Journal of Environmental Economics & Chemical Processes (JEECP)

https://doi.org/10.22034/jeecp.2026.2090285.1038

Journal Homepage: WWW.JEECPJournal.com

Journal of Environmental Economics & Chemical Processes (JEECP)

ORIGINAL RESEARCH PAPER

Adaptive Data-Consistency Governance for Sustainable Chemical Industries: Economic, Environmental, and Regulatory Implications of Safety-Constrained Q-Learning

Hadi Davardoost [a, *]; Julia Peryazeva [b]; Hawezy Las Jamal Alan

[a] *Department of Computer Science and Digital Modeling, Saint Petersburg Mining University of Empress Catherine II , Saint-Petersburg, Russia.*
[b] *Department of Computer Science and Digital Modeling, Saint Petersburg Mining University, Saint Petersburg, Russia.*
[c] *Department of International Business Development, Peter the Great St. Petersburg Polytechnic University, Saint Petersburg, Russia.*

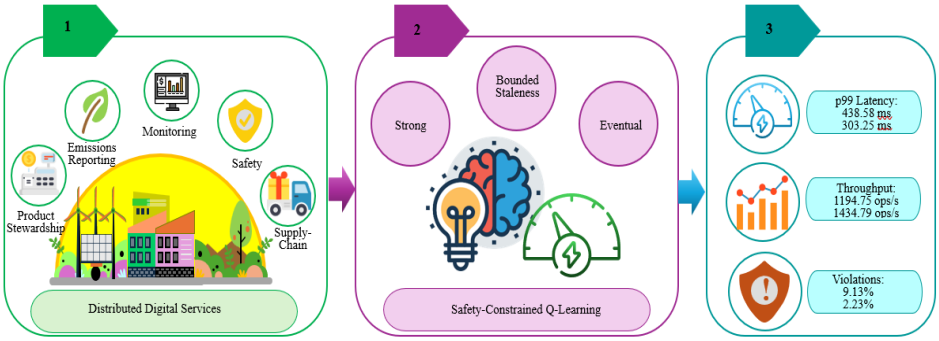
P A P E R I N F O

Paper history:
Received: 08/07/2026
Revised: 10/04/2026
Accepted in revised form: 10/14/2026
Published: 03/30/2027

***Keywords*:**
Adaptive Consistency;
Environmental Policy;
Economic Policy;
Data Governance;
Q-Learning;
Industrial Emissions;
Bounded Staleness.

A B S T R A C T

Industrial chemical production increasingly depends on geographically distributed digital services for process monitoring, emissions reporting, product stewardship, safety management, and supply-chain coordination. These services must deliver timely information without imposing unnecessary communication, computing, and compliance costs. This article presents a policy-oriented secondary analysis of a three-region simulation in which a safety-constrained Q-learning controller selected among strong, bounded-staleness, and eventual consistency modes. The technical results are reinterpreted through economic, environmental, and chemical-governance lenses rather than presented as a new database benchmark. During a 600-second burst interval, compared with static strong consistency, the adaptive controller reduced p99 latency from 438.58 to 303.25 ms, increased throughput from 1,194.75 to 1,434.79 operations per second, and reduced latency-service violations from 9.13% to 2.23%. Across six scenarios, communication volume was 31.4–60.0% lower than under static strong consistency, while the maximum staleness constraint was maintained. These findings support a tiered governance model: strong consistency for emergency shutdown, toxic-release, and legally decisive records; bounded staleness for environmental monitoring and operational compliance; and carefully supervised eventual consistency for low-risk analytical and reporting workloads. The article develops equations for total policy cost, communication-related energy and emissions proxies, compliance risk, and a multi-criteria policy score. It concludes that adaptive consistency should be governed as an industrial policy instrument, with auditable limits, human override, fallback rules, and data-class-specific service objectives. The environmental interpretation remains a proxy because direct electricity and carbon measurements were not available.



1. Introduction

Chemical manufacturing is becoming a data-intensive industrial activity. Production planning, distributed control, laboratory information management, emissions monitoring, hazardous-material inventories, maintenance records, and supply-chain documentation are increasingly connected across plants, corporate data centers, cloud platforms, laboratories, regulators, and external partners. This digital integration can improve decision speed and traceability, yet it also creates a governance problem: not every item of industrial data requires the same level of freshness, coordination, or cross-region replication[1,2]. Chemical and environmental policy is moving toward lifecycle accountability, verifiable emissions information, faster accident response, and more structured data sharing. The Global Framework on Chemicals calls for coordinated management of chemicals and waste across their life cycles, while the OECD chemical-accident guidance emphasizes reliable information for prevention,

* Corresponding Author: Hadi Davardoost
E-mail address:
ISSN: 3116-0050
URL: https://www.jeecpjournal.com/article_736321.html

Please cite this article as: H. Davadoost, J. Peryazeva, H. Las Jamal Alan, *Journal of Enviromental Economics & Chemical Processes (JEECP)*, 2027, *4*(1), 23-30. https://doi.org/10.22034/jeecp.2026.2090285.1038

/© 2027 The Author(s). Published by Power of New Vision Movement Institute.

preparedness, response, investigation, and organizational learning [3] In the European context, the revised Industrial Emissions Directive broadens attention to resource efficiency, circularity, decarbonization, and environmental performance, while the Corporate Sustainability Reporting Directive, Data Act, and Ecodesign for Sustainable Products Regulation increase expectations concerning data availability, sustainability disclosure, access rights, and digital product information [4].

These developments imply that industrial data architecture is not a purely technical concern. Excessively strong synchronization can increase delay, inter-region traffic, and infrastructure cost. Excessively weak synchronization can expose operators and regulators to outdated information, inconsistent inventories, delayed alarms, or inaccurate environmental records. The appropriate policy is therefore conditional on the consequences of stale data. A reactor trip signal and an aggregated monthly sustainability dashboard should not be governed by the same consistency rule [5].

The distributed-systems literature has long described the trade-off among consistency, availability, latency, and partition tolerance [6,7]. Systems such as Dynamo emphasized availability and conflict resolution, whereas globally coordinated systems such as Spanner pursued strong transactional guarantees across regions [8,9]. Later work formalized numerous consistency semantics and quantified staleness as a measurable outcome rather than a vague property [10,11]. Adaptive approaches, including specifically, PCAP adapts to probabilistic consistency or latency SLAs, whereas OptCon uses predictive learning to select consistency settings under user-defined latency and staleness thresholds [12,13]. These frameworks address service-level constraints but do not develop a chemical-safety governance framework linking data classes to operational consequences, audit responsibilities, and human override. The present secondary analysis addresses this governance gap by connecting an independent supervisor that rejects actions predicted to exceed a staleness limit with class-specific policy requirements, fallback procedures, and accountability mechanisms; this supervisory approach does not constitute a formally verified safety guarantee.

The sustainability dimension adds another layer. Data centers and networked computing consume material and energy resources, and data-center electricity demand is expected to grow as digital services expand [14]. Efficiency gains have historically moderated the relationship between digital-service demand and electricity use, which means that software and architecture choices can influence environmental performance even when they do not directly measure electricity at the application level [15]. Replication policy is especially relevant because cross-region messages, retransmissions, coordination waits, and redundant data movement are all affected by consistency design [16,17,18]. This article does not attempt to relabel a technical database study as a chemical-policy experiment. Instead, it develops a distinct secondary analysis that asks how selected simulation results can inform governance of data services used by chemical and environmentally regulated industries. The analysis is intentionally policy-oriented: it translates latency, throughput, staleness, and communication outcomes into economic, environmental, safety, and regulatory decision criteria.

The study addresses three policy questions:

- How should industrial data be classified according to the economic, environmental, and safety consequences of staleness?
- When can adaptive consistency reduce coordination and communication burden without undermining chemical-safety or regulatory obligations?
- Which governance controls are required before a learning-based consistency controller can be used in a regulated industrial environment?

The first question is addressed through the risk-based data classification in Section 2.3 and its policy interpretation in Section 4.5. The second is examined through the analytical framework in Section 3 and the scenario results and economic and environmental interpretations in Section 4. For the third question, Section 5.1 establishes the governance principles, while Section 5.4 translates them into an organizational implementation and assurance roadmap. Section 6 synthesizes the answers into policy recommendations.

1.2. Policy and Industrial Context

1.2.1. Data freshness as an industrial-policy variable

Data freshness is often treated as a technical quality attribute, but in regulated industry it has distributive and legal consequences. Freshness affects who receives information first, how quickly a deviation is recognized, whether a record can support an audit, and whether an operator can rely on a remote replica during an emergency [19,20]. A consistency policy therefore allocates risk among operational units, corporate management, regulators, communities, and digital-service providers. Data freshness is often treated as a technical quality attribute, but it also has governance implications in regulated industry. Established data-quality literature identifies timeliness and accessibility as distinct dimensions of information fitness for use [21], while the OECD's data-governance principles connect access and sharing arrangements with stakeholder rights, responsibilities, and accountability [22]. Applied to geographically distributed industrial services, these principles suggest that differences in replication delay can create unequal access to current information among otherwise authorized users, potentially affecting when operators, managers, and regulators recognize and respond to a deviation. Consistency policy therefore helps determine how the risks associated with outdated information are distributed across stakeholders.

For policy purposes, the important question is not whether data are globally identical at every instant. The relevant question is whether each data class remains sufficiently current, complete, and traceable for its intended decision. This distinction aligns with risk-based chemical management. The OECD guidance does not require every informational resource to have identical technical architecture; it requires systems that support prevention, preparedness, response, consequence assessment, investigation, and communication [3].

Adaptive consistency can support this approach by assigning strict coordination to high-consequence records and more flexible replication to low-consequence workloads. However, the adaptive mechanism must be transparent. A regulator or plant manager should be able to determine which rule was active, why a weaker mode was selected, what safety limit applied, and whether a fallback was triggered [23].

1.2.2. Environmental and economic relevance of digital coordination

In a geo-distributed service, strong coordination can require multiple cross-region acknowledgements before an operation completes. The direct economic burden may include bandwidth charges, cloud charges, additional replica activity, queueing delay, transaction timeout, and lost productivity. The indirect burden may include delayed environmental reporting, slower laboratory release, and lower resilience during network impairment.

The environmental burden is more difficult to quantify. Communication volume is not equivalent to electricity consumption or greenhouse-gas emissions. Conversion requires network-energy intensity, data-center power data, marginal electricity factors, equipment utilization, and temporal and regional boundaries. Nevertheless, transferred data volume is a defensible intermediate indicator because it records a physical service demand that must be supported by networking and computing equipment. The analysis therefore treats communication reduction as an environmental proxy, not as measured carbon reduction [24].

This distinction follows broader sustainability practice. Chemical sustainability assessment increasingly combines lifecycle data, model-based estimates, and uncertainty disclosure rather than presenting a single proxy as a complete environmental result [25,26]. Decision frameworks for sustainable chemical processes also emphasize the joint consideration of economic feasibility, technology maturity, policy alignment, and supply-chain constraints [27].

1.2.3. Chemical-policy data classes

The policy analysis distinguishes five data classes. The classification is functional rather than technology-specific and can be adapted to petrochemical plants, pharmaceutical production, bulk-chemical logistics, environmental laboratories, and industrial waste-management systems [28].

Table 1. Risk-based classification of industrial chemical data

Data class	Illustrative records	Policy consequence of staleness	Recommended consistency posture
Class A: immediate safety control	Emergency shutdown commands, toxic-gas alarms, pressure excursions, interlock status	Near-zero tolerance for stale or conflicting values	Local fail-safe protection independent of cross-region connectivity; replicated records are supplementary. During a network partition, local protective functions remain active and follow the predefined safety response. No autonomous weakening or dependence on remote acknowledgements for protective action.
Class B: emergency coordination	Incident logs, evacuation status, hazardous-material location, response-resource availability	Extremely low tolerance; bounded by emergency-response procedure	Strongly consistent authoritative coordination records with quorum-controlled geographic failover. During a partition, only a partition retaining quorum may commit authoritative updates; other replicas

			suspend authoritative writes. If quorum is unavailable, invoke predefined emergency communication procedures.
Class C: environmental compliance	Continuous emissions monitoring, discharge records, waste-transfer documentation, legally reportable inventories	Defined temporal bound linked to permit or reporting rule	Bounded-staleness mode with audit trail and automatic escalation
Class D: production and stewardship	Batch genealogy, quality release, maintenance status, product-composition, and stewardship data	Moderate and context-dependent tolerance	Bounded staleness; eventual consistency only for non-decisive views
Class E: analytics and planning	Historical dashboards, scenario models, aggregate sustainability analytics, non-operational reporting caches	Higher tolerance when source records remain authoritative	Supervised eventual consistency may be acceptable

The operational boundary between Classes A and B concerns the role of the distributed service during a network partition. Class A covers immediate protective functions, which must operate locally without waiting for cross-region agreement. Loss of connectivity should trigger the response defined by the facility's hazard assessment and safety requirements, including transition to a predefined safe state where required; it should not disable local alarms or interlocks. Replicated Class A records support monitoring and subsequent audit but do not authorize or delay protective action.

Class B covers shared emergency-coordination records rather than direct protective actuation. Geographic fallback means controlled transfer of authority, not simultaneous acceptance of conflicting authoritative updates in disconnected regions. For a deployment with one voting replica in each of three regions, a connected two-region majority may continue committing updates, while an isolated replica must not commit authoritative writes. Failover must prevent the former authority from issuing accepted updates.

If no majority is available, authoritative updates are suspended and personnel use predefined emergency communication procedures. Disconnected sites may preserve timestamped local observations for later reconciliation, but these must be identified as unconfirmed rather than presented as current shared state.

This restriction applies to the shared record and must not prevent local emergency response. These partition-handling rules are proposed deployment requirements; their implementation and effectiveness were not established by the simulation reported here.

2. Method

2.1. Study design and transparency boundary

The secondary policy analysis covers six scenarios representing complementary operating conditions: normal operation provides a baseline; burst demand examines temporary congestion; read-intensive skew and write-intensive demand examine workload composition; WAN degradation examines impaired communication; and partial communication failure examines reduced connectivity. These conditions are relevant to the reliability and efficiency of distributed industrial information services, although they were not calibrated using chemical-plant observations. The analysis interprets the reported output through economic, environmental, and safety-governance frameworks. Table 2 summarizes the data sources, analytical procedures, and evidence boundaries.

Table 2. Technical evidence base used in the secondary policy analysis

Component	Description
Data source	Numerical outputs from the authors' previously developed three-region discrete-event simulation, with five consistency policies and 30 stochastic replications per policy–scenario combination.
Scenarios covered	Normal operation, burst demand, read-intensive skew, write-intensive demand, WAN degradation, and partial communication failure.
Rationale for scenario coverage	Examine baseline operation, temporary congestion, differences in workload composition, and communication impairment relevant to distributed industrial information services.
Evidence included	Reported latency, throughput, service-level violations, staleness, inter-region communication volume, and switching behavior relevant to the policy questions.
Economic analysis	Interpret delay, service-level violations, and communication demand using the proposed cost framework; do not assign monetary savings without organization-specific coefficients.
Environmental analysis	Compare communication volumes and calculate avoided data transfer; treat energy and emissions estimates as conditional on documented conversion coefficients.
Safety and governance analysis	Assess reported staleness against the simulation's 250 ms limit and interpret findings through the proposed industrial data classes and governance requirements.
Excluded conclusions	Direct electricity or carbon savings, quantified monetary benefits, and demonstrated chemical-plant safety outcomes, because these were not measured.
Evidence limitations	Simulation-based findings; reconstructed communication totals are estimates; governance recommendations require validation before deployment.
Data source	Numerical outputs from the authors' previously developed three-region discrete-event simulation, with five consistency policies and 30 stochastic replications per policy–scenario combination.
Scenarios covered	Normal operation, burst demand, read-intensive skew, write-intensive demand, WAN degradation, and partial communication failure.
Rationale for scenario coverage	Examine baseline operation, temporary congestion, differences in workload composition, and communication impairment relevant to distributed industrial information services.
Evidence included	Reported latency, throughput, service-level violations, staleness, inter-region communication volume, and switching behavior relevant to the policy questions.
Economic analysis	Interpret delay, service-level violations, and communication demand using the proposed cost framework; do not assign monetary savings without organization-specific coefficients.
Environmental analysis	Compare communication volumes and calculate avoided data transfer; treat energy and emissions estimates as conditional on documented conversion coefficients.
Safety and governance analysis	Assess reported staleness against the simulation's 250 ms limit and interpret findings through the proposed industrial data classes and governance requirements.
Excluded conclusions	Direct electricity or carbon savings, quantified monetary benefits, and demonstrated chemical-plant safety outcomes, because these were not measured.
Evidence limitations	Simulation-based findings; reconstructed communication totals are estimates; governance recommendations require validation before deployment.
Data source	Numerical outputs from the authors' previously developed three-region discrete-event simulation, with five consistency policies and 30 stochastic replications per policy–scenario combination.
Scenarios covered	Normal operation, burst demand, read-intensive skew, write-intensive demand, WAN degradation, and partial communication failure.

The six evaluated scenarios cover distinct operating conditions relevant to distributed industrial information services. Normal operation provides a baseline; burst demand examines temporary congestion; read-intensive skew and write-intensive demand distinguish workload composition effects; WAN degradation examines impaired communication; and partial communication failure examines reduced connectivity. Their methodological value lies in covering complementary sources of operational stress. Their relevance to chemical industries is interpretive: they represent generic digital-service conditions rather than empirically calibrated chemical-plant workloads.

The article uses a policy-oriented secondary analysis of numerical outputs generated by a previously developed discrete-event simulation. The underlying technical model represented a database replicated across three geographical regions and compared five policies: static strong consistency, static bounded-staleness consistency, static eventual consistency, a threshold-based adaptive policy, and a safety-constrained Q-learning policy. Thirty stochastic replications were used for each policy–scenario combination in the source analysis.

The present article does not claim a new production deployment, chemical-plant experiment, direct electricity measurement, or monetary cost study. The simulation evidence is retained only where it supports policy reasoning. No original figure, literature-review table, or full technical derivation is reproduced. The contribution is the economic–environmental–regulatory interpretation, the data-class framework, the policy equations, and the implementation recommendations.

This boundary is important for publication ethics. If the technical manuscript is submitted or published elsewhere, the policy manuscript should be disclosed to the editor as a related work and should cite the technical publication when bibliographic details become available.

2.2. Technical evidence retained for policy analysis

The simulation included normal operation, burst demand, read-intensive skew, write-intensive demand, wide-area network degradation, and partial communication failure. The controller observed workload, read–write composition, latency, replication lag, staleness, service-level violations, network delay, and the previous consistency mode. The controller evaluated consistency actions at 5 s intervals [29]. separate safety supervisor rejected actions predicted to exceed the maximum staleness limit of 250 ms Hysteresis, minimum holding time, and rollback were used to reduce unstable switching. Q-learning was selected because it supports sequential decisions under uncertain state transitions [30], while the external supervisor reflects safe-reinforcement-learning practice in which constraints are not left to the reward function alone [31].

Table 3. Technical evidence base used in the secondary policy analysis

Element	Specification
Regions and replication factor	3 regions; replication factor 3
Decision interval	5 s
Evaluation duration	3,600 s, with 300 s warm-up
Maximum staleness limit	250 ms
Consistency actions	Strong, bounded staleness, eventual
Comparison policies	Static strong, static bounded, static eventual, threshold adaptive, proposed Q-learning
Independent replications	30 per policy–scenario combination
Reported policy indicators	p99 latency, throughput, SLA violations, peak staleness, inter-region communication, switching

The burst results describe the staleness indicator as “mean peak staleness.” A mean of run-level maxima must be distinguished from the maximum observed across all 30 runs: an average below 250 ms does not establish that every run remained below this threshold. Verification of observed compliance requires the maximum staleness across all runs and the number of runs containing at least one exceedance. Consequently, the reported summary values alone cannot substantiate a claim that the staleness limit was never exceeded.

2.3. Economic policy model

The economic model treats consistency as a portfolio of operational burdens rather than a single cloud-computing expense. For decision interval t, total policy cost is defined as the sum of network, delay, service-level, switching, and compliance components:

$$Ctotal(t) = Cnet(t) + Cdelay(t) + CSLO(t) + Cswitch(t) + Ccomp(t). \quad (1)$$

The terms can be monetized when organization-specific parameters are available. Network cost is represented as the product of inter-region data volume and the applicable transfer price:

$$Cnet(t) = pV \times VWAN(t). \quad (2)$$

Delay cost is expressed as the number of affected requests multiplied by the exceedance above the target latency and the estimated cost per unit of delay. Delay cost is represented by a tail-latency proxy that scales the exceedance of p99 latency above the target by the request volume and a monetary weighting coefficient:

$$Cdelay(t) = pL \times Nreq(t) \times max[0, Lp99(t) - Ltarget]. \quad (3)$$

Here, $Nreq(t)$ is the total number of requests in interval (t), and pL is a monetary weighting coefficient expressed in currency per request per unit of latency. Because p99 latency characterizes the tail of the latency distribution rather than the delay experienced by each request, Equation (3) is a comparative policy-cost proxy, not a measurement of aggregate delay cost. Estimating aggregate delay cost would instead require request-level latency exceedances and an appropriate valuation of their operational consequences.

Compliance cost may include investigation, correction, reporting delay, contractual consequences, or regulatory exposure. Because those values depend on jurisdiction and facility, the manuscript does not assign monetary figures. Instead, it identifies the conditions under which the term becomes non-zero.

2.4. Environmental proxy model

Communication volume is converted into energy and greenhouse-gas proxies only when local coefficients are supplied. The network-energy proxy is:

$$Enet(t) = VWAN(t) \times \varepsilon net \quad (4)$$

where εnet is the measured or contractually documented electricity intensity of transferred data within the defined system boundary. The associated greenhouse-gas proxy is:

$$GHGnet(t) = Enet(t) \times EFgrid(t) \quad (5)$$

where $EFgrid(t)$ is the relevant location- and time-dependent electricity emission factor. Equations (4) and (5) are intentionally conditional. Without credible εnet and $EFgrid(t)$ values, the appropriate result is avoided communication volume rather than avoided emissions.

2.5. Chemical-safety and compliance-risk model

For data class k, expected compliance and safety risk is represented as the probability of an unacceptable information state multiplied by its impact and an enforceability factor:

$$Rchem, k = P(Sk > Smax, k) \times Ik \times Fk \quad (6)$$

Sk is observed staleness, $Smax, k$ is the class-specific maximum, Ik represents the consequence of using outdated information, and Fk represents legal, contractual, or organizational enforceability. The admissibility condition for an adaptive action a in state x is:

$$Accept(a \mid x, k) \; only \; if \; \hat{S}upper(x, a) \leq Smax, k \quad (7)$$

Equation (7) screens only actions already permitted for the relevant data class; it does not itself establish permission to weaken consistency. For Class A, weakening actions must be excluded from the controller’s authorized action set and independently blocked at the architectural level, even when their predicted staleness satisfies the numerical bound.

2.6. Multi-criteria policy score

A normalized policy-suitability index is proposed for comparing admissible configurations. Higher throughput receives a positive weight, while latency, communication volume, service-level violations, staleness risk, and policy instability receive negative weights:

$$PSI = wT \times Tnorm - wL \times Lnorm - wV \times Vnorm - wO \times Onorm - wR \times Rnorm - wW \times Wnorm \quad (8)$$

Here, $Vnorm$ is normalized inter-region communication volume, $Onorm$ is normalized SLO-violation rate, $Rnorm$ is normalized chemical-safety and compliance risk, and $Wnorm$ is normalized switching frequency. The weights must be approved by the organization and documented for each data class. The index is applied only after non-negotiable safety and legal constraints have been satisfied. This prevents a large performance gain from compensating mathematically for an unacceptable chemical-safety outcome.

3. Results of the Policy-Oriented Reanalysis

3.1. Burst-demand evidence

During the 600-second high-load interval, static strong consistency produced a p99 latency of 438.58 ms, throughput of 1,194.75 operations per second, and a latency-service violation rate of 9.13%. The safety-constrained controller produced 303.25 ms, 1,434.79 operations per second, and 2.23%, respectively.

These differences correspond to a 30.86% reduction in $p99$ latency, a 20.09% increase in throughput, and a 75.56% reduction in service-level violations.

The proposed controller produced the same aggregate burst-period results as static bounded-staleness consistency. This is an important policy finding rather than a weakness to conceal. Under the defined 250 ms limit, bounded staleness was the highest-performing admissible policy throughout the burst. Static eventual consistency was faster, but its mean peak staleness reached 275.55 ms, 10.22% above the limit. The threshold-based policy reached 257.14 ms, 2.86% above the limit.

The policy interpretation is that unconstrained performance optimization is inappropriate for regulated records. A configuration may be economically attractive in terms of latency and throughput yet remain inadmissible because it crosses a data-freshness boundary linked to safety or legal reporting.

The 600-second burst interval used for the latency and throughput comparisons in Section 4.1 is distinct from the complete burst-scenario run used for communication-volume reporting. The estimated avoided volume of 3.210 GB refers to the complete burst run, whereas 9.915 GB refers to the WAN-degradation run.

Table 4. Policy interpretation of the 600-second burst interval

Policy	$p99$ latency (ms)	Throughput (ops/s)	SLA violations (%)	Peak staleness (ms)	Policy interpretation
Static strong	438.58	1,194.75	9.13	43.22	Admissible but high coordination and delay burden
Static bounded staleness	303.25	1,434.79	2.23	192.57	Admissible under 250 ms limits
Static eventual	158.28	1,590.14	0.10	275.55	Not admissible for data governed by the 250 ms limits
Threshold adaptive	265.65	1,450.76	2.94	257.14	Marginal limit exceedance; requires stronger guardrail
Safety-constrained Q-learning	303.25	1,434.79	2.23	192.57	Admissible; selected bounded staleness during the burst

3.2. Communication-volume implications

Across the six scenarios, the proposed controller reduced inter-region communication volume by approximately 31.4 – 60.0 % relative to static strong consistency. The largest proportional reduction occurred in the read-intensive skew scenario, where a larger share of reads could be served without synchronous cross-region coordination. The largest absolute avoided volume occurred during WAN degradation because strong coordination generated high communication and retransmission demand.

Table 4. Derived communication-volume implications relative to static strong consistency

Scenario	Adaptive volume (GB/run)	Reported reduction (%)	Estimated strong-consistency volume (GB/run)	Avoided volume (GB/run)
Normal operation	4.758	40.0	7.930	3.172
Complete burst run	5.443	37.1	8.653	3.210
Read-intensive skew	4.511	60.0	11.278	6.767
Write-intensive workload	7.138	34.6	10.914	3.776
WAN degradation	14.873	40.0	24.788	9.915
Partial communication failure	7.847	31.4	11.439	3.592

The strong-consistency volumes in Table 5 are back-calculated from the adaptive volume and reported percentage reduction. They should be replaced by direct raw-output values if the simulation archive is prepared for publication. The avoided volumes are not reported as electricity or carbon savings. Equations (4) and (5) show how a facility or cloud provider could complete that conversion using auditable coefficients.

3.3. Scenario-specific policy implications

Normal operation: The controller reduced p99 latency by 33.7% and communication volume by about 40% relative to static strong consistency. Because workload conditions were stable, the controller remained conservative rather than pursuing the absolute minimum latency. For environmental and economic policy, this behavior is desirable when the marginal benefit of further weakening is small.

Read-intensive skew: The controller achieved 113.98 ms p99 latency and 4.511 GB/run communication volume while peak staleness reached 240.30 ms. This scenario demonstrates the largest performance opportunity but also the smallest freshness margin. It is suitable for Class D or Class E workloads only when the safety supervisor uses conservative estimates and continuous monitoring.

Write-intensive workload: The controller remained in bounded-staleness mode and matched the static bounded policy. The absence of switching indicates that adaptation does not require constant change. A well-governed controller may determine that the current policy is already the best admissible option.

WAN degradation: The adaptive policy reduced communication volume by about 40% relative to strong consistency while preserving the staleness limit. This result is economically important in cloud or cross-border deployments in which network impairment increases retry, timeout, and transfer costs.

Partial communication failure: The controller accepted lower throughput and higher latency to preserve freshness, using safety and rollback mechanisms. For chemical emergency systems, the willingness to sacrifice performance is a required feature, not an optimization failure.

3.4. Economic interpretation

The burst results suggest that adaptive consistency can reduce congestion-related delay without automatically selecting the weakest policy. A 30.86% reduction in p99 latency and a 75.56% reduction in service-level violations can lower the probability of timeout, retry, delayed workflow completion, and contractual non-performance. The actual monetary value depends on request criticality and operational context.

Communication reductions can also lower network-transfer and cloud-egress costs when providers charge by transferred data. For example, the estimated avoided volume during WAN degradation was approximately 9.915 GB per evaluated run. Multiplying that figure by the organization's transfer price provides the network component in Equation (2). The result should not be generalized across providers because pricing, private-network contracts, and data-compression policies differ.

The largest economic risk is not necessarily the largest bandwidth bill. For Class A–C records, a freshness violation can create investigation, remediation, reporting, production-interruption, or liability costs that dominate ordinary infrastructure expenses. This is why the compliance term in Equation (1) must operate as a constraint or high-priority penalty rather than a minor efficiency adjustment.

3.5. Environmental and chemical-policy interpretation

The communication results support a narrow environmental conclusion: the adaptive controller required less cross-region data transfer than static strong consistency in every tested scenario. This is consistent with research showing that replication and synchronization choices affect bandwidth and infrastructure activity [16,18]. It does not demonstrate a measured reduction in facility electricity or greenhouse-gas emissions.

For chemical policy, the most useful outcome is the data-class model. Environmental permits, hazardous-material inventories, and accident-response records can receive explicit staleness limits, while non-operational analytics can use a more flexible mode. The revised Industrial Emissions Directive and corporate sustainability reporting requirements increase the value of traceable and timely environmental data, but they do not imply that every replica must be synchronously coordinated for every analytical request [32].

Digital product passports and lifecycle chemical information create an additional application. Product stewardship systems may contain both legally authoritative records and read-heavy public or supply-chain views. The authoritative record can remain strongly or boundedly consistent, while derivative catalog and analytics views can be replicated more flexibly. This separation supports data access and product sustainability without weakening the source of truth [33].

3.6. Discussion

3.6.1. Adaptive consistency as governance, not automatic optimization

The central policy lesson is that adaptive consistency should be treated as a governed decision system. The controller does not merely choose a faster database setting; it decides how much temporary divergence is acceptable for a specific industrial purpose. That decision requires authority, accountability,

monitoring, and documented limits.

The NIST AI Risk Management Framework distinguishes governance, mapping, measurement, and management functions for trustworthy AI systems [34]. A similar structure is appropriate here. Governance defines data classes and non-negotiable constraints. Mapping identifies stakeholders and failure consequences. Measurement covers latency, staleness, communication, overrides, and near misses. Management defines fallback, rollback, human intervention, audit, and periodic review.

A controller should therefore be unable to weaken a Class A or Class B policy solely because a reward function predicts lower latency. For Class C records, the system may switch only within a documented bond. For Class D and E workloads, the controller can receive greater flexibility, provided that the authoritative source remains identifiable and reconciliation procedures are assessed.

3.6.2. Regulatory alignment

REACH, industrial-emissions regulation, chemical-accident guidance, sustainability reporting, and emerging product-information rules share an underlying requirement: organizations must know which information is authoritative, who can access it, how it is updated, and how errors are corrected. Adaptive consistency can support these objectives only when it is embedded in records-management and compliance architecture [32].

The OECD guidance places particular emphasis on accident prevention, emergency preparedness, communication, incident investigation, consequence assessment, and sharing lessons learned [3] . These functions require different temporal guarantees. An alarm or isolation status requires immediate reliability; an investigation repository requires integrity, provenance, and completeness; a lessons-learned database may tolerate delayed synchronization as long as the final record is accurate and traceable.

The UNEP Global Framework on Chemicals emphasizes lifecycle management and multi-stakeholder participation [35]. In practice, this means that data may cross organizational and national boundaries. The Data Act and related governance instruments reinforce the need for access, portability, contractual clarity, and interoperability. Consistency policy should therefore be included in data-sharing agreements rather than left as an undocumented internal engineering decision.

3.6.3. Relationship to sustainable chemistry and environmental analytics

Environmental and chemical decision-making increasingly uses machine learning, lifecycle assessment, digital monitoring, and data-driven screening [26,36]. The quality of these analyses depends not only on model design but also on the timing, completeness, and provenance of the underlying records. An advanced carbon-footprint model cannot compensate for an undocumented or inconsistent source database.

Chemical sustainability metrics are most useful when they integrate safety, health, waste, environmental performance, and economic practicality. The solvent eco-impact approach, for example, illustrates the value of combining several operational dimensions rather than relying on one environmental indicator [37]. The PSI in Equation (8) follows the same principle for digital infrastructure, but it retains a strict distinction between weighted performance criteria and non-negotiable safety limits.

The proposed framework also complements absolute environmental sustainability approaches that define a safe operating space for chemical pressure [25]. In both cases, optimization is subordinate to an explicit boundary. The analogy should not be overstated the units and consequences are different but the policy logic is similar: improved efficiency does not justify crossing a defined protection threshold [38].

3.6.4. Organizational implementation model

A regulated organization can implement adaptive consistency through a five-stage governance cycle.

- Inventory and classify data. Identify authoritative records, users, legal obligations, safety consequences, geographical locations, retention rules, and acceptable recovery times.
- Define service objectives by class. Specify maximum staleness, latency target, availability requirement, reconciliation time, and whether automated weakening is permitted.
- Validate offline. Train and evaluate the controller using workload, network, failure, and attack scenarios that are broader than the expected operating range. Record state-action coverage and uncertainty.
- Deploy with supervision. Use deterministic evaluation, independent safety filters, minimum holding periods, rollback, human override, and immutable decision logs.
- Audit and revise. Review overrides, limit exceedances, near misses, cost performance, energy coefficients, and regulatory changes. Reapprove weights and thresholds after material system changes.

Table 6. Recommended implementation and assurance roadmap

Stage	Go/no-go criterion
1. Governance approval	Go: Data owners, process-safety personnel, and compliance representatives approve the data classification, staleness limits, performance targets, fallback procedures, and decision authority. No-go: Any safety-critical requirement or responsibility remains unresolved.
2. Offline validation	Go: All mandatory workload and failure tests pass against predefined acceptance criteria; supervisor rejection, rollback, and fallback behavior are verified. No-go: An unresolved unsafe action, failed fallback, or material gap in test coverage remains.
3. Controlled pilot	Go: The predefined observation period and minimum workload coverage are completed; no observed class-specific staleness-limit breaches or unauthorized weakening occur; every required fallback and override test passes; decision-log completeness is 100%; and the predefined performance target is met relative to the static baseline. No-go: Any safety gate fails or the evidence is insufficient to assess the acceptance criteria.
4. Supervised deployment	Go: Monitoring confirms continued compliance with approved limits, complete decision logging, and successful scheduled recovery tests; incidents are resolved and the responsible functions approve continued operation or expansion. No-go: A limit breach, failed protective mechanism, or unresolved material incident requires suspension of adaptation and use of the approved fallback.
5. Periodic assurance	Go: The audit confirms that thresholds, data classifications, model behavior, and fallback procedures remain appropriate; material findings are closed and authorization is renewed. No-go: Unresolved material findings or significant unvalidated changes prevent renewal.

The transition criteria in Table 6 are proposed governance requirements, not outcomes demonstrated by the simulation. Before a controlled pilot begins, the functions responsible must approve its duration, minimum request volume, workload and failure coverage, class-specific staleness limits, and performance acceptance thresholds. Pilot assessment should report the maximum observed staleness, the number and proportion of observations exceeding the applicable limit, unauthorized policy changes, fallback and override test success rates, and decision-log completeness. Log completeness is defined as the proportion of controller decisions for which all mandatory audit fields are recorded. Performance changes relating to the static baseline should be reported with confidence intervals and assessed against the predefined target. Progression requires all safety and assurance gates to pass; performance gains cannot compensate for a failed safety gate. Zero observed violations during a pilot provide evidence only for the conditions tested and does not establish a universal safety guarantee.

3.6.5. Limitations

The evidence is simulation-based and does not demonstrate production performance in a commercial database, chemical plant, or regulatory platform. Processing times, message behavior, and network disturbances were controlled assumptions. The policy analysis should therefore guide pilot design rather than justify immediate deployment.

The environmental assessment is limited to transferred data volume. Electricity, cooling, embodied infrastructure, equipment utilization, and marginal grid emissions were not measured. Communication reduction may correlate with lower resource demand, but the strength of that relationship is deployment-specific.

The economic model is parametric. No cloud price, labor cost, production-loss function, fine, or liability value was assigned. Organizations should populate Equations (1)–(3) with their own auditable values and perform sensitivity analysis.The safety supervisor used learned staleness distributions and operational guardrails, not formal verification. Small, realized exceedances may remain possible when prediction is imperfect. Safety-critical chemical control should continue to rely on local fail-safe design and certified protection layers; database consistency control is not a substitute for process safety engineering. The policy recommendations reflect international and European instruments but do not constitute legal advice. Jurisdiction-specific requirements, sector permits, data-localization rules, cybersecurity obligations, and contractual terms must be assessed before implementation.

3.6.6. Policy Recommendations

• Adopt a formal industrial data-class policy. Every replicated dataset should have an owner, authoritative source, maximum staleness, availability target, reconciliation rule, and emergency fallback.

• Enforce an architectural prohibition on autonomous weakening for Class A data. Within the proposed framework, this prohibition is a mandatory system-design requirement, not a configurable policy preference or reward penalty. The learning controller must have no authority to relax Class A consistency requirements, bypass local protective functions, or modify their safety limits. Enforcement must reside outside the learning component through restricted control interfaces and independently enforced permissions, with hardware separation where required by the facility's safety design. Local protective action must remain independent of cross-region connectivity.

- Use bounded staleness as the default for environmental-compliance services. The bound should be linked to permit, reporting, inspection, or operational requirements rather than selected only from technical benchmarks.
- Reserve eventual consistency for derivative and low-consequence workloads. Analytical caches, historical dashboards, and public catalog views may use weaker consistency when they cannot alter the authoritative record.
- Require an independent safety supervisor. The supervisor should use conservative uncertainty estimates, insufficient-data rules, rollback, and a safe fallback that the learning component cannot override.
- Create an immutable policy-decision log. Each switch should record the state, proposed action, accepted action, predicted staleness, active limit, override reason, and recovery outcome.
- Measure environmental performance correctly. Report avoided communication volume first; convert to energy and emissions only with documented network-energy and electricity-emission coefficients.
- Integrate consistency policy into data-sharing contracts. Cross-border industrial and chemical data agreements should specify source-of-truth rules, permitted divergence, reconciliation time, access rights, and incident communication.
- Pilot in shadow mode. Before enforcement, run the controller alongside the existing policy and compare recommendations, safety-filter interventions, and realized outcomes.
- Disclose related technical work. If the underlying simulation is published separately, the policy article should cite it and clearly describe the secondary nature of the analysis.

4. Conclusion

This secondary analysis positions adaptive data consistency as a governance decision for regulated chemical industries. Its principal contribution is a framework linking permissible data staleness to operational consequences, organizational responsibility, and assurance requirements. Across the evaluated scenarios, the adaptive controller reduced inter-region communication volume by approximately 31.4–60.0% relative to static strong consistency, indicating an opportunity to reduce coordination burden. These reductions do not establish electricity, carbon, or monetary savings without deployment-specific measurements and coefficients.

The policy implication is that consistency requirements should follow the purpose and consequences of the data. Immediate protective functions must remain locally operable, with autonomous weakening prohibited through architectural controls. Emergency-coordination records require controlled authority and explicit partition-handling procedures. Environmental-compliance and operational records require documented freshness limits, while greater flexibility may be appropriate for low-consequence analytical views.

Deployment should therefore depend on independent action screening, auditable decision records, tested fallback procedures, human oversight, and explicit go/no-go criteria. Performance improvements cannot compensate for failure to meet an applicable safety requirement. The simulation results inform this framework but do not demonstrate chemical-plant safety or establish universally valid staleness limits.

Future validation should combine reproducible controller specifications and run-level uncertainty analysis with industrial pilots, direct resource measurements, and stakeholder review. The practical objective is to reduce unnecessary coordination while preserving authoritative information, local protective capability, and regulatory accountability.

4.1. Transparency and Prior-Work Statement

This manuscript is written as a distinct policy-oriented secondary analysis of simulation outputs developed by the authors. It is not intended to conceal textual or substantive overlap with a related technical manuscript. Before submission, the authors should disclose any related manuscript to the target journal and replace this statement with a citation to the technical publication if it has been published or accepted. No production of chemical-plant data, human-subject data, or direct energy measurements are reported.

4.2. Data and Code Availability

The numerical values summarized in this article originate from the authors' simulation study. For reproducibility, the final submission should be accompanied by the scenario definitions, random seeds, policy parameters, result-processing scripts, and raw policy–scenario outputs in a public repository or journal supplementary archive. Any proprietary information should be removed without altering the reported calculations.

Acknowledgments

The authors would like to thank the reviewers and the Editor for their valuable contributions to enhancing the scientific quality of this research.

Author Contributions

Hadi Davardoost: Conceptualization, methodology, formal analysis, and writing—original draft. Julia Peryazeva: Methodology, validation, supervision, writing—review and editing. Alan Las Jamal Hawezy: Economic and policy interpretation and writing—review and editing. All authors read and approved of the final manuscript.

Conflict of Interest

The authors, while observing publication ethics in the referencing, declare the absence of interest of conflict.

AI Use Statement

The author retains full responsibility for the accuracy, originality, integrity, legality, and ethical compliance of all text, data, tables, figures, analyses, references, and conclusions, and have verified all AI-assisted outputs. Any fabrication, falsification, manipulation, plagiarism, copyright infringement, or other misuse of AI remains solely the responsibility of the authors and does not transfer or diminish their accountability for the published work.

References

[1] O. Afanaseva, M. Afanasyev, S. Neyrus, D. Pervukhin, D. Tukeev, *Inventions*, 2024, **9(6)**, 1151-20. https://doi.org/10.3390/inventions9060115

[2] D. Pervukhin, L. Tang, *Operational Research in Engineering Sciences: Theory and Applications*, 2024, **7(3)**,144-163. https://doi.org/10.5281/zenodo.15093139

[3] Inter-Organization Programme for the Sound Management of Chemicals, *OECD Guiding Principles for Chemical Accident Prevention, Preparedness and Response*, Paris: OECD Publishing, 2003. https://www.oecd.org/en/publications/oecd-guiding-principles-for-chemical-accident-prevention-preparedness-and-response_9789264101821-en.html

[4] European Parliament and the Council of the European Union, *Regulation (EU) 2024/1781 of the European Parliament and of the Council of 13 June 2024 Establishing a Framework for the Setting of Ecodesign Requirements for Sustainable Products, Amending Directive (EU) 2020/1828 and Regulation (EU) 2023/1542 and Repealing Directive 2009/125/EC (Text with EEA Relevance)*, Luxembourg: Publications Office of the European Union, 2024. http://data.europa.eu/eli/reg/2024/1781/oj

[5] O. Afanaseva, S. Neyrus, V. Navatskaya, A. Perezhogina, *Risk Assessment of Investment Projects Using the Simulation Decomposition Method*. Springer Nature Switzerland, 2023, pp. 776-785. https://doi.org/10.1007/978-3-031-36960-5_88

[6] D. Pervukhin, H. Davardoost, D. Kotov, Y. Ilyukhina, K. Hasanov, *Advanced Mathematical Models & Applications*, 2023, **8(3)**,502-528.

[7] D. Pervukhina, H. Davardoost, E. Gasimovb, A. Hawezyc, *International Journal of Engineering*, 2025, **38(8)**, 1909-1921. https://doi.org/10.5829/ije.2025.38.08b.15

[8] J. C. Corbett, J. Dean, M. Epstein, A. Fikes, C. Frost, J. J. Furman, S. Ghemawat, A. Gubarev, C. Heiser, P. Hochschild, *ACM Transactions on Computer Systems (TOCS)*, 2013, **31(3)**, 1-22. https://doi.org/10.1145/2491245

[9] G. DeCandia, D. Hastorun, M. Jampani, G. Kakulapati, A. Lakshman, A. Pilchin, S. Sivasubramanian, P. Vosshall, W. Vogels, *ACM SIGOPS Operating Systems Review*, 2007, **41(6)**, 205-220. https://doi.org/10.1145/1294261.1294281

[10] P. Bailis, , S. Venkataraman, M. J. Franklin, , J. M. Hellerstein, I. Stoica, *Communications of the ACM*, 2014, **57(8)**, 93-102. https://doi.org/10.1145/2632792

[11] P. Viotti, M. Vukolić, *ACM Computing Surveys (CSUR)*, 2016, **49(1)**, 1-34. https://doi.org/10.1145/2926965
[12] M. R. Rahman, L. Tseng, S. Nguyen, I. Gupta, N. Vaidya, *ACM Transactions on Autonomous and Adaptive Systems (TAAS)*, 2017, **11(4)**, 1-36. https://doi.org/10.1145/299765
[13] S. Sidhanta, W. Golab, S. Mukhopadhyay, S. Basu, *2016 16th IEEE/ACM International Symposium on Cluster*, Cloud and Grid Computing (CCGrid), Cartagena, Colombia, 2016. https://doi.org/10.1109/CCGrid.2016.9
[14] International Energy Agency (IEA), *Energy and AI*, Paris: IEA, 2025. https://www.iea.org/reports/energy-and-ai
[15] E. Masanet, A. Shehabi, N. Lei, S. Smith, J. Koomey, *Science*, 2020, **367(6481)**, 984-986. https://doi.org/10.1126/science.aba3758
[16] D. Boru, D. Kliazovich, F. Granelli, P. Bouvry, A. Y. Zomaya, *Cluster Computing*, **18(1)**, 2015, 385-402. https://doi.org/10.1007/s10586-014-0404-x
[17] V. Fetisov, A. M. Gonopolsky, E. A. Mazlova, R. M. Behbahani, H. Davardoost, *Environmental Modeling & Assessment*, 2026, ***31*(1)**, 159-170. https://doi.org/10.1007/s10666-025-10077-3
[18] D. Xu, T. Li, Z. Sun, Z. Chen, W. Zhou, Y. Zhang, W. Lu, X. Du, *Proceedings of the ACM on Management of Data*, 2026, ***4*(1)**, 1-28. https://doi.org/10.1145/3786677
[19] F. Dashtroo, V. Varahrami, *Journal of Environmental Economics and Chemical Processes*, 2026, **3(2)**, 115-121. https://doi.org/10.22034/jeecp.2026.735188
[20] A. Emami Meybodi, F. Maralkhani Azar, *Journal of Environmental Economics and Chemical Processes*, 2025, **2(3)**, 103-106. https://doi.org/10.22034/jeecp.2025.735018
[21] R. Y. Wang, D. M. Strong, *Journal of Management Information Systems*, 1996, **12(4)**, 5-33. https://doi.org/10.1080/07421222.1996.11518099
[22] Organisation for Economic Co-operation and Development, Recommendation of the Council on Enhancing Access to and Sharing of Data, Paris: OECD Publishing, 2021. https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0463?
[23] M. Nasr Esfahani, R. Montazer, E. Mehranpour, *Journal of Environmental Economics and Chemical Processes*, 2025, ***2*(4)**, 143-150. https://doi.org/10.22034/jeecp.2025.735025
[24] S. A. Razavi, & M. Nasr Esfahani, *Political Strategic Studies*, 2020, **9(35)**, 101-134. https://doi.org/10.22054/qpss.2020.42777.2320
[25] M. B. Kosnik, M. Z. Hauschild, P. Fantke, *Environmental Science & Technology*, 2022, ***56*(8)**, 4776-4787. https://doi.org/10.1021/acs.est.1c06098
[26] D. Zhang, Z. Wang, C. Oberschelp, E. Bradford, S. Hellweg, *ACS Sustainable Chemistry & Engineering*, 2024, ***12*(7)**, 2700-2708. https://doi.org/10.1021/acssuschemeng.3c07038
[27] A. Silveira Sbrice Pinto, N. Gulpinar, E. A. Gibson, L. Fuller, P. Souter, *ACS Sustainable Resource Management*, 2026, **3(6)**, 1875-1886. https://doi.org/10.1021/acssusresmgt.6c00052
[28] H. Mofid, M. R. Khosravi-Nikou, S. A. Razavi, M. Mohammadian, *Applied Thermal Engineering*, 2024, **245**, 1228101-17. https://doi.org/10.1016/j.applthermaleng.2024.122810
[29] S. A. Razavi, S. Pirani, *Petroleum Business Review*, 2019, **3(4)**, 28-36. https://doi.org/10.22050/pbr.2019.119151
[30] C. J. Watkins, P. Dayan, *Machine Learning*, 1992, ***8*(3)**, 279-292. https://doi.org/10.1007/BF00992698
[31] J. Garcıa, F. Fernández, *Journal of Machine Learning Research*, 2015, ***16*(1)**, 1437–1480. https://dl.acm.org/doi/10.5555/2789272.2886795
[32] European Parliament and the Council of the European Union, *Directive (EU) 2022/2464 of the European Parliament and of the Council of 14 December 2022 Amending Regulation (EU) No 537/2014, Directive 2004/109/EC, Directive 2006/43/EC and Directive 2013/34/EU, as Regards Corporate Sustainability Reporting (Text with EEA Relevance)*, Luxembourg: Publications Office of the European Union, 2022. http://data.europa.eu/eli/dir/2022/2464/oj
[33] European Parliament and the Council of the European Union, *Directive (EU) 2024/1785 of the European Parliament and of the Council of 24 April 2024 Amending Directive 2010/75/EU of the European Parliament and of the Council on Industrial Emissions (Integrated Pollution Prevention and Control) and Council Directive 1999/31/EC on the Landfill of Waste (Text with EEA Relevance)*, Luxembourg: Publications Office of the European Union, 2024. http://data.europa.eu/eli/dir/2024/1785/oj
[34] E. Tabassi, *Artificial Intelligence Risk Management Framework (AI RMF 1.0), Gaithersburg*, MD: National Institute of Standards and Technology, 2023. https://doi.org/10.6028/NIST.AI.100-1
[35] I. Andersen, *A Planet Free of Harm from Chemicals and Waste*, United Nations Environment Programme, 2023. https://www.unep.org/en/publications/report/global-framework-chemicals-planet-free-harm-chemicals-and-waste
[36] Z. J. Ren, G. V. Lowry, W. A. Arnold, B. W. Brooks, P. Gago-Ferrero, J. M. Garcia, K. C. Leonard, M. Mills, J. F. Serrano, S. Wang, *ACS Sustainable Chemistry & Engineering*, 2022, ***10*(44)**, 14392-14393. https://doi.org/10.1021/acssuschemeng.2c06310
[37] M. Ferrara, F. Della Negra, M. Verzini, *Organic Process Research & Development*, 2025, ***29*(4)**, 1110-1124. https://doi.org/10.1021/acs.oprd.4c00546
[38] S. A. Razavi, H. Motavali, *Journal of Industrial Engineering and Management Studies*, 2023, **10(2)**, 106-115. https://doi.org/10.22116/jiems.2023.399150.1516