposal comes from the Ada Lovelace/AI Now/OGP study, which urges governments to embed “reason APIs” so that every automated decision releases the rule that fired, the data used, and the threshold crossed (Ada Lovelace Institute et al., 2021, p. 36). They argue that this design converts the classical right to reason into a “machine-readable due process,” a formulation that travels readily across jurisdictions because it aligns with the administrative law axiom that power must be both authorized and explained.

Williams (2022) extended this logic by welding the European proportionality doctrine to the code itself. He insists that “structured proportionality can be operationalized in the algorithmic environment” because each rung suitability, necessity, minimal impairment, and balance maps onto specific points in a model pipeline (p. 478). Suitability translates into feature relevance tests, necessity into sparsity constraints, minimal impairment into calibrated thresholds, and balancing into post-hoc fairness audits. The elegance of mapping lies in its promise that constitutional scrutiny does not require slow computation; the audit string rides alongside the prediction as a checksum travel with a data packet.

However, sophisticated abstractions seldom survive first contact with sectoral practices. In policing, Green (2024) shows that predictive-harm indices created for bail hearings in England “invert the traditional hearing model” by rendering judgment before a defendant speaks (p. 4). Courts, she writes, can no longer rely on cross-examination to surface error; they must instead require “traceable logic pathways” disclosed *ex ante*. Australian experience corroborates this shift. Allars (2024) catalogues cases in which tribunals demanded “coherent, transparent, and bounded” explanations, even from black-box systems (p. 1158), confirming that the right to reason binds regardless of the medium. However, the need for litigation signifies that the policies have lagged behind the code.

Regulators began to close this gap by treating algorithms as normative acts. Boix-Palop Devotees (2020) criticizes that a model in which the agent distributes legal impacts should be published within an agency, reviewed and corrected, as in any other regulation. (p. 78). The European Parliament’s Science-and-Technology Panel has a similar position, in that it cautions that a lack of transparency affects legitimacy. Since citizens have no right to challenge a rule, the contents of which are unfamiliar to them (European Parliamentary Research Service, 2019, p. 6). Russell, Judge, and Nitzberg (2024), in their turn, warn that code is not law” unless it is anchored to popular values, calling for a change in the efficiency measure to democratic oversight (p. 8). Nordell (2021) taps into the conflict between philosophy and says that the code de-territorializes law and law attempts to reterritorialize the code (p. 2).

However, visibility alone does not guarantee fairness. The Egmont Group (2025) reports that financial intelligence units armed with transparent rules still freeze accounts disproportionately when the data quality is poor (p. 12). AFME and EY (2021) reached a similar conclusion for transaction-monitoring algorithms: “rule-based systems may over block low-risk customers,” generating a compliance burden with scant crime-fighting benefits (p. 22). The FATF (2025) therefore urges regulators to adopt “tiered transparency,” releasing summary logic to the public while reserving sensitive thresholds for supervised auditors (p. 17). This calibrated approach aligns with Murray’s (2024) argument that regulation must “minimise collateral harm without neutering innovation” (p. 4).

The comparative literature also evaluates these remedies. The Commonwealth Ombudsman's *Automated Decision-Making: Better Practice Guide* (2025) lists 20 checkpoints ranging from data-source validation to human-in-the-loop escalation that agencies must certify before deployment (pp. 8–12). The Human Technology Institute submission to the Australian Attorney-General (2025) adds an “AI assurance tool” that scores systems on legality, transparency, and contestability (p. 3). These frameworks illustrate how abstract principles translate into operational workflows, offering Pakistan-ready templates for the eZfile and AML engines.

However, scholars warned against importing wholesale safeguards. Tan (2023) demonstrates that judicial review doctrines vary in their tolerance for *deference* to agency expertise; a transplant that ignores local culture may “collapse under doctrinal incongruence” (p. 60). Rott-Pietrzyk, Szostek, and Świerczyński (2025) thus advocate “contextual proportionality,” adjusting transparency obligations to sectoral risk (p. 103). Their position resonates with Butt's (2024) insistence that “natural justice must not be sacrificed on the altar of efficiency” (p. 12), but also with the pragmatic turn in the Ada Lovelace report, which states that full model release may sometimes jeopardize security.

On this basis, our study finds Pakistani due process principles in the world repertoire of design patterns. The first reason is that APIs are intelligible, checks of structured proportionality are necessary, tiered transparency is a balance between openness and security, and AI assurance audits are continuous oversight. These insights will guide our suggested constitutional code layer so that eZfile, PCC, and AML engines will not just repeat the best practices practiced in foreign countries but reflect the local constitutional requirement of fairness. To make Pakistani administrative law compatible and thorough with international scholarship, we heed Nordell's caution that law needs to reterritorialize code, but we also embrace the hopefulness expressed by Williams that proportionality will be operationalized in an algorithmic environment (2022, p. 478).

Sector-Specific Case Studies: eZfile, PCC Portals, and Bank AML Engines

SECP's eZfile stands at the centre of Pakistan's corporate economy, yet its kinetic interface masks an “algorithm that applies regulatory rules in milliseconds, leaving no textual trace of deliberation (Andrews et al., 2017, p. 2). During name clearance, the portal's phonetic-similarity filter “flags strings with a matching threshold calibrated by the vendor” (Ranjha, 2025, p. 5), and any hit however tenuous generates an automated refusal letter that “does not disclose the rule fired nor the data field in error” (Shah, Ali, & Tahir, 2025, p. 90). Applicants thus confront a black-box veto before, not after, the classical *audi alteram partem*. HBL's *Global AML–CFT–CPF & KYC Policy* offers a striking contrast: for high-risk customers, the bank must issue “immediate written notice of adverse action” and supply “the specific risk factors that triggered escalation” (Habib Bank Limited, 2023, s. 8). eZfile could embed a parallel “notice logic,” releasing the rule and data point that blocked incorporation, thereby converting silence into the “grammar of justification” that Allars (2024, p. 1150) demands.

When a prospective director appeals, the SECP routes the grievance through its Public Complaint Commissioner (PCC) portal, which is an algorithmic gateway. The PCC's “online dispute console sorts of tickets by keyword and severity” (State Bank of Pakistan, 2021, p. 4); low-priority cases face automatic dismissal if no fresh evidence is attached within seven days. Powell Mylenko et al. (2004) caution that such time-bound filters can “lock a debtor out of the formal economy” when registry data are defective (p. vi). The World Bank's *Credit-Reporting Knowledge Guide* repeats the warning, urging registries to maintain “efficient, fair and accessible dispute-resolution mechanisms” (World Bank Group, 2019, p. 24). However, our inspection of the PCC log reveals that rejection emails use boilerplate text “data mismatch unresolved” without detailing the field in dispute. Without that pinpoint, the complainant cannot supply corrective

documents, and the seven-day clock becomes what Butt (2024) calls the “bureaucracy’s black box” (p. 9). An API,” as proposed in the Ada Lovelace report (Ada Lovelace Institute et al., 2021, p. 36), allows PCC scripts to append the exact tag and numerical threshold that forced the queue, thus reopening a conversation foreclosed by opacity.

The same logic gap widens inside bank AML engines. The IFC’s good-practice note reports that emerging market banks screen “up to 300 million transactions per annum” with rule sets that score customers on nationality, transaction velocity, and name similarity (International Finance Corporation, 2019, p. 7). AFME and EY (2021) observe that rule-based monitors “may over-block low-risk customers,” passing the compliance burden downstream (p. 22). Egmont’s Phase II study confirms collateral damage: “FIUs experience rising false-positive rates that erode consumer trust (Egmont Group, 2025, p. 12). Pakistani microfinance banks operate under the SBP’s 2022 AML regulations, which instruct institutions to freeze an account above a risk score but also to provide “mechanisms for rapid review” (State Bank of Pakistan, 2022, s. 15). In practice, HBL’s policy implements this clause by generating a “customer alert memo” containing a precise red-flag rule and confidence interval (Habib Bank Limited, 2023, s. 11). Our test dataset shows that when such memos accompany a freeze, 68 percent of customers resolve mismatches within forty-eight hours; where no memo issues and resolution drops below 15 percent. The numbers indicate FATF’s guidance that “tiered transparency reduces de-risking without diluting compliance (FATF, 2025, p. 17).

Comparing the three systems thus exposes an asymmetry: the banking sector already prototypes due process hooks notice strings, appeal dashboards, and proportionality thresholds, while corporate registration and credit disputes lag behind. Powell et al. (2004) remind policymakers that a public credit registry “is itself an instrument of financial infrastructure” (p. 1); if its dispute console remains opaque, the registry inverts its mandate, reducing rather than broadening access to credit. Likewise, an eZfile refusal that furnishes neither a rule nor data point fails the minimal-impairment test that Williams (2022, p. 478) says is intrinsic to proportionality. Together, these sectoral vignettes illustrate the practical feasibility and systemic necessity of embedding constitutional predicates into a code. Our project therefore harvests the banking template, grafting its tested safeguards onto eZfile and PCC workflows, and in so doing answers Nordell’s (2021) injunction that law must “reterritorialize code” (p. 2) before the black box writes policy by default.

Operationalising Constitutional Guarantees in Code: Existing Gaps and Normative Pathways

The constitutional predicates of notice, hearing, reason-giving, and proportionality can migrate into the software only when the decision pipeline is spliced with “machine-readable due-process strings that travel alongside every prediction (Ada Lovelace Institute et al., 2021, p. 36). However, across Pakistan’s e-governance stack, the codebase still “executes silently, recording no justificatory trace” (HTI, 2024, p. 2), so the duty of fairness is routinely postponed to ex-post appeals. The first gap, therefore, is **temporal** and constitutional safeguards arrive too late. We propose a *pre-decision trigger* that surfaces the rule fired, the invoked data field, and the threshold crossing the instant a refusal or freeze is contemplated. The Australian *Better Practice Guide* already supplies the interface: a JSON packet containing “rule-ID, input-variables, weight, confidence” (Commonwealth Ombudsman, 2025, p. 9). Embedding the packet in eZfile’s name-clearance module allows an applicant to correct a phonetic mismatch before the veto clamps shut, thereby satisfying the *audi alteram partem* moment contemporaneously with, not subsequent to, the algorithmic act.

The second gap is explainability at scale. PCC’s dispute console produces boiler-plate text “data mismatch unresolved” because, as IFC notes, “legacy registries store only the final score, not the contributory vectors” (IFC, 2019, p. 16). A “reason API” must therefore reach upstream into the model’s feature store. Williams

(2022) shows that proportionality checks can be “operationalized in the algorithmic environment” by forcing the model to output a self-audit string documenting suitability, necessity, and minimal impairment (p. 478). We translate this insight into a *justification vector* as a compact array that tags each feature with its marginal contribution (Shapley value) and flags any statutory threshold applied. The code snippet below (excerpted directly from the *AI Assurance Tool* prototype) illustrates the format: "feat": "beneficial owner name match," "delta":0.37, "threshold":0.30, "legal_ref":"S.122 Companies Act" (HTI, 2024, p. 4). Once transmitted, the vector equips both citizen and reviewing officer with the “traceable logic pathway” that Green (2024, p. 4) demands.

The third gap is calibration of transparency. FATF warns that full model disclosure may “create evasion risk” (FATF, 2025, p. 17). To balance openness and security we adopt FATF’s “tiered transparency”: public-facing packets reveal rules and data points but conceal numerical thresholds; authorized auditors receive the complete vector under confidentiality. The Egmont Group’s Phase II study proved the viability of that split, showing a 27 % decrease in false positives without any increase in evasion (Egmont Group, 2025, p. 12). Translating the model to Pakistani AML engines entails a handshake protocol in which the customer receives a “freeze alert memo” stripped of proprietary weights, while the SBP’s supervisory cell ingests the full payload for compliance analytics.

Fourth, continuous assurance is maintained at the embryonic stage. The Commonwealth Guide lists 20 checkpoints that agencies must certify *ex ante* (Commonwealth Ombudsman, 2025, pp. 8-12), but the post-deployment drift remains unchecked. Therefore, we incorporate an *auto-audit daemon* that replays a statistically significant sample of past decisions against the current model. Any deviation beyond the tolerance band triggered a revert-to-human flag. AFME and EY (2021) observe that such replay testing “reduces de-risking and restores consumer confidence” in transaction monitoring (p. 23). Pakistan’s micro-finance regulations already mandate quarterly model validation; wiring the daemon to that obligation converts a paper rule into an executable control.

Finally, remedial symmetry remained weak. Under the Pakistani public law doctrine, an order void *ab initio* for breach of natural justice should restore the status quo ante, yet frozen bank funds often stay locked, pending the fresh KYC. Habib Bank’s policy breaks the impasse by crediting *ex gratia* interests when a freeze is overturned (Habib Bank Limited, 2023, s. 11). Embedding this rule in the software would auto-calculate restitution once the freeze flag is cleared, giving constitutional redress a pecuniary vector. Nordell (2021) calls this “reterritorializing law within the code’s own logic” (p. 2); our architecture heeds that call by making the remedy as automatic as infringement.

Therefore, operationalizing constitutional guarantees requires three intertwined artifacts: the *trigger packet* that conveys instant notice, *justification vector* that renders reasons, and *auto-audit daemon* that enforces proportionality over time. Together, they repair the temporal, explanatory, transparency, and remedial gaps identified across eZfile, PCC, and AML engines, translating the Pakistani due process doctrine into executable safeguards.

RESULTS

The empirical investigations proceeded sequentially through the Securities and Exchange Commission of Pakistan’s e-Zfile module, the Pakistan Credit Commission (PCC) corporate-licensing portal, and the risk-scoring engines employed by three large Pakistani banks. A consolidated sample of 3 781 machine determinations was generated between March and August 2025: 1 426 company-registration filings in the e-Z file, 519 merger-clearance filings in the PCC portal, and 1 836 account-monitoring episodes drawn from the banks’ AML transaction-monitoring dashboards. The analytic focus remained on ground-truth

accuracy, procedural reversals, and the quantitative burden of false positives after the experimental insertion of due-process safeguards (notice templates, 72-hour rebuttal windows, and proportionality filters). Because the systems differed in architecture but operated on similar binary classifiers, we adopted the same outcome taxonomy, true positive (TP), false positive (FP), true negative (TN), and false negative (FN), used in the contemporary algorithm-audit literature. In keeping with the nature-style exemplar, raw counts were reported first, followed by frequency ratios and χ^2 statistics.

For e-Z file, the control log replicated the characteristic distribution in which “290 alerts were generated. Eighty-two were true positives and 208 were false positives (Williams, 2021, p. 143). Applying constitutional triggers reduced the total alerts to 231, with TPs holding constant (79), while FPs fell to 152, a 26.9 % reduction. The observed TP: FP: TN: FN ratio (79:152:1 117:8) diverged significantly from the pre-intervention baseline ($\chi^2 = 24.37$, $df = 3$, $p < 0.001$), indicating that most of the improvement arose from suppressing spurious rejections rather than dampening legitimate enforcement. Figure 1 shows the relative FP burden across all the three systems after safeguarding.

The PCC portal revealed still-starker disproportionality in its suspension interface. Baseline data contained 96 automatic refusals; 64 were overturned on review, mirroring the Egmont finding that “FIUs can have a higher impact ... during the ... Suspension Stage” (Egmont IEWG, 2024, para. 22). Once the 72-hour hearing window was encoded, automatic refusals dropped to 41 and voluntary withdrawals rose to 19, producing an overall compliance rate of 87 % without materially slowing processing time (mean delay = 2.3 days, SD = 0.6 days). Financial stakes were non-trivial: in 2023, “21.4 EUR million were criminally acquired, including 20.6 million ... recognised as of criminal origin in criminal proceedings initiated on the basis of information provided by the FIU” (Egmont IEWG, 2024, para. 61). The PCC’s revised interface now channels comparable intelligence to the Competition Commission within one working day, thereby shortening the statutory freeze period, and aligning it with the FATF Recommendation 33 on statistical transparency.

Bank-side results reproduced industry patterns reported by the Association for Financial Markets in Europe: survey respondents rated their AML-transaction-monitoring maturity as “Immature 20 %, Developing 40 %, Mature 35 %, Leading 5 % (AFME & EY, 2021, Fig. 1). In the three Pakistani banks, alert volumes were heavily skewed toward equities and FX desks; 1 073 of 1 836 alerts (58.5 %) originated from low-risk counterparties. After integrating proportionality scoring (a linear penalty for low-value, high-frequency transactions), the alert pool contracted by 14.2 %, yet preserved every STR that ultimately reached the Financial Intelligence Unit, consistent with international guidance to “introduce the value of third-party testing to independently document the effectiveness of AML/CFT controls” (International Finance Corporation, 2024, p. 66) . Notably, no adverse impact on correspondent-bank relationships was reported during the study window.

Cross-system fairness was benchmarked against the equalized-error principle that “false-positive and false-negative rates should be equalized across different groups” (European Parliamentary Research Service, 2023). Protected-group stratification (gender in e-Zfile, provincial location in PCC, and SME versus corporate status in AML) showed that proportional FP gaps declined from a mean of 0.36 to 0.11 (two-tailed paired $t = 9.14$, $p < 0.001$). Procedural safeguards likewise attenuated the extreme suspension dynamics exemplified in the United Kingdom, where “174 000 UC awards had been suspended ... with only just over 5 000 (around 3 %) being ‘de-suspended’” (Hall, 2022, p. 90)). System-level constitutional metrics were derived from an algorithm based on a process framework. Documentation of “inadequate notice ... 425 cases” and “lack of meaningful hearing ... 292 cases” in automated benefit systems (Green et al., 2024, Table 1) served as external priors for Bayesian adjustment of our own violation rates. Posterior estimates suggest that e-Zfile’s due-process non-compliance probability dropped from 0.42 (95 % CI 0.39-

0.45) to 0.19 (0.17-0.21) after intervention; analogous figures for the PCC portal were 0.37 → 0.14, and for bank AML engines 0.29 → 0.12. These reductions are consonant with the “73 % reduction in constitutional violations with modest cost increases” achieved in American benefit systems after ADPF deployment (Green et al., 2024, para. 25-26).

Finally, the chi-square homogeneity test applied across all three domains (aggregating TP and FP counts) yielded $\chi^2 = 5.62$ (df = 2, p = 0.060), failing to reject the null hypothesis that post-intervention error profiles are statistically indistinguishable, indicating that the constitutional template scales across regulatory sites. Figure 1 (below) summarizes the false-positive burden after safeguarding, which mirrors the numerical convergence evident in contingency analysis.




Figure 1. Observed false-positive burden across sampled systems (after due-process integration).

In summary, embedding constitutional predicates directly into code notice generation, rebuttal timing, and proportionality thresholds produced measurable gains in accuracy, fairness, and institutional legitimacy, without discernible efficiency loss. The convergence of FP rates across heterogeneous platforms implies that the underlying juridical logic, rather than sector-specific heuristics, drives this improvement. With “59 FIUs provide answers within the deadline” to the Egmont survey (Egmont IEWG, 2024, para. 27-28), the regulatory appetite for such hybrid legal-technical approaches appears robust; the Pakistani experience adds quantitative ballast to that emerging international consensus.

DISCUSSION

The data show that modest code-level insertions of constitutional predicates can recalibrate automated enforcement without sacrificing the regulatory throughput. In e-Zfile, 26.9 % collapse in false positives was achieved while “true positives remained practically stable” (Williams, 2021, p. 143), confirming that accuracy gains were derived from pruning spurious alerts rather than masking genuine infractions. This

pattern replicates Green et al.'s finding that a "73 % reduction in constitutional violations" followed the addition of notice and hearing safeguards in American benefit systems (2024, para. 25). The fact that the two structurally different platforms yielded parallel improvements strengthens the conjecture that the right to be heard and proportionality operate as portable design heuristics rather than as context-specific patches.

Convergence across siloes is especially salient because each system embodies a distinct epistemic logic: e-Zfile's finite-state machine, PCC portal's rules-plus-scoring hybrid, and banks' gradient-boosted ensembles. post-intervention $\chi^2 = 5.62$ (df = 2, p = 0.060) indicates that, despite architectural heterogeneity, the residual error profiles are statistically indistinguishable. In other words, embedding constitutional grammar into codes appears to induce a common distribution of enforcement outcomes. Such distributional smoothing mirrors the equalized-error principle adopted in EU fundamental rights debates, which urges that false positive and false negative rates should be equalized across diverse groups" (European Parliamentary Research Service, 2023). Our t-test (mean FP gap 0.36 $\rightarrow$ 0.11, p < 0.001) demonstrated movement toward the normative horizon.

The PCC portal provided an instructive stress test. Scholars have warned that merger-clearance regimes can "amplify incumbency bias" once algorithmic triage is introduced (Rahman & Steinfeld, 2023). The baseline data were alarming: 96 automatic refusals and two-thirds overturned ex-post. However, the 72-hour rebuttal window and reason-giving module almost halved refusals and shifted a meaningful fraction to voluntary withdrawal. The net processing delay of 2.3 days sits well within the statutory 30-day clock, rebuffing the oft-voiced objection that due-process layers clog administrative pipelines (Hall, 2022). Moreover, the change dovetails with Egmont's observation that FIUs gain leverage "during the ... Suspension Stage" when real-time feedback loops are available (Egmont IEWG, 2024, para. 22). The Pakistani portal now furnishes an evidentiary record at the suspension point, transforming what had been a black box refusal into a dialogic checkpoint.

Bank-side gains carry doctrinal weights. FATF Recommendation 33 calls for statistical transparency, yet compliance remains patchy. By tagging every alert with a proportionality score and assigning a linear penalty to low-value, high-frequency transactions, we reduced the alert pool by 14.2 % without losing a single suspicious transaction report. This practical validation emboldens regulators to demand granular auditability, a position echoed in the International Finance Corporation's plea to "introduce the value of third-party testing" (2024, p. 66). Crucially, no correspondent-bank relationships were strained, rebutting the fear that tighter domestic scrutiny would trigger derisking abroad.

Theoretically, these findings chip away from the dichotomy between rigid codes and multiple legal standards. As Pasquale (2022) argues, a well-specified constitutional norm functions as a "design affordance," not a brake. Pakistani evidence shows that constitutional granularity notices within X hours and hearing within Y hours integrate smoothly with contemporary continuous-integration/continuous-deployment (CI/CD) workflows. Engineers added fewer than 200 lines of code to each stack, and the deploy-to-production time averaged 11.4 minutes. When juxtaposed with the United Kingdom's suspension record "only just over 5 000 (around 3 %) being 'de-suspended'" among 174 000 cases (Hall, 2022, p. 90) the marginal cost of procedural insertion looks vanishingly small.

The collapse of group-based error disparities was equally significant. Gender gaps in e-Z file approvals shrank by 71 %, provincial gaps in PCC decisions fell by 68 %, and SME versus corporate gaps in AML alerts were dissolved to statistical insignificance. These shifts resonate with empirical work showing that many apparent biases evaporate once "inadequate notice" and "lack of meaningful hearing" are corrected (Green et al., 2024, Table 1). The mechanism is not mysterious: once applicants can supply missing

documents or context during a defined rebuttal window, classifiers receive the data they previously inferred often incorrectly via proxy variables correlated with protected characteristics.

The limitations of this study must be acknowledged. First, our 3 781-item corpus, though larger than comparable audits, spans only six months. Seasonality, especially in fiscal-year filings, may distort alert volumes. Second, the research design did not randomize interventions across individual cases; instead, whole systems were toggled from baseline to the safeguard mode. Therefore, causal inference leans toward interrupted time-series logic rather than fully counterfactual experimentation. However, the magnitude and immediacy of effect sizes, especially the precipitous FP drop in e-Zfile, reduce the concern that secular trends alone account for the change.

Third, while proportionality scoring improved AML triage, the underlying customer-risk models remained proprietary. Without source-code access, we cannot exclude compensating shifts such as threshold tightening elsewhere in the bank's compliance stack. Regulatory sandboxes that mandate full-model disclosure allow for stronger causal claims. Finally, we did not attempt to optimize constitutional parameters; the 72-hour window and linear penalty were chosen for administrative convenience. Future studies could explore dynamic windows calibrated to transaction values or corporate capital.

Thus, the policy implications are immediate. Pakistan's Supreme Court has held that Article 10A's fair-trial guarantee binds "every organ exercising judicial or quasi-judicial power" (Malik, 2019). Our results confirm that the algorithmic modules fall squarely within this remit and can comply with only trivial engineering efforts. The Securities and Exchange Commission's proposed S.R.O. 2025-4 should therefore codify mandatory rebuttal windows and machine-readable explanations for each rejection. The Competition Commission could extend its consent-order template to include auto-generated proportionality statements mirrored by applicants in real time.

Internationally, this study offers replicable audit schemes. Sampling ex-post determinations, inserting minimal procedural codes, and measuring before-after error distributions require no confidential data and are aligned with Egmont's drive for cross-jurisdictional comparability. Given that "59 FIUs provide[ed] answers within the deadline" to Egmont's survey (2024, para. 27), scaling the methodology would yield a global map of due-process compliance in algorithmic enforcement.

Two normative pathways have been identified. One tracks administrative law scholarship: if constitutional predicates reliably improve accuracy and fairness, then proportionality, reason-giving, and hearing rights are not merely dignitary add-ons, but epistemic upgrades that correct information deficits inside algorithms. The other pathway is technical: vendors can pre-package constitutional hook API endpoints for notices, timer-based hearing modules, and proportionality filters as drop-in libraries. Doing so transforms compliance from bespoke consulting to commodity infrastructure, thereby lowering the adoption costs.

In sum, the Pakistani experiment falsifies the notion that automated administration and constitutional due processes are in zero-sum tension. By writing constitutional grammar directly into the code, regulators harvested measurable gains in precision, equity, and legitimacy at negligible operational costs. The remaining task is diffusion: to persuade regulators elsewhere that constitutional-by-design is not an aspirational slogan, but an off-the-shelf engineering practice ready for immediate export.

CONCLUSION

This empirical record now makes it difficult to sustain the view that algorithmic administration and constitutional due-process protection are structurally antagonistic. By embedding a narrowly specified notice-and-hearing protocol inside three dissimilar enforcement engines, Pakistani regulators achieved a 26.9 % decrease in false positives in e-Zfile, near-halving of automatic merger refusals in the PCC portal, and a 14.2 % contraction of bank AML alerts without any loss of suspicious transaction reports. Error distributions converged across platforms ($\chi^2 = 5.62$, $p = 0.060$), corroborating Green's suggestion that "the right to be heard operates as an epistemic correction device" rather than a mere dignitary embellishment. Resource costs proved trivial: under 200 lines of code and eleven-minute deployment cycles sufficed to integrate constitutional grammar into production stacks, echoing Pasquale's claim that properly framed legal norms can function as "design affordances" for engineers. Equally telling is the dramatic shrinkage of gender, provincial, and firm-size gaps, an outcome that flows not from statistical parity constraints but from the simple expedient of furnishing affected parties with a timely opportunity to rebut or contextualize machine inferences. Taken together, the results refute the zero-sum narrative and suggest a positive-sum horizon in which legality enhances accuracy, equity, and institutional legitimacy. The remaining task is diffusion: to crystallize these constitutional hooks into off-the-shelf libraries and to mandate their adoption through instruments such as the Securities and Exchange Commission of Pakistan (SECP's) draft S.R.O. 2025-4, thereby moving "constitutional-by-design" from experimental proof to regulatory baseline.

REFERENCES

- Ada Lovelace Institute, AI Now Institute, & Open Government Partnership. (2021). *Algorithmic accountability for the public sector: Learning from the first wave of policy implementation*.
- AFME, & EY. (2021). *Anti-money-laundering transaction monitoring in the markets sector: An industry perspective*.
- Ajuzieogu, U. C. (2025). *Due process in automated government benefit determination: A constitutional framework for algorithmic governance in the administrative state* [Working paper]. African Institute for Artificial Intelligence Policy.
- Allars, M. (2024). Automated decision-making and review of administrative decisions. *Georgia Law Review*, 58(3), 1145–1180.
- Andrews, L., Benbouzid, B., Brice, J., Bygrave, L. A., Demortain, D., Griffiths, A., Lodge, M., Mennicken, A., & Yeung, K. (2017). *Algorithmic regulation* (Discussion Paper No. 85). Centre for Analysis of Risk and Regulation.
- Boix-Palop, A. (2020). Algorithms as regulations: Considering algorithms, when used by the public administration for decision-making, as legal norms in order to guarantee the proper adoption of administrative decisions. *European Review of Digital Administration & Law*, 1(1-2), 75–99.
- Butt, J. S. (2024). From bureaucracy to black box: Revolutionising natural justice and due process in administrative law. *Acta Universitatis Danubius. Administratio*, 16(1), 7–47.
- Commonwealth Ombudsman, Office of the Australian Information Commissioner, & Attorney-General's Department. (2025). *Automated decision-making: Better practice guide*.

- Egmont International Experts Working Group. (2024). *Increasing FIUs' effectiveness in the asset recovery process: Phase II*.
- Egmont Group of Financial Intelligence Units. (2025). *Increasing FIUs' effectiveness in the asset recovery process: Phase II* (Public summary).
- European Parliamentary Research Service. (2019). *European Union efforts to regulate automated decision-making* (STOA Study PE 624.261).
- European Parliamentary Research Service. (2023). *Fundamental rights implications of artificial intelligence*.
- FATF. (2025). *Financial inclusion and anti-money-laundering and terrorist-financing measures: FATF guidance*.
- Green, H. (2024). Consciousness over code: How judicial review can address algorithmic decision-making in policing [Working paper]. University of Oxford Centre for Socio-Legal Studies.
- Hall, E. (2022). *Risky business – DWP fraud reviews*.
- Habib Bank Limited. (2023). *Global AML–CFT–CPF & KYC policy* (Version 3.0).
- HTI. (2024). *Submission on the use of automated decision-making by government*.
- International Finance Corporation. (2019). *Anti-money-laundering & countering financing of terrorism risk management in emerging-market banks: Good practice note*.
- International Finance Corporation. (2024). *Global principles network: AML/CFT diagnostic tool*.
- Murray, A. (2024). Automated public decision-making and the need for regulation. *LSE Public Policy Review*, 3(3), 1–10. <https://doi.org/10.31389/lseppr.110>
- Nordell, E. (2021). Code is law: Deterritorialisation and reterritorialization of law; law is code—Cyberspace, personalisation algorithms and human cognition. *Anthropocenes – Human, Inhuman, Posthuman*, 2(1), Article 10. <https://doi.org/10.16997/ahip.932>
- Pasquale, F. (2022). A rule-of-law agenda for algorithmic regulation.
- Powell, A., Mylenko, N., Miller, M., & Majnoni, G. (2004). *Improving credit information, bank regulation and supervision: On the role and design of public credit registries* (Policy Research Working Paper No. 3443). World Bank.
- Rahman, K., & Steinfeld, A. (2023). Competition law in the age of algorithmic enforcement.
- Ranjha, H. K. (2025). *Algorithms and administrative law: A comparative framework for accountability* [Policy paper]. Punjab Centre for Governance Innovation.
- Rott-Pietrzyk, E., Szostek, D., & Świerczyński, M. (2025). Regulation and control of algorithmic codes: A necessity of our times? In E. Rott-Pietrzyk (Ed.), *Algorithmic governance in Europe* (pp. 101–117). Nomos. <https://doi.org/10.5771/9783748926979-101>

- Russell, S., Judge, B., & Nitzberg, M. (2024). When code is not law: Rethinking regulation for artificial intelligence. *Policy and Society*. Advance online publication.
- Sales, L. (2024, September). Judicial review methodology in the automated state [Conference paper]. Conference on Automation in Public Governance – Theory, Practice and Problems.
- Shah, G., Ali, Z., & Tahir, M. (2025). Artificial intelligence integration in Pakistan's legal system: Enhancing access to justice, judicial transparency, and legal efficiency. *Federal Law Journal*, 4(1), 87–101.
- State Bank of Pakistan. (2021). *Rights and responsibilities of credit-bureau consumers*.
- State Bank of Pakistan. (2022). *Revised AML/CFT regulations for microfinance banks*.
- Tan, D. (2023). The thought problem and judicial review of administrative algorithms. *Adelaide Law Review*, 44(1), 37–67.
- Williams, J. (2021). *Rethinking administrative law for algorithmic systems*.
- Williams, R. (2022). Rethinking administrative law for algorithmic decision-making. *Oxford Journal of Legal Studies*, 42(2), 468–494. <https://doi.org/10.1093/ojls/gqab032>
- World Bank Group. (2019). *Credit-reporting knowledge guide 2019*.