Journal of Secure Digital Technologies
Volume 1 (№2) September 2026

ISSN (E): 3136-9340
Sertificate № KZ19VPY00128988, date of issue: 25.12.2025

The journal is published 4 times a year
(March, June, September, December)

EDITORIAL TEAM
Editor-in-Chief
Dr. Leila Rzayeva, Association prof., PhD, Digital government support center (Kazakhstan)

Deputy Editor-in-Chief
Dr. Ali Myrzatay, PhD, Digital government support center (Kazakhstan)

Executive Secretary
Dr. Abdenova Zhuldyz, PhD, Digital government support center (Kazakhstan)

EDITORIAL BOARD MEMBERS (KAZAKHSTAN)
Seitkulov E.N. – Candidate of Physical and Mathematical Sciences, Professor Computer Science, Computer Engineering and Control
Abdenova G.A. – PhD in Automation and Control, Candidate of Technical Sciences, Professor L.N. Gumilyov Eurasian National University
Atanov S.K. – Doctor of Technical Sciences, Professor.
Sergazin G.K. – PhD, Associate Professor Computer Science, Computer Engineering and Control
Satybaldina D.Zh. – Candidate of Physical and Mathematical Sciences, Professor Director, Research Institute of Information Security and Cryptology, L.N. Gumilyov Eurasian National University

INTERNATIONAL EDITORIAL BOARD MEMBERS
Ibraheem Shayea – PhD in Electrical and Electronic Engineering Istanbul Technical University (Turkey)
Yadigar Nasiboglu Imamverdiyev – DSc in Technical Sciences, Associate Professor Information Security, Computer Science and Computer Technology (Azerbaijan)
Erdal Irmak – PhD in Electrical and Electronics Engineering, Professor Gazi University (Turkey)
Korhan Kayisli – PhD in Electrical Electronics Engineering, Associate Professor Gazi University (Turkey)
Yonghao Wang – PhD in Electronic Engineering, Associate Professor Queen Mary, University of London (UK)
Amir Mosavi – PhD in Applied Informatics, Professor of AI Obuda University (Hungary)
Gerald Feldman – PhD in General Computing, Associate Professor Birmingham City University (UK)
Mohamed Ragab – PhD in Computer Science, Assistant Professor Birmingham City University (UK)
Seref Sagiroglu – PhD in System Engineering, Professor Gazi University (Turkey)

MANAGING EDITOR
Dr. Ali Myrzatay, PhD, Digital government support center (Kazakhstan)

TECHNICAL EDITOR
Dr. Ali Myrzatay, PhD, Digital government support center (Kazakhstan)

CONTENTS

Cyber Incident Damage Assessment for Public Digital Systems and Critical Infrastructure 3

Cyber Incident Damage Assessment for Public Digital Systems and Critical Infrastructure 37

Verifiable but Not Correct: Blockchain Trust Boundaries for Autonomous Agents 71

Calibrated Digital Twins and Geometry-based Perception for Automated Driving in Heavy Rain and Snowfall: Evidence From Research Testing in China and Implications for Kazakhstan. 88

Personal Data Protection in Kazakhstan's AI Enabled Public Services 109

Asking Another Jurisdiction What It Holds Without Telling It What You Hold: Private Artifact Matching with a Ledger-Anchored Transcript 131

On-Chain and Off-Chain Data Placement in Hybrid Blockchain Health Systems 143

An Integrated Project Management Maturity Model for Digital Qazaqstan Initiatives: Embedding Cybersecurity, AI Governance and Forward-Looking Cyber-Risk Requirements into ST RK ISO 21502-2022..... 158

Appendix A. Indicators and evidence for the baseline practice groups 181

Privacy Risks and Data Protection Readiness in AI Marketing in Kazakhstan..... 183

Supplementary Materials..... 205

Leila Rzayeva
PhD, Associate Professor
e-mail: l.rzayeva@govtec.kz
ORCID: 0000-0002-3382-4685
Digital Government Support Center, Astana, Republic of Kazakhstan

Aina Tazhiyeva*
Master of Science, Researcher
e-mail: a.tazhiyeva@govtec.kz
ORCID: 0009-0008-3211-7840
Digital Government Support Center, Astana, Republic of Kazakhstan

Iskander Tazhibekov
Master of Business Administration
e-mail: i.tazhibekov@mdai.gov.kz
ORCID: 0009-0003-6517-9355
Information security Committee of the Ministry of Artificial Intelligence
and Digital Development of the Republic of Kazakhstan

CYBER INCIDENT DAMAGE ASSESSMENT FOR PUBLIC DIGITAL SYSTEMS AND CRITICAL INFRASTRUCTURE

Abstract: Cyber incident damage estimates can change materially when overlapping costs, normative coefficients and classification thresholds are combined without a common measurement boundary. This study integrates a comparative narrative review of 52 sources with formal accounting and sensitivity analysis for public digital services and critical infrastructure. An observation-to-loss model identifies the exact inflation caused by repeated items. A monotonicity criterion determines when a financial category is stable over a declared uncertainty set, while operational and safety escalation remain separate. The analytical conditions and reproducible calculations are applied to two demonstrations from unpublished national drafts and to a controlled two-sector propagation model. In the original portal case, the reputation convention accounts for 90.96% of the total and has elasticity 0.90956; the critical financial threshold is reached at 42.905% of its original magnitude. Removing the convention changes financial severity from critical to moderate without necessarily changing operational escalation. Exact scenario bounds are 20.8265–34.109 million KZT for the revised wider portal account and 67.912–83.720 million for the mine account. In the propagation model, a cap tying downstream loss to primary loss understates the total by 20% at coefficient 0.60 and 60% at 0.80. A thirteen-case reported-loss table is strongly concentrated and is used descriptively rather than for calibration. The contribution is a set of testable consistency and robustness conditions; independently verified incident-based validation remains necessary.

Keywords: cyber incident damage; loss estimation; public digital services; critical infrastructure; double counting; uncertainty; cascading effects; severity assessment

Оценка ущерба от киберинцидентов в государственных цифровых системах и критической инфраструктуре

Аннотация: Оценки ущерба от киберинцидентов существенно меняются при объединении пересекающихся затрат, нормативных коэффициентов и порогов классификации без единых границ измерения. В работе сравнительный нарративный обзор 52 источников дополнен формальным анализом учёта и чувствительности для государственных цифровых услуг и критической инфраструктуры. Модель связи исходных потерь с отчётными величинами определяет точное завышение из-за повторного учёта. На основе монотонности сформулировано условие устойчивости финансовой категории при заданной неопределённости; операционная тяжесть и безопасность рассматриваются отдельно. Аналитические условия и воспроизводимые расчёты применены к двум демонстрационным примерам из проектов национальной методики и условной двухсекторной модели распространения ущерба. В исходном портальном примере репутационная составляющая определяет 90,96% суммы и имеет эластичность 0,90956; критический финансовый порог достигается при 42,905% её исходной величины. Исключение составляющей снижает финансовую категорию с критической до средней, не обязательно изменяя операционный уровень. Точные сценарные границы составляют 20,8265–34,109 млн тенге для расширенной оценки портального случая и 67,912–83,720 млн для шахтного. Ограничение каскадных потерь величиной первичного ущерба занижает итог условной модели на 20% при коэффициенте 0,60 и на 60% при 0,80. Таблица тринадцати заявленных потерь используется для описательного анализа, а не калибровки. Результат работы — проверяемые условия согласованности расчётов и устойчивости классификации; необходима независимая проверка на первичных материалах инцидентов.

Ключевые слова: ущерб от киберинцидентов; оценка потерь; государственные цифровые услуги; критическая инфраструктура; двойной учёт; неопределённость; каскадные эффекты; оценка тяжести последствий

Мемлекеттік цифрлық жүйелер мен аса маңызды инфрақұрылымдағы киберинциденттерден келетін залалды бағалау

Аңдатпа: Киберинцидент залалының бағасы қайталанатын шығындар, нормативтік коэффициенттер және жіктеу шектері бірыңғай өлшеу шекараларынсыз біріктірілгенде елеулі өзгереді. Жұмыста 52 дереккөздің салыстырмалы нарративтік шолуы мемлекеттік цифрлық қызметтер мен аса маңызды инфрақұрылым үшін есепке алу мен сезімталдықтың формалды талдауымен толықтырылған. Бастапқы шығындарды есептік шамалармен байланыстыратын модель қайталап есептеуден туындайтын нақты артық соманы анықтайды. Монотондылық негізінде берілген белгісіздік жағдайында қаржылық санаттың тұрақтылық шарты тұжырымдалады; операциялық салдар мен қауіпсіздік бөлек қарастырылады. Аналитикалық шарттар және қайта шығарылатын есептер ұлттық әдістеме жобаларындағы екі демонстрациялық мысалға және шартты екі секторлы таралу моделіне қолданылады. Бастапқы портал мысалында беделдік құрамдас жалпы соманың 90,96%-ын құрайды, оның икемділігі 0,90956; критикалық қаржылық шек осы құрамдастың бастапқы шамасының 42,905%-ында орындалады. Оны алып тастау қаржылық санатты критикалықтан орташаға төмендетеді, бірақ операциялық деңгейді міндетті түрде өзгертпейді. Порталдың кеңейтілген бағасының нақты сценарийлік шектері 20,8265–34,109 млн теңге, шахта бағасының шектері 67,912–83,720 млн теңге. Каскадтық шығындарды бастапқы залалмен шектеу шартты модельдің жалпы сомасын коэффициент 0,60 болғанда 20%-ға, ал 0,80 болғанда 60%-ға кемітеді. Мәлімделген он үш шығын кестесі калибрлеуге емес, сипаттамалық талдауға қолданылады. Нәтиже — есептеулердің сәйкестігі мен жіктеудің тұрақтылығын тексеру шарттары; оларды инциденттердің бастапқы материалдары негізінде тәуелсіз тексеру қажет.

Түйін сөздер: киберинциденттерден келетін залал; шығындарды бағалау; мемлекеттік цифрлық қызметтер; аса маңызды инфрақұрылым; қосарланған есеп; белгісіздік; каскадтық әсерлер; салдардың ауырлығын бағалау

1 Introduction

Cybersecurity incidents can generate expenditure for the affected organisation, disrupt dependent services and impose costs on people who are not parties to the original system. An assessment intended to support recovery funding, regulatory reporting or national statistics must therefore specify whose loss is measured, over what period and on what evidence. IBM reports an average data breach cost of USD 4.99 million for its 2026 study [1]. This benchmark describes the surveyed population; it is not a valuation rule for a particular public digital service or industrial installation.

The distinction matters because cyber risk quantification and incident damage assessment answer different questions. Risk analysis estimates the distribution of possible future loss, whereas incident assessment reconstructs the consequences of an event that has

occurred. Both may use monetary models, but only the former necessarily includes event frequency. Woods and Böhme [2] emphasise the importance of causal interpretation in empirical cyber risk research. Cremer et al. [3] identify limitations in accessible loss datasets, while event-level studies demonstrate substantial variation in reported costs [4], [5]. These findings support transparent assumptions and uncertainty reporting rather than the direct transfer of a global mean to a national setting.

The relevant literature comprises several complementary traditions. Investment models relate security expenditure to expected loss reduction [6], [7]. ISO/IEC 27005 and NIST SP 800-30 establish risk assessment processes [8], [9]. Open FAIR provides a taxonomy for analysing loss frequency and magnitude [10]. NIS2 and DORA establish criteria for regulatory significance [11], [12]. Business continuity, operational technology safety and infrastructure interdependency research address consequences that extend beyond information assets. Composite indicator methods offer a way to summarise heterogeneous impacts, subject to explicit assumptions about scaling and weights [13]. Their outputs cannot be compared without first aligning the object of measurement.

For government systems, the central difficulty is not that existing methods are inapplicable to non-commercial organisations. Response expenditure, replacement costs and productivity losses can be estimated without sales revenue. The harder problem is separating these costs from interrupted public service, loss of strategic autonomy, safety consequences and effects on other sectors. A budget allocation is not automatically a loss, a service delay is not necessarily permanent loss of output, and a critical safety consequence need not produce a large financial bill. An assessment that obscures these distinctions can misdirect both compensation and recovery priorities.

This study combines a comparative review with formal consistency and sensitivity analysis. It asks which measurement objects the main method families support, how overlapping losses can be reconciled, and when uncertainty in valuation changes an incident's financial or overall severity. The analytical contribution has three parts: an explicit observation-to-loss reconciliation condition; a criterion for category stability over a declared uncertainty set; and exact parameter-response relations for two draft-method demonstrations and a controlled propagation model. The conditions are derived and numerically checked. They are not claims of a new general theory of risk or of validated national loss coefficients.

Section 2 describes the source-based comparison and reproducible analytical procedure. Section 3 derives the accounting and classification conditions, reconstructs the two demonstration cases and presents quantitative sensitivity results. Section 4 discusses their practical interpretation and validation requirements. Section 5 states the conclusions. S1 documents source identities and evidence links; S2 and the accompanying scripts reproduce the calculations and figures.

2 Review design and analytical method

2.1 Scope and source selection

The study uses a structured comparative narrative review of 52 identified publications and instruments, followed by formal and numerical analysis of selected measurement problems. Selection is purposive and covers loss definition, monetary estimation, regulatory significance, service and physical consequences, interdependencies and indicator construction. S1 lists every reference and its role. The review is not an exhaustive systematic search and does not estimate the prevalence of defects across the literature.

Research studies support empirical and modelling claims; standards define processes and terminology; legal instruments specify rules within their stated scope; and industry reports provide contextual benchmarks. Principal source claims were checked against issuing-body and publisher records on 21 October 2026. S1 distinguishes inspectable passages from catalogue descriptions and supporting references. The cited edition defines each instrument; full-text access to every proprietary source is not assumed.

Sources restricted to attack detection are excluded unless they connect an incident to measurable consequences. Two unpublished national drafts are objects of critical analysis: the earlier Appendix 2.6 and working revision 2.2 [14], [15]. S1 supplies their exact file identities and section locators. The earlier Appendix B defines the demonstration inputs; revision 2.2 changes accounting and classification rules and clarifies the status of thirteen reported-loss entries. These drafts are not independent validation sources or adopted national standards.

2.2 Units of analysis and comparison criteria

The source-level unit is a named instrument or study, identified in S1. Families group sources with related purposes. Table 1 gives representative members; Tables 2 and 3 identify source-linked measurement routes and consistency checks. No numerical source ratings are averaged into family scores. A capability is attributed to a specified source and use; an unreported feature remains undetermined rather than being scored as absent.

Table 1. Analytical families and representative sources

Family	Principal output	Representative anchors
Risk management processes	Risk assessment structure and decision criteria	ISO/IEC 27005 [8]; NIST SP 800-30 [9]; CSF 2.0 [16]
Monetary risk and investment models	Loss distributions, expectations or investment criteria	Open FAIR [10], [17], [18]; ROSI [7], [19]; CyVaR [20], [21]
Empirical incident studies	Observed costs and associations in a defined sample	IBM [1]; Eling and Wirfs [4]; Romanosky [5]; Aldasoro et al. [22]
Regulatory instruments	Reporting thresholds and consequence categories	NIS2 [11]; DORA [12]; GDPR [23]; delegated and implementing acts [24], [25]

Family	Principal output	Representative anchors
Continuity and OT safety	Service impact, recovery priorities and physical consequences	ISO 22301 [26]; ISO/TS 22317 [27]; IEC 62443-3-2 [28]; NIST SP 800-82 [29]
Interdependency models	Propagation of disruption between systems or sectors	Leontief [30]; Rinaldi et al. [31]; ENISA [32]
Composite indicators	A score or ranking under stated aggregation rules	OECD/JRC [13]; AHP [33]

Comparison uses four criteria: the meaning and units of the output; the treatment of uncertainty; the prevention of overlapping loss entries; and the declaration of the organisational, temporal and economic boundary. Nine consequence dimensions are examined: financial expenditure, material damage, operational disruption, workforce productivity, reputation, public service continuity, technological dependency, cross-sector propagation and human safety. These dimensions are descriptive tags. They are not nine amounts that can automatically be added, because one consequence can carry several tags.

2.3 Synthesis and reproducibility

The analysis links each comparative claim to a named source, an identifiable passage or public scope statement, and a stated inference. S1 distinguishes a source statement, an analytical inference and a numerical demonstration. The synthesis rule is non-statistical: a family description identifies a capability illustrated by its cited members; it neither averages ordinal ratings nor interprets one positive example as universal coverage. For example, NIST SP 800-30 discusses uncertainty explicitly, while the OECD/JRC handbook addresses both overlap and sensitivity. These observations preclude family-wide claims that process standards or composite indicator methods lack those features [9], [13].

The evidence map was assembled through AI-assisted extraction and editorial analysis with OpenAI ChatGPT. No independent human double coding or inter-rater reliability estimate is claimed. Conflicting readings are resolved by narrowing the interpretation to the passage that can be supported. S1 records the source-to-claim links and access limitations. This interpretive evidence is distinguished from the algebraic results and reproducible numerical calculations below.

2.4 Analytical and computational procedure

The analytical unit is an incident account with a fixed owner, valuation perspective, time horizon and currency basis. The procedure first reconstructs the draft’s original entries, then changes one accounting convention or a declared parameter set at a time. Loss components are retained separately so that changes in a total can be attributed to a specific assumption. Derivatives quantify local sensitivity; monotonicity gives exact interval endpoints; classification rules map the feasible ratios to severity bands. An explicit

symmetric two-sector system isolates the effect of propagation strength. Finally, descriptive statistics examine the concentration of the thirteen reported amounts without treating them as a representative sample.

Figures use deterministic grids: 301 reputation multipliers, two 101 × 101 parameter surfaces and 451 propagation coefficients. These grids visualise response functions; the interval bounds and threshold-crossing conditions follow analytically rather than from sampling. No random seed or probability distribution is required. Numerical checks compare original totals with their source, verify the fixed-point relation and interval endpoints, and compare an analytical elasticity with a central finite difference. The supplied CSV files contain every plotted point.

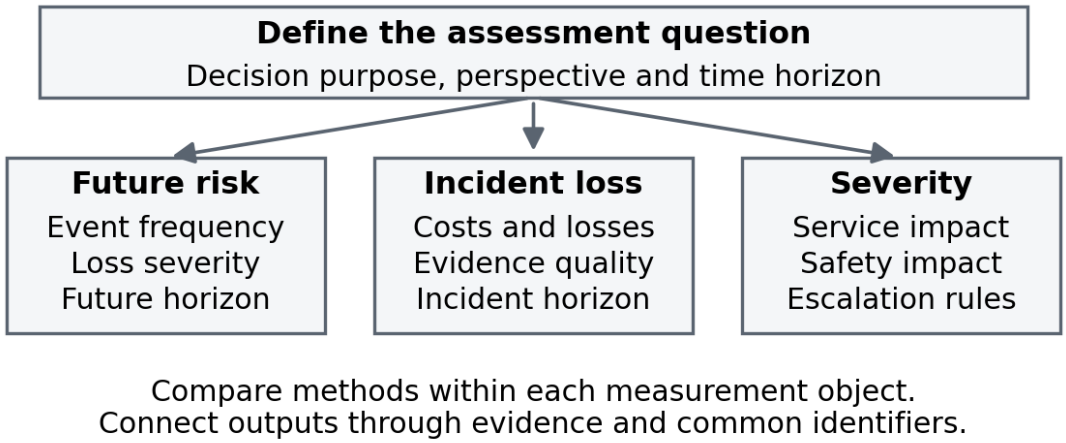


Figure 1. Measurement objects distinguished in the comparative analysis. Currency units alone do not make the outputs interchangeable.

3 Comparative analysis and assessment design

3.1 Risk management process standards

ISO/IEC 27005 defines guidance for information security risk management, while NIST SP 800-30 organises risk assessment around threats, vulnerabilities, likelihood and impact [8], [9]. CSF 2.0 places cybersecurity outcomes within six functions, including Govern [16]. ISO 31000 and IEC 31010 provide broader risk management principles and assessment techniques [34], [35]. These instruments support consistent processes, but their adoption alone does not establish common monetary assumptions across organisations.

The resulting limitation is one of specification. Two agencies may apply the same process and legitimately choose different impact scales, time horizons or valuation perspectives. Their risk assessments can be useful locally without being suitable for aggregation into national loss statistics. A national procedure should retain the governance

and documentation functions of these standards while separately defining the minimum incident accounting record. Lack of a universal currency formula is not a defect in a standard whose primary purpose is risk governance.

This distinction should not obscure the safeguards already present in process guidance. NIST SP 800-30, Section 2.3.1, discusses uncertainty; Section 2.3.2 distinguishes assessment scales; and Task 2-6 calls for uncertainty to accompany risk determinations [9]. The present contribution is the operationalisation of compatible incident accounting records, not the discovery of uncertainty as a missing concept.

3.2 Monetary loss and investment models

Open FAIR distinguishes loss event frequency from loss magnitude and organises losses through primary and secondary consequences [10], [17], [18]. Its loss forms can support identification of response, replacement, productivity and other costs in a public organisation. The primary–secondary distinction concerns the causal pathway, including reactions by stakeholders; it should not be treated as an exact synonym for direct–indirect accounting categories. The cited O-RT and O-RA editions are Version 3.1 and Version 2.1, respectively, published in 2025.

The FAIR Institute model specification independently identifies productivity, response and replacement loss forms in Sections 3.1–3.3 [36]. For a public agency, response and replacement expenditure remain meaningful even when forgone sales are zero. Salary-based productivity valuation requires evidence of displaced useful work and is not automatically an incremental cash expense. Reputation-related revenue loss may be inapplicable to a particular public service, but that limitation concerns a loss pathway, not the validity of the entire framework.

Annualised loss expectancy expresses expected aggregate loss over a defined horizon. Under the additional assumptions that event losses are identically distributed, have finite means and are independent of the event count, the compound loss model has the form in Equation (1).

$$\mathrm{E}[S_H] \;=\; \mathrm{E}\left[\sum_{k=1}^{N_H} X_k\right] \;=\; \mathrm{E}[N_H]\,\mathrm{E}[X] \quad (1)$$

Here S_H is aggregate loss over horizon H, N_H is the number of loss events and X_k is the loss of event k, all monetary quantities being expressed on a consistent price basis. Independence of count and severity should be checked rather than assumed for common-cause or cascading incidents. For an incident that has already occurred, estimation of its damage is conditional on the observed event; it does not require estimating how often that event will occur next year.

ROSI and security investment models evaluate a counterfactual reduction in future loss relative to protection expenditure [6], [7], [19]. Cyber value-at-risk specifies a loss

quantile for a stated horizon and confidence level [20], [21]. These quantities inform investment or capital decisions, but neither a quantile nor expected annual loss estimates the incurred cost of one realised event. Sparse data limit calibration, yet distributions may also be elicited from expert judgement. Their provenance and uncertainty must then be disclosed. Frequency data are a constraint on prospective risk estimation, not a prerequisite for recording documented recovery costs.

3.3 Empirical cost studies and benchmarks

Empirical studies supply evidence about the dispersion of incident costs and the influence of organisational and event characteristics [4], [5], [22], [37]. Industry reports serve different purposes: a breach cost survey estimates costs in a sampled population, whereas the Verizon Data Breach Investigations Report primarily characterises incidents and attack patterns [1], [38]. Their denominators, inclusion rules and reporting mechanisms must be examined before results are combined.

A national assessment should use external benchmarks to test plausibility and to inform explicitly labelled estimates where local evidence is limited. Direct substitution is problematic when wage levels, liability rules, service portfolios and reporting incentives differ. Currency conversion alone cannot correct these differences. Neither a sample mean nor a loss-per-record rate supplies a causal estimate for an individual event. A benchmark-based component should retain the source population, adjustment assumptions and a sensitivity range, and should be replaced by incident-specific evidence as it becomes available.

3.4 Regulatory significance and economic impact

NIS2 and DORA link incident reporting to the significance of consequences within their respective scopes [11], [12]. Their subordinate acts contain more economic detail than a simple ordinal severity scale. Article 7 of Commission Delegated Regulation (EU) 2024/1772 specifies direct and indirect cost and loss categories for determining economic impact under DORA [24]. Commission Implementing Regulation (EU) 2024/2690 specifies significant-incident criteria for listed digital service providers and discusses relevant financial loss categories [25]. It would therefore be inaccurate to state that regulatory instruments provide no guidance on monetary loss.

The remaining distinction concerns purpose and boundary. A reporting threshold answers whether a statutory notification condition is met. A financial estimate supports accounting or resource allocation under a specified valuation perspective. A regulatory total may prescribe a different treatment of recoveries or forgone revenues from an economic resource-loss estimate. These totals require a documented reconciliation rather than forced numerical equality. The turnover-based maximum fines in GDPR are legal sanctions, not estimates of the harm suffered by individuals [23]. The ENISA personal data breach methodology provides a severity assessment approach [39], while the US cyber incident severity schema supports coordination and escalation [40].

There are also concrete alternatives to an isolated national total. The European Supervisory Authorities' Joint Guidelines JC 2024 34 provide a reference-year approach to aggregated annual incident costs and losses and separate financial recoveries [41]. These instruments demonstrate existing monetary reporting practice. Adapting their boundary to public services remains a design choice rather than evidence that no alternative approach exists.

Version control is also material. NIST SP 800-61 Revision 3, issued in April 2025, supersedes Revision 2 and aligns incident response guidance with CSF 2.0 [42]. Older lifecycle descriptions may remain relevant when explicitly discussed historically. Their use should not be represented as application of the current edition. More generally, the assessment record should identify the rule and edition that governed classification at the time of reporting.

3.5 Business continuity and operational technology safety

Business impact analysis links disrupted activities to organisational objectives, consequences over time and recovery priorities [26], [27]. It can support valuation of interruption, but does not guarantee that a validated hourly loss rate exists. Such a rate may depend on transaction timing, spare capacity, manual alternatives and whether delayed work is subsequently completed. Multiplying annual operating expenditure by a fraction of unavailable hours measures an allocation of service resources; without additional assumptions it does not measure incremental damage.

For operational technology, IEC 62443-3-2 and NIST SP 800-82 support assessment of cyber-related consequences in systems that interact with physical processes [28], [29]. Functional safety and integrated safety–security approaches add analysis of hazardous scenarios [43], [44]. Actual injuries, hazardous releases and exposure duration should be recorded separately from the probability of a hypothetical accident. Medical expenditure and repair costs may enter the relevant monetary account. Injury severity and risk to life should remain visible as separate outcomes even when a jurisdiction permits economic valuation of health effects. A low estimated financial loss cannot cancel a safety-critical escalation condition.

3.6 Interdependencies and cascading consequences

Input–output models and infrastructure interdependency research describe how disruption can propagate beyond its initial location [30], [31]. ENISA and OECD materials provide complementary guidance on critical functions and infrastructure resilience [32], [45]. Supply chain guidance addresses dependency on technology providers and services [46], [47]. Applying these approaches to a realised cyber incident requires an event boundary, evidence of affected links and a distinction between observed downstream losses and modelled scenarios.

For an explicitly defined linear loss-propagation approximation, let d be a vector of initial losses and A a non-negative matrix whose element a_{ij} maps a unit of loss in sector j into an additional loss in sector i over the same horizon. A finite non-negative infinite-series solution for arbitrary non-negative d is obtained when the spectral radius ρ(A), the largest eigenvalue magnitude, is below one. Equation (2) states the resulting model.

$$x = d + Ax = (I - A)^{-1}d, \quad A \geq 0, \quad \rho(A) < 1 \quad (2)$$

Here x contains total sector losses and I is the identity matrix. Strictly sub-unit row or column sums are sufficient conditions for convergence, but are not necessary. Sums no greater than one are not sufficient when ρ(A) equals one. Nor is there a general economic requirement that downstream losses be smaller than the initial loss. For example, a hypothetical two-sector matrix with both diagonal elements equal to zero and both off-diagonal elements equal to 0.6 has ρ(A) = 0.6. With d = (1, 0) million KZT, Equation (2) gives x = (1.5625, 0.9375) million KZT. Total loss is 2.5 million KZT, including 1.5 million of propagated loss. This is a mathematical counterexample to the proposed cap, not an empirical estimate.

Equation (2) is an illustrative abstraction, not a validated cybereconomic model or a direct reuse of national input–output coefficients. Its coefficients require causal interpretation and calibration; near-unit spectral radii can amplify parameter uncertainty. If x already includes d, the two vectors must not be added again. Empirical implementation should also use consistent value-added or resource-cost concepts to avoid counting intermediate transactions repeatedly. Where these conditions cannot be met, report observed downstream costs and separate bounded scenarios rather than an unsupported multiplier.

Where observations replace model outputs, maintain a sector-by-item reconciliation: each downstream item has an affected owner, time interval and causal link to the initiating event. A modelled sector total and the observed expenses it already represents cannot both enter the same consolidated total. A finite propagation-horizon model can instead sum d, Ad, …, A raised to K times d; the infinite-series convergence condition then does not apply, although causal validity and avoidance of repeated valuation remain necessary. No coefficient restriction can substitute for this accounting reconciliation.

3.7 Composite indicators and monetary valuation

Composite indicators combine selected variables through normalisation, weighting and aggregation [13], [33]. Their interpretation depends on the measurement scales and on whether compensation between dimensions is acceptable. ISO/IEC 27004 supports a disciplined approach to information security measurement and evaluation [48]. A severity score may be useful for escalation, but the score alone does not establish an exchange rate between public trust, safety and money.

Multiplying a dimensionless quantity by a financial base is not inherently invalid. For example, multiplying a calibrated physical damage ratio by an appropriate replacement value can be meaningful. The problematic step is treating an uncalibrated multi-criteria severity index as if it were an empirically established loss fraction. Such conversion requires an identified monetary target, supporting data, calibration, validation and uncertainty analysis. In their absence, index values and monetary totals should remain separate.

A numerical distinction is useful. A calibrated 0.20 physical damage fraction applied to a 10 million KZT replacement value produces a 2 million KZT estimate for the specified asset. An ordinal severity score of 0.20 applied to the same base has no equivalent meaning without calibration. Neither calculation permits a second entry for the same damaged asset. ENISA's personal data breach method defines severity levels in Section 3.2, not monetary conversion factors [39].

The proposed architecture therefore retains a monetary account and a consequence profile as linked outputs with a common incident identifier (Figure 2). Critical safety or essential-service conditions should be capable of triggering escalation independently. If an aggregate score is used for other purposes, its component values, weights and sensitivity should remain available; averaging must not conceal a non-compensable consequence.

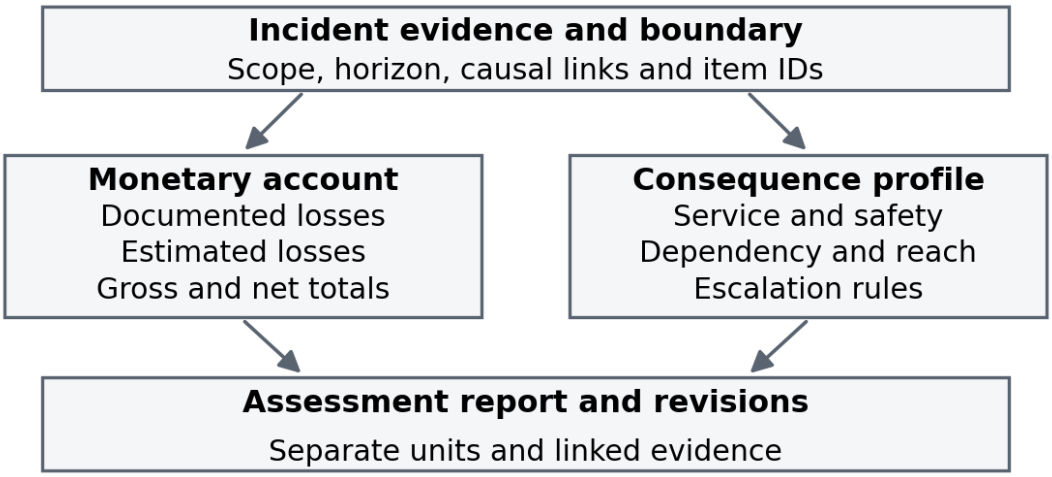


Figure 2. Proposed assessment architecture. Financial loss and consequence severity share evidence and an incident boundary but remain distinct outputs.

3.8 Comparison across consequence dimensions

Table 2 distinguishes feasible valuation routes from consequences that require non-monetary reporting. The comparison does not establish that a dimension is absent from every member of a family. It identifies the evidence and boundary conditions needed when representative methods are used together. In particular, technological replacement costs can

be measured without claiming to monetise technological sovereignty itself, and public service interruption can be measured in service units without treating the delivery budget as lost revenue.

Table 2. Consequence dimensions and defensible measurement routes

Dimension	Measurement route and anchors	Boundary condition
Financial expenditure	Incident response, legal and recovery entries [10], [42], [24]	Identify incremental costs; distinguish payments, provisions and recoveries.
Material damage	Repair or replacement cost for damaged assets [28], [29]	Avoid adding full replacement cost and full asset write-off for the same loss.
Operational disruption	Unrecovered output or contribution margin; additional work [26], [27]	Account for catch-up, saved costs and workarounds.
Workforce productivity	Attributable work hours and a declared valuation basis [10], [18]	Separate opportunity cost from additional payroll; exclude overlaps.
Reputation	Incremental churn, contract loss or mitigation expenditure [4], [5], [10]	Establish attribution; avoid adding market-value loss to overlapping cash flows.
Public service continuity	Failed or delayed requests, affected users, backlog and alternatives [26], [27]	Allocated budget is a resource proxy; citizen loss needs separate evidence.
Technological dependency	Incident-induced migration, substitution and integration costs [46], [47]	Exclude unrelated modernisation; report residual dependency separately.
Cross-sector propagation	Observed downstream costs or a specified dependency model [30], [31]	Use a common boundary; reconcile transfers and duplicated outputs.
Human safety	Observed harm, exposure and scenario consequences [29], [43], [44]	Retain non-monetary severity even when some costs are monetised.

Three integration problems follow from this comparison. First, one underlying consequence can appear in several categories: an interrupted activity may already incorporate idle staff, contract penalties and recovery expenditure. Second, estimates with different purposes may share currency units while representing different objects. Third, costs borne by one entity can be transfers received by another. A national aggregation cannot be obtained by adding all organisational bills without specifying the economic perspective.

The evidence supports specific mechanisms, not a prevalence estimate. The OECD/JRC handbook explicitly discusses collinear indicators that can give excess weight

to one underlying dimension and recommends uncertainty and sensitivity analysis [13], Sections 1.6–1.7. This is an identified aggregation risk and an existing safeguard, rather than evidence that the handbook commits an error. NIST's qualitative and quantitative scales [9], Section 2.3.2, ENISA's severity levels [39], Section 3.2 and the cost categories in Regulation 2024/1772 [24], Article 7 illustrate different measurement objects. The invalid step would be combining those outputs without reconciling their meanings. The cases below identify concrete budget-proxy, sector-uplift and severity-feedback problems in the earlier national draft [14]. These examples support source-specific criticism, not a claim that such defects recur at a measured rate across the corpus.

Table 3. Consistency checks for a combined assessment

Issue	Diagnostic question	Required treatment
Risk versus realised damage	Is the output an event loss, an annual expectation or a quantile?	Label the object and horizon; compare only like quantities.
Duplicate loss	Can two entries refer to the same resource or output?	Use a unique item identifier and an exclusive accounting allocation.
Score converted to money	Has the index been calibrated to a defined monetary target?	Keep uncalibrated severity scores outside the loss total.
Unstated denominator	Does the base represent annual flow, asset value or lifecycle cost?	State base, year and scope; avoid cross-base rankings.
Transfers and recoveries	Does the amount change resource loss or only who bears it?	Report organisational gross and net loss; reconcile national totals.
Uncertainty hidden in a point estimate	Which inputs are observed, estimated or conventional?	Disclose evidence status, ranges and sensitivity.

3.9 Requirements for a national assessment procedure

Kazakhstan provides a relevant application context because public digital services and critical infrastructure combine organisational, public and cross-sector consequences. The Digital Code of the Republic of Kazakhstan, adopted on 9 January 2026, and legislation on state services provide the legal context [49], [50]. The architecture below is a methodological proposal; it is not represented as an adopted legal requirement or an independently validated national instrument.

The first step is to define an incident record: affected entities and services, assessment date, start and end of the loss horizon, currency and price basis, and perspective. At least three perspectives should be distinguishable: the operator's organisational loss, the fiscal effect on public bodies and the wider economic or social consequence. These views can overlap and should not be added together as independent layers. A penalty paid by an

operator to the state is a cost to the operator but a transfer within a consolidated national account; the associated administrative resources are a separate matter.

The monetary account should start from distinct loss items rather than from broad labels such as direct, indirect and strategic. Let c_j be the monetary value of item j and a_j indicate whether that item is included within the declared boundary. For every included item, a category-allocation variable z_{jk} assigns it to exactly one monetary category. Equation (3) defines the aggregation and exclusivity condition.

$$U_{gross} = \sum_{j=1}^{n} a_j c_j = U_D + U_E, \quad \sum_{k=1}^{K} z_{jk} = 1 \quad (3)$$

U_D is the subtotal supported by incident-specific documentary evidence and U_E is the subtotal that depends on explicit estimation assumptions. Both require attribution to the incident; a paid invoice alone does not establish that the entire amount is attributable. The allocation rule prevents duplicate entries when identifiers are maintained, but it cannot by itself detect two differently named records of the same loss. Reconciliation of invoices, staff time and output measures remains necessary. An item can have several descriptive consequence tags while retaining one monetary allocation.

Gross loss is reported before insurance or third-party recoveries. For an organisational account, received recoveries R_{rec} can be disclosed separately and reconciled to net loss as in Equation (4). Expected, disputed or conditional recoveries should retain their uncertainty and not be presented as already received.

$$U_{net} = U_{gross} - R_{rec} \quad (4)$$

An interruption estimate must follow the observed service trajectory. Let q(t) be the fraction of normal useful service delivered at time t, bounded between zero and one, and let b(t) be the marginal value of output permanently lost after allowing for subsequent catch-up and avoided variable costs, expressed in KZT per hour. A simplified operational loss model is given in Equation (5).

$$U_{op} = \int_0^T [1 - q(t)]\, b(t)\, dt + C_{alt} + C_{backlog} + C_{pen} \quad (5)$$

C_{alt} denotes additional workaround costs, $C_{backlog}$ additional costs of clearing delayed work and C_{pen} attributable contractual penalties under the chosen perspective. Each term must exclude amounts already captured in b(t) or another account. This model is suitable only when b(t) has a defensible basis. For a public service without a meaningful monetary output value, report service disruption directly and quantify evidenced additional resources. An annual service budget B divided by planned service hours H, multiplied by

equivalent interruption hours T_{eq}, may be disclosed as the resource proxy V_{proxy} in Equation (6), but should not automatically enter the loss ledger.

$$V_{proxy} = \frac{B}{H} T_{eq}, \quad T_{eq} = \int_0^T [1 - q(t)]\, dt \quad (6)$$

The consequence profile should retain service scope, geographic reach, population affected, safety outcomes and residual technological dependency. Escalation rules may use these variables even when a monetary estimate is incomplete. Restoration, migration, integration and retraining expenses can be included when attributable to the incident; a general coefficient for technological sovereignty is not a substitute for identifying those expenses. This separation allows financially small but safety-critical incidents to remain visible.

Table 4. Minimum information needed for national implementation

Record element	Required information	Purpose
Incident boundary	Identifiers, entities, affected services, causal links and time horizon	Link primary and downstream records without duplicate incidents.
Monetary item	Item ID, amount, units, owner, period and category	Reconstruct totals and enforce exclusive allocation.
Evidence status	Documented, estimated or excluded proxy; source and rationale	Expose the support for each included value.
Valuation perspective	Organisational, fiscal or wider economic account	Identify transfers and prevent incompatible aggregation.
Uncertainty	Plausible ranges or distributions; dependence assumptions	Identify which missing evidence changes the estimate most.
Consequence profile	Service disruption, safety, dependency and territorial reach	Support escalation without forcing every consequence into money.
Revision history	Assessment date, changed inputs and reasons	Separate preliminary estimates from settled outcomes.

3.9.1 Loss reconciliation and identifiable overlap

A category list alone does not prevent double counting. Consider n distinct underlying loss items with non-negative values c, all within the same assessment boundary. Let H be a binary observation matrix: $H_{ij} = 1$ when reported subtotal i includes item j. Assume every item is represented by at least one row. The reported subtotals are $r = Hc$ and the multiplicity vector $m = H^T 1$ counts how often each item appears. Equation (7) gives the exact inflation produced by summing the reports without reconciliation.

$$r = Hc, \quad m = H^{T}1, \quad \Delta_{dup} = 1^{T}r - 1^{T}c = (m - 1)^{T}c \quad (7)$$

Proposition 1. Under these assumptions, the raw sum equals the unique-item total if and only if no positive-valued item has multiplicity greater than one. Proof: Equation (7) is an identity, and every term $(m_j - 1)c_j$ is non-negative. The discrepancy vanishes exactly when all duplicated items have zero value. For example, items (1, 2, 3) and reports containing items {1, 2} and {2, 3} give subtotals 3 and 5; their sum 8 overstates the unique total 6 by the repeated item 2. The condition is independent of category names.

The accounting implication is operational: a component record must identify what an aggregate hourly rate or invoice already contains. When H is unknown, the discrepancy cannot be estimated merely from the number of categories. In the mine case, adding staff time to a downtime rate is valid only if that rate excludes the same displaced work. A sector multiplier or a cap on total loss cannot establish this exclusion. Equation (3) allocates reconciled items to categories; Equation (7) addresses the prior task of detecting repeated observations.

3.9.2 Parameter uncertainty and sensitivity

Let θ collect valuation inputs and let Θ specify their admissible joint values. The fixed expense subtotal D is distinguished from modelled components $f_j(\theta)$. Equation (8) defines an uncertainty envelope for the monetary total. The set Θ may encode dependence or exclusion constraints; it is not automatically a Cartesian product of independent marginal intervals.

$$U(\theta) = D + \sum_{j=1}^{J} f_j(\theta), \quad U^{-} = \min_{\theta \in \Theta} U(\theta), \quad U^{+} = \max_{\theta \in \Theta} U(\theta) \quad (8)$$

If each uncertain input has a non-negative partial derivative and the feasible set is a rectangle, the lower and upper corners attain the exact bounds. No Monte Carlo approximation is needed in that case. If combinations are infeasible, corner calculations provide only outer bounds until the joint constraint is imposed. These are uncertainty sets, not confidence intervals: a probability statement requires a defensible probability model in addition to admissible ranges.

For a positive total and a positive parameter p, the local elasticity in Equation (9) measures the proportional change in total loss caused by a small proportional change in p, with other inputs fixed. When a component is proportional to p, its elasticity equals that component's share of the total. This provides a direct test of whether the monetary conclusion is governed by a documented expense or by a convention.

$$E_p = \frac{p}{U}\frac{\partial U}{\partial p} \quad (9)$$

For a component containing several multiplied inputs, each input can have the same component-share elasticity; these input elasticities must not be summed as if they were disjoint cost shares. Zero-valued inputs require absolute derivatives instead. The sensitivity of an estimate to a parameter is also distinct from uncertainty in that parameter: a large derivative is consequential only when the parameter is poorly constrained over a material range.

3.9.3 Robustness of financial and overall classification

Let $B > 0$ be the declared comparison base and let $\tau_1 < \tau_2 < \tau_3$ be three percentage thresholds. Equation (10) maps a financial ratio to ordinal levels 1–4. The indicator is one when its condition is true and zero otherwise. The inclusive convention at a boundary is used here for the specified public-service scenarios; a different reporting rule must state its own convention.

$$r = \frac{100U}{B}, \quad k_F(r) = 1 + \sum_{h=1}^{3} 1\{r \geq \tau_h\} \quad (10)$$

Suppose U lies in $[U^-, U^+]$ and B in $[B^-, B^+]$, with positive denominators. Equation (11) bounds the possible ratio. Proposition 2. The financial category is robust whenever both endpoint ratios fall in the same band. This condition is also necessary when both opposite corners of the joint feasible set can occur. Proof: U/B is increasing in U and decreasing in B; the band function is non-decreasing. Therefore every feasible ratio lies between the endpoints and has the same category when their categories agree.

$$r^- = \frac{100U^-}{B^+}, \quad r^+ = \frac{100U^+}{B^-}, \quad k_F(r^-) = k_F(r^+) \quad (11)$$

When the endpoints give different categories but the joint corners are not feasible, instability has not yet been proved; the feasible set must be evaluated more tightly. Likewise, if B also enters the loss formula, U(B) must be recomputed jointly rather than varying the numerator and denominator independently. The maximum-rule overall level is $\max\{k_F, k_O, k_H, \ldots\}$, where the additional terms denote separately mapped operational and safety levels. Consequently a robust critical operational condition can make overall severity robust even when financial severity changes across several bands. This is a classification property, not evidence that monetary valuation is accurate.

3.9.4 Propagation strength and finite assessment horizons

To separate model behaviour from a supposed downstream-loss cap, consider A with zero diagonal and equal off-diagonal entries a, where $0 \leq a < 1$, and initial loss $d = (d_0, 0)$, $d_0 > 0$. Let K count the number of propagation steps, with step zero representing the initial loss. Since $1^T A = a1^T$, the aggregate multiplier is the geometric sum in Equation (12).

$$M_K(a) = \sum_{k=0}^{K} a^k = \frac{1-a^{K+1}}{1-a}, \quad M_\infty(a) = \frac{1}{1-a} \quad (12)$$

It follows that the infinite-series propagated loss divided by initial loss equals a/(1 − a), which exceeds one precisely when a > 0.5. A rule capping downstream loss at d_0 would then limit total loss to $2d_0$ and understate the model total by the fraction 2a − 1. The result is exact for this constructed model. It does not estimate a national propagation coefficient: repeated propagation must represent distinct marginal consequences, not repeated valuation of the same consequence. Finite K can represent a declared truncation of dependency paths; interpreting steps as elapsed time requires an additional dynamic model.

3.10 Reconstructing the draft-method demonstrations

Appendix B of the earlier national draft [14] explicitly describes its two cases as demonstrations whose parameters require factual verification before real-incident use. Its portal total is 187,230,752 KZT and its mine total 80,040,000 KZT; these reproduce the rounded 187.2 and 80.0 million in the earlier draft. The calculations below recover those figures from the actual inputs, without treating scenario costs as observed invoices. S2 identifies every input and coefficient source. Monetary amounts are KZT; tables use million KZT. No recovery or inflation adjustment is specified by the source, so no net-loss or constant-price cross-incident comparison is inferred.

3.10.1 Government portal

The source assumes 18 hours of complete unavailability, a 600 million KZT annual service budget and 8,760 annual operating hours. Recovery comprises restoration 8 million, additional personnel 3 million, emergency arrangements 2 million and audit 1 million KZT. Staff productivity uses 35 employees, 12 hours, 4,500 KZT per hour, efficiency reduction 0.60 and sector multiplier 1.50. The reputation convention uses 75,000 affected users, a draft-assumed monthly calculation index of 4,325 KZT, sector coefficient 0.70 and severity coefficient 0.75. The latter coefficients come from the draft's own scales and Appendix C, not an external empirical valuation study.

The portal case can be written as Equation (13), in million KZT. Here s is the productivity sector multiplier, z is a binary inclusion flag for the allocated-budget proxy, and λ scales the reputation convention. The original case is (s, z, λ) = (1.50, 1, 1); the reconciled organisational scenario is (1, 0, 0). This parametrisation separates a change in the valuation rule from a change in observed service disruption.

$$U_P(s, z, \lambda) = 14 + 1.134s + 1.232877z + 170.296875\lambda \quad (13)$$

Table 5. Reconstruction and evidential status of the original portal calculation

Component	Calculation in source units	Million KZT	Interpretation
Recovery	8,000,000 + 3,000,000 + 2,000,000 + 1,000,000	14.000000	Scenario expense subtotal; supporting invoices not supplied
Productivity	35 × 12 × 4,500 × 0.60 × 1.50	1.701000	Model estimate including an uncalibrated sector uplift
Service continuity proxy	18 × (600,000,000 / 8,760) × 1.00	1.232877	Allocated budget; rounded to whole KZT
Reputation convention	75,000 × 4,325 × 0.70 × 0.75	170.296875	Normative scenario amount, not measured reputation loss
Original total	14,000,000 + 1,701,000 + 1,232,877 + 170,296,875	187.230752	Reproduces the draft; does not establish actual loss

The hourly budget proxy is calculated from the unrounded quotient before rounding the final component to whole KZT; multiplying the displayed rounded hourly value of 68,493 would give a slightly different answer. The reputation amount is 90.96% of the total. The remaining 16.933877 million is not a documented-cost baseline: it still contains a budget allocation and a sector-weighted productivity estimate. Removing the reputation convention, excluding the budget proxy and setting the productivity sector multiplier to one yields an organisational scenario account of 14 + 1.134 = 15.134 million KZT. Even its 14 million expense subtotal requires underlying records before it can be called documented incident loss.

This separation is consistent with revision 2.2 [15], Appendix B.1, which uses 15.134 million KZT for the organisational account and excludes the budget proxy. It additionally estimates citizen-time costs as 75,000 × 0.30 × 1 hour × 506 KZT/hour = 11.385 million KZT. The rate 506 is the draft's rounded 85,000/168 hourly proxy; the repeat-contact share and additional time are scenario inputs, not observations available to this review. The wider total is 26.519 million KZT. Varying the repeat-contact share over 0.20–0.40 and time over 0.75–1.25 hours gives citizen-time costs of 5.6925–18.975 million and a wider total of 20.8265–34.109 million. These are deterministic assumption bounds. Third-party time costs should remain separate from the organisation's own expenditure and estimates.

With r denoting repeat-contact share and t denoting additional hours per affected citizen, the wider account has the bilinear form in Equation (14). Its positive derivatives justify using the joint endpoints to obtain the reported bounds. The coefficient 37.95 is

75,000 × 506 / 1,000,000 and has units of million KZT per hour; the share r is dimensionless and expressed as a fraction.

$$U_{soc}(r,t) = 37.95rt, \quad U_{wide}(r,t) = 15.134 + 37.95rt \quad (14)$$

3.10.2 Mine ventilation control

The earlier draft assumes a 12-hour complete stoppage and a downtime rate of 3 million KZT per hour. The rate's cost composition is not documented in the earlier case. Recovery, response, legal support and replacement are separate scenario amounts; safe shutdown/restart and process losses are added separately. Productivity uses 80 employees, 8 hours, 5,500 KZT per hour, efficiency reduction 0.80 and a sector multiplier of 2.50. Appendix C gives a 2.00–3.00 range for critical objects, but neither that range nor the selected 2.50 has an empirical calibration supplied with the case [14].

Equation (15) gives the mine total in million KZT when c is the downtime rate in million KZT per hour, κ the efficiency-reduction factor and s the productivity sector multiplier. The 37 million subtotal comprises 26 million direct expenditure and 11 million process-related costs. The original parameters are c = 3, κ = 0.80 and s = 2.50. The model is conditional on exclusion of these separate losses from c.

$$U_M(c,\kappa,s) = 37 + 12c + 3.52\kappa s \quad (15)$$

Table 6. Reconstruction of the original mine calculation

Component	Calculation in source units	Million KZT	Required qualification
Downtime	12 × 3,000,000 × 1.00	36.000	Hourly rate must exclude separately counted losses
Direct expenditure	12,000,000 restoration + 8,000,000 response + 1,000,000 legal + 5,000,000 replacement	26.000	Scenario amounts; actual financial records unavailable
Process losses	7,000,000 safe shutdown/restart + 4,000,000 process losses	11.000	Must not be embedded in hourly downtime cost
Productivity	80 × 8 × 5,500 × 0.80 × 2.50	7.040	Sector uplift changes cost without evidence of additional lost work
Original total	36,000,000 + 26,000,000 + 11,000,000 + 7,040,000	80.040	Conditional on disjoint components
Safety index, separate	0.50 × 0.75 × 0.75 × 0.50 = 0.140625	Not money	Rounded HSI = 0.141; no addition to total

The four safety factors represent the draft's assumed hazard probability, harm severity, exposure and barrier degradation. Their product is a constructed index, not an empirically established probability of injury. The earlier draft places 0.141 in its moderate safety band; revision 2.2 places the same number in a high band. This change illustrates why an index value without its edition-specific thresholds is insufficient. The source's 2,500 million KZT comparison base is conditional; revision 2.2 identifies it as asset value. Its 3.2016% original loss ratio must not be directly compared with the portal's ratio to an annual budget.

Removing only the 2.50 productivity uplift while retaining the original efficiency factor 0.80 reduces the mine total to 75.816 million KZT. Revision 2.2 instead sets the efficiency factor to 1.00 for complete staff stoppage and uses no sector uplift, giving productivity of 3.520 million and a total of 76.520 million. It explicitly excludes labour, restoration and process losses from the 3 million hourly rate and separates an expense subtotal of 37 million from estimates of 39.520 million. Those exclusions make the ledger internally reconstructible; the demonstrations still do not supply primary documents proving them [15], Appendix B.2.

An additional stress test, introduced here rather than attributed to the draft, varies the downtime rate by ±20% (2.4–3.6 million KZT/hour) and efficiency reduction from 0.60 to 1.00, with no sector uplift and the 37 million expense subtotal fixed. The total ranges from 67.912 to 83.720 million KZT. These are conditional endpoints, not confidence limits. If the hourly rate already includes staff or process losses, the relevant separate entries must be removed, not merely assigned narrower ranges. A verified actual case also needs the assessment horizon, recoverable output, avoided costs and later adjustments.

3.11 Reputation sensitivity and severity classification

For the earlier portal case, let λ multiply only the reputation convention. The exact reconstructed sensitivity function in million KZT is $U(\lambda) = 16.933877 + 170.296875\lambda$. The comparison ratio is $100U(\lambda)/600$. Table 7 uses the earlier draft's financial bands: below 1% low, 1–5% moderate, 5–15% high and above 15% critical [14], Section 7.3. Its boundary rule selects the higher band for public services; none of the tabulated scenarios lies on a boundary. These are the source's proposed bands, not validated or adopted national thresholds.

For a threshold τ and fixed base B, Equation (16) gives the multiplier at which $L + \lambda R$ reaches the threshold, where L = 16.933877 and R = 170.296875 million KZT. A threshold is reachable within the specified multiplier interval only if λ^* lies in that interval. This supplies a quantitative robustness margin rather than simply declaring an estimate sensitive.

$$\lambda^* = \frac{\tau B/100 - L}{R}, \quad R > 0 \quad (16)$$

Table 7. Portal sensitivity under the earlier draft's financial classification

Reputation multiplier λ	Total, million KZT	Annual-budget ratio	Financial band
0.00 — excluded	16.933877	2.8223%	Moderate
0.25	59.508096	9.9180%	High
0.50	102.082315	17.0137%	Critical
1.00 — original	187.230752	31.2051%	Critical
1.50	272.379190	45.3965%	Critical

The financial band changes from critical to moderate when reputation is excluded; the 15% threshold is reached at $\lambda \approx 0.42905$. This does not automatically lower overall incident severity. Sections 7.2–7.3 of the earlier draft classify a stopped key service as critical and take the highest applicable consequence level. Given the scenario's complete portal outage, that operational rule keeps the overall classification critical at every multiplier. Even if an assessor did not identify the portal as a stopped key service, the supplied GSCI of 0.68 would independently imply a high level. Thus the monetary total is strongly convention-dependent while operational escalation can remain unchanged.

The original reputation coefficient also creates a feedback risk: it is set to 0.75 for a high incident, yet the resulting financial total is critical. Replacing it by the draft's critical coefficient of 1.00 increases reputation to 227.0625 million and the total to 243.996377 million KZT. A procedure must either freeze an independently determined non-monetary severity input before valuation or avoid using the resulting money-based category to set its own monetary coefficient. This is an identifiable problem in the earlier draft, not evidence that all monetary or index-based approaches are defective.

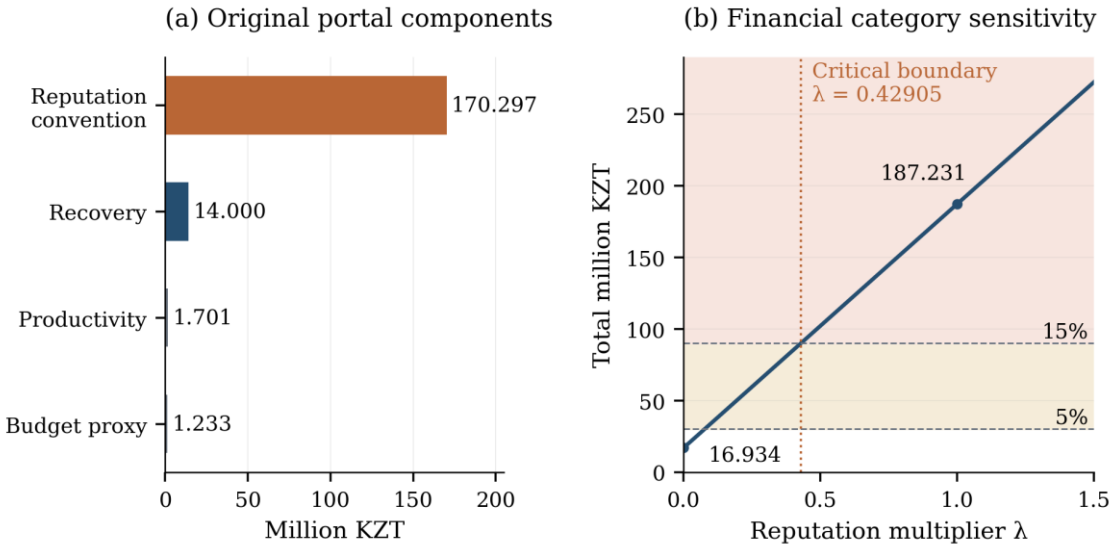


Figure 3. Portal component dominance and reputation sensitivity. Panel (a) reconstructs the earlier draft's components; panel (b) varies only the reputation multiplier at B = 600 million KZT. Dashed horizontal lines mark the draft's 5% and 15% financial thresholds. The operational classification is held separate. Source inputs: [14]; derived curves: accompanying scripts.

Revision 2.2 changes the financial threshold basis to the documented organisational subtotal and uses bands below 0.1%, 0.1–1%, 1–5% and at least 5% [15], Sections 6.1–6.3. Within its demonstration, the portal expense subtotal 14/600 = 2.3333% and the mine subtotal 37/2,500 = 1.48% both give high financial severity. Overall severity remains critical because the assumed service stoppages exceed their recovery time objectives: 18 versus 12 hours and 12 versus 4 hours. These results describe the draft rules conditional on their assumed evidence. Actual documented status cannot be inferred from a label in a demonstration. Threshold changes and accounting changes must be reported separately; lower revised totals alone do not demonstrate more accurate valuation.

3.12 Provenance and limits of the thirteen-case table

Appendix G (Russian designation Г) of the earlier draft [14] contains thirteen anonymous sector/scenario entries. Revision 2.2 [15], Appendix G, pp. 109–110, explicitly states that their amounts and valuation logic were reported by organisations and were not calculated using the methodology's formulas. S2 reproduces all thirteen entries and their stated valuation descriptions. The arithmetic is 167.7 million KZT in total, a mean of 12.9 million and a range of 1–90 million. This establishes the immediate documentary source and reported status, not independent authenticity: incident dates, named organisations, primary records, price bases and a collection protocol are not provided.

The table supports a limited check of whether the proposed categories can describe the reported consequences. It cannot establish accuracy, calibration, representativeness or

successful end-to-end field application. Four entries mention ransom; two contain only ransom amounts of 3.7 and 2.0 million, while the transport and media entries do not separate ransom from other losses. Consequently neither payment verification nor consistent recalculation is possible from the table alone. The present study reconstructs demonstrations and examines a secondary table of reported losses; it performs no independently verified application to primary incident records.

3.13 Quantitative results and interpretation

3.13.1 Sensitivity and classification stability

At the original portal point, the reputation-multiplier elasticity is 0.90956. Thus a 1% change in that multiplier changes the total by approximately 0.91% locally; the relation is exactly linear for a fixed baseline. In contrast, equal proportional changes in the recovery, productivity and allocated-budget components have total elasticities 0.07477, 0.00909 and 0.00658. The monetary estimate is therefore dominated by the reputation convention. The critical financial boundary occurs at $\lambda = 0.42905$, so reducing that convention by about 57.1% from its original level reaches the 15% threshold. Eliminating it crosses both the critical and high boundaries.

Figure 4 shows the distinct sources of uncertainty after removing the unsupported sector uplifts. The wider portal total increases bilinearly with repeat-contact share and citizen time. At $r = 0.30$ and $t = 1$ hour, an increase of 0.10 in r adds 3.795 million KZT; an increase of 0.25 hour in t adds 2.84625 million. These are local one-input changes, not additive estimates for a simultaneous change because the bilinear interaction must also be included. In the mine account, a 0.1 million KZT/hour change in downtime rate changes loss by 1.2 million, whereas a 0.10 change in efficiency reduction changes loss by 0.352 million. Under the specified ranges, the downtime-rate assumption accounts for 14.4 million of the 15.808 million total interval width.

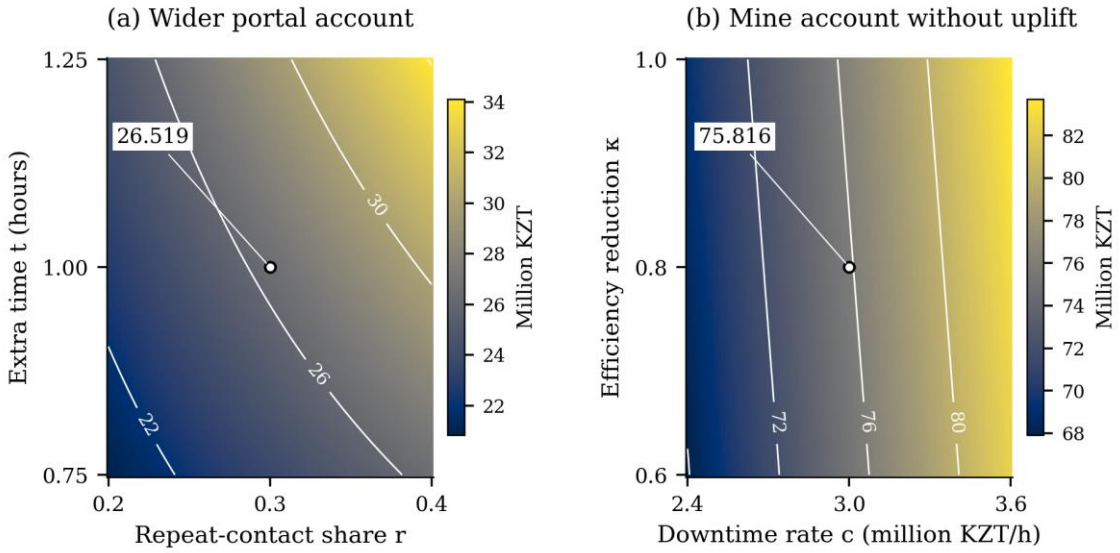


Figure 4. Deterministic monetary response surfaces. Panel (a) uses U = 15.134 + 37.95rt and marks the revision-2.2 portal assumptions; panel (b) uses U = 37 + 12c + 3.52κ and marks the earlier mine efficiency value with the sector uplift removed. Colour and contour labels show million KZT. The grids visualise analytical bounds and are not statistical confidence surfaces.

Table 8. Valuation sensitivity and category robustness under explicit rules

Scenario and varied inputs	Money interval million KZT	Financial category	Overall consequence level
Earlier portal, λ from 0 to 1.5; B = 600	16.933877–272.379190	Moderate–critical; not robust	Critical if the key-service stoppage condition is retained
Revision 2.2 portal, r = 0.20–0.40 and t = 0.75–1.25	20.8265–34.109 for the wider account	High from fixed expense subtotal 14; robust to these inputs	Critical from 18-hour stoppage exceeding 12-hour RTO
Revision 2.2 mine accounting, c = 2.4–3.6 and κ = 0.60–1.00; s = 1	67.912–83.720 for organisational account	High from fixed expense subtotal 37; robust to these inputs	Critical from 12-hour stoppage exceeding 4-hour RTO

The stability in the latter two rows follows from the revision's choice of financial numerator, not from precise knowledge of all economic losses. Estimated productivity, operations and citizen time are outside that numerator. If the supposed expense documentation or the comparison base changes, the financial result must be reevaluated. Table 8 therefore demonstrates robustness to the specified parameters only. It also shows

why replacing one draft's total and thresholds with another draft's rules in the same calculation would produce an uninterpretable comparison.

3.13.2 Cascading amplification and the cost of a cap

Figure 5 compares exact finite-step sums with the convergent limit. At a = 0.60, three propagation steps already give a multiplier of 2.176; the infinite limit is 2.5. At a = 0.80, the limit is 5. A downstream-loss cap equal to initial loss would understate these infinite-horizon totals by 20% and 60%, respectively. The comparison refutes a universal cap within a model whose assumptions are fully specified; it does not show that every real incident exhibits such amplification. As a approaches one, the infinite estimate becomes highly sensitive and finite-horizon interpretation becomes increasingly material.

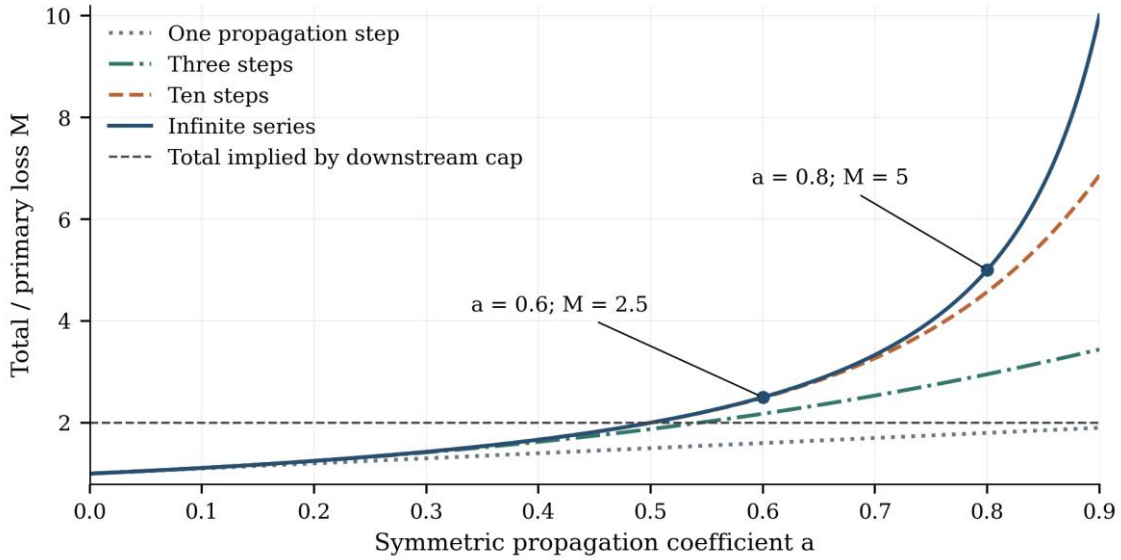


Figure 5. Total-to-primary-loss multiplier in the symmetric two-sector model. K denotes propagation steps, not elapsed hours. The horizontal line at two is the maximum total implied by capping downstream loss at primary loss. All curves are calculated from Equation (12); coefficients are controlled mathematical inputs.

3.13.3 Concentration in the secondary reported-loss table

The thirteen reported amounts have a median of 4.0 million KZT and mean of 12.9 million. The largest reported case, 90 million, contributes 53.67% of the 167.7 million total; the three largest contribute 71.56%. Figure 6 displays every amount with its S2 row ID. A leave-one-largest-out descriptive stress check gives a mean of 6.475 million, 49.81% below the full-table mean. This does not justify discarding the largest event. It shows that the aggregate benchmark is dominated by one record whose primary evidence is unavailable.

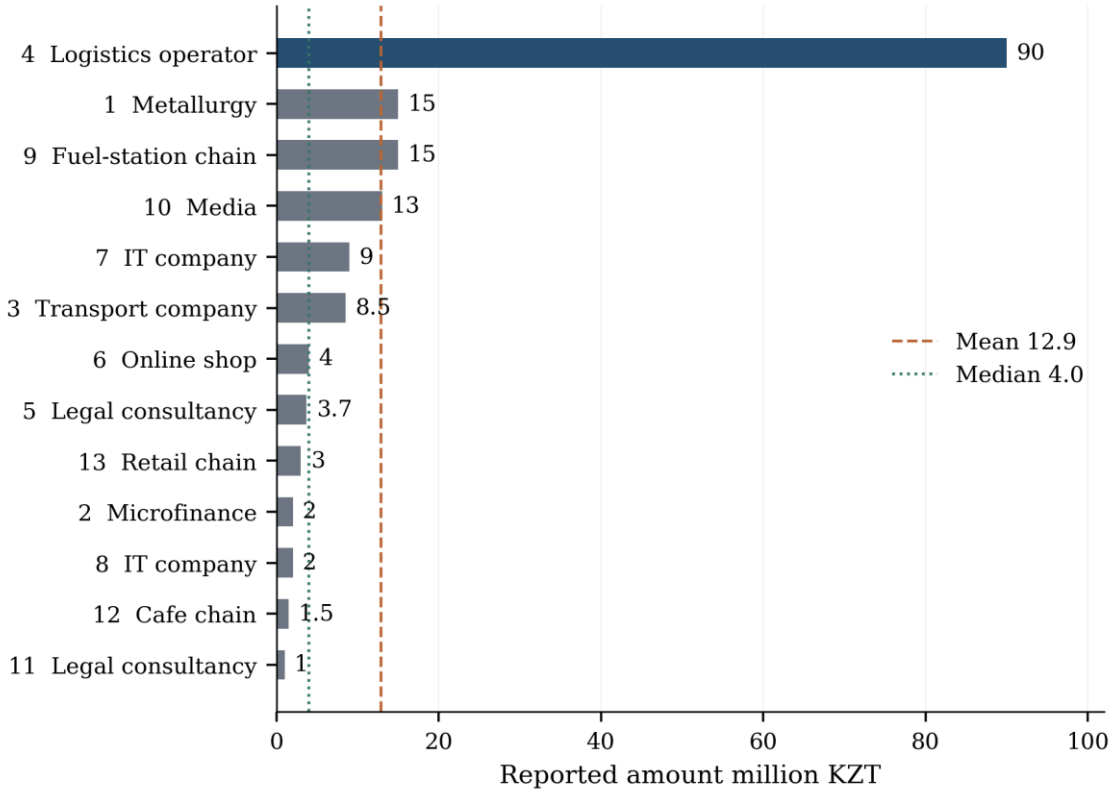


Figure 6. The thirteen organisation-reported amounts in the national draft, ordered by magnitude. IDs match S2.7. Mean and median describe this table only; the amounts have not been independently verified, price-normalised or generated by the proposed methodology. Sources: [14], [15], Appendix G.

Neither the mean nor a fitted distribution from this small secondary table would provide a defensible national coefficient. Sector labels mix different event mechanisms and accounting practices, while incident dates and sampling coverage are absent. The relevant numerical result is the observed concentration of the supplied table, not an estimate of tail risk in a population. This distinguishes a useful descriptive check from an unsupported calibration claim.

4 Discussion and validation requirements

4.1 Interpretation of the comparative findings

The comparative and numerical results support different levels of inference. Source passages establish the purposes and safeguards of particular instruments. The reconciliation and category-stability propositions establish conditional mathematical results. Case reconstructions show how those conditions operate with the supplied draft inputs. The secondary thirteen-case table provides only descriptive evidence. Treating these levels as

interchangeable would reproduce the same measurement problem that the assessment architecture is intended to prevent.

The strongest practical finding is that uncertainty in monetary valuation and stability of an escalation decision can coexist. In the portal case, removing a convention changes the financial category but need not change the operational decision. Conversely, a stable category based on a fixed expense subtotal does not establish that the wider economic total is accurate. Reports should therefore show both the valuation interval and the rule responsible for escalation; one overall score cannot convey both properties.

The formulas also localise the evidence needed to improve an estimate. For the portal, the valuation basis of reputation dominates the original total, while repeat-contact counts and additional time govern the revised third-party estimate. For the mine, the composition and level of the hourly downtime rate dominate the stated uncertainty width. In a dependency model, the admissibility of propagated items and the assessment horizon matter as much as convergence. These findings give concrete priorities for collecting evidence instead of suggesting that all missing information has equal value.

4.2 Validation and implementation

The first validation stage should assess whether independent analysts can reconstruct the same loss ledger from the same incident materials. Analysts should record disagreements about eligibility, causal attribution and valuation before reconciliation. Agreement can be assessed for categorical decisions, while absolute and relative discrepancies can describe monetary differences. A single reconciled total should not conceal initial disagreement. This stage tests reproducibility and clarity of the rules, rather than predictive accuracy.

The second stage should compare preliminary estimates with later documented outcomes. For a prospective pilot, record assessments at fixed milestones, for example 30, 90 and 180 days, while allowing a longer horizon for unresolved liabilities. Evaluate component-level errors and the coverage of any stated prediction intervals. Analyse sectors separately where consequence mechanisms differ. No particular sample size is asserted to be sufficient here: it should follow from the intended precision, variability and representation of sectors, including incidents with substantial downstream effects.

The third stage should evaluate whether non-monetary classifications support the intended decisions. A monetary threshold must not systematically suppress essential-service or safety escalation. Tests should vary valuation assumptions, index weights where used, and missing-data handling. A controlled comparison can ask whether the proposed ledger reduces duplicate entries or assessor disagreement relative to an existing procedure. Demonstrating these improvements would support an operational contribution beyond the conceptual synthesis presented here.

A national incident-loss register could support calibration if it records consistent scopes, revisions and evidence quality. Mandatory reporting can improve coverage, but does

not by itself remove reporting thresholds, selection bias or under-recorded social harms. A phased implementation should begin with auditable costs and clearly labelled estimates, then add calibrated sector-specific models as evidence accumulates. Access controls and disclosure rules are necessary because the underlying incident materials may contain commercially sensitive or personal information. Insurance research also highlights the importance of consistent information for loss assessment and risk transfer [51], [52].

4.3 Limitations

The source comparison is purposive and non-exhaustive, with 52 references and no independent human double coding. It cannot support a prevalence estimate for methodological defects. The two national drafts are unpublished objects of analysis with an institutional connection to the authors. The mathematical conditions are elementary conditional results; their contribution here is to make the accounting and classification requirements explicit and testable, rather than to claim that sensitivity analysis or interval classification is itself new.

No independently verified field application is reported. Reconstructed demonstration expenses are not inspected invoices, and the thirteen secondary reported amounts lack primary records and a collection protocol. The rectangular uncertainty sets are declared scenarios; dependencies could narrow them or make some combinations infeasible. The symmetric propagation model isolates one mechanism and does not reproduce a national economy. Greater algebraic detail and numerical reproducibility do not remove these limits; incident-based validation remains necessary to establish valuation accuracy and operational benefit.

5 Conclusion

Cyber incident damage assessment becomes more defensible when it begins with a declared measurement object and an auditable set of loss items. Existing standards, monetary models, regulations and infrastructure methods provide complementary components. Their integration requires explicit choices about perspective, time horizon, causal attribution and the treatment of uncertainty. The principal contribution of this review is a set of consistency requirements that connects these choices to evidence collection and reporting for public digital systems and critical infrastructure.

The quantitative analysis yields three operational conclusions. First, reconciliation must occur at the underlying loss-item level because category totals can conceal repeated observations. Second, monetary category stability can be tested against an explicit uncertainty set and must be distinguished from stability of non-monetary escalation. The portal's original monetary result has reputation elasticity 0.90956 and crosses the critical boundary at $\lambda = 0.42905$, while a key-service condition can preserve critical overall severity. Third, convergent propagation can produce downstream loss greater than primary loss; a universal cap would understate the controlled model by 20% at $a = 0.60$ and 60% at $a = 0.80$. These results establish reproducible consistency tests and priorities for further evidence

collection, while empirical accuracy remains to be evaluated against primary incident records.

Acknowledgments

The authors thank the specialists of the Digital Government Support Center who contributed to discussions of cyber incident damage assessment.

Author Contributions

Conceptualization: L.R. and A.T.; Methodology: L.R.; Formal analysis: A.T.; Investigation: L.R. and A.T.; Data curation: A.T.; Writing – original draft preparation: L.R.; Writing – review and editing: L.R. and A.T.; Visualization: A.T.; Supervision: L.R.

Conflicts of Interest

The authors are affiliated with the Digital Government Support Center and have an institutional interest in the development of national cybersecurity assessment practices. The proposed national application is presented as a methodological contribution rather than an independent evaluation of an adopted instrument. The authors declare no additional conflicts of interest.

Data Availability Statement

No primary incident dataset was collected. S1 supplies the complete reference inventory, source roles and evidence locators. S2 contains draft-case inputs, coefficient provenance, all thirteen reported-loss rows, the analytical conditions and computational settings. The accompanying Python scripts reproduce every numerical result and export all plotted data as CSV; figure generation uses Matplotlib and NumPy. Exact draft identities and source sections are listed in S1. The historical 68-source register and database screening history are not claimed.

Use of artificial intelligence

OpenAI ChatGPT was used for editorial restructuring, source-assisted evidence mapping, calculation scripting and figure preparation in this revision. This assistance is disclosed separately from the authors' contributions and does not constitute independent human double coding or empirical validation.

References

[1] IBM Security. (2026). Cost of a data breach report 2026. IBM Corporation. https://www.ibm.com/reports/data-breach

[2] Woods, D. W., & Böhme, R. (2021). SoK: Quantifying cyber risk. In 2021 IEEE Symposium on Security and Privacy (SP) (pp. 211-228). IEEE. https://doi.org/10.1109/SP40001.2021.00053

[3] Cremer, F., Sheehan, B., Fortmann, M., Kia, A. N., Mullins, M., Murphy, F., & Materne, S. (2022). Cyber risk and cybersecurity: A systematic review of data availability. The Geneva Papers on Risk and Insurance - Issues and Practice, 47(3), 698-736. https://doi.org/10.1057/s41288-022-00266-6

[4] Eling, M., & Wirfs, J. (2019). What are the actual costs of cyber risk events? European Journal of Operational Research, 272(3), 1109-1119. https://doi.org/10.1016/j.ejor.2018.07.021

[5] Romanosky, S. (2016). Examining the costs and causes of cyber incidents. Journal of Cybersecurity, 2(2), 121-135. https://doi.org/10.1093/cybsec/tyw001

[6] Gordon, L. A., & Loeb, M. P. (2002). The economics of information security investment. ACM Transactions on Information and System Security, 5(4), 438-457. https://doi.org/10.1145/581271.581274

[7] Sonnenreich, W., Albanese, J., & Stout, B. (2006). Return on security investment (ROSI): A practical quantitative model. Journal of Research and Practice in Information Technology, 38(1), 45-52.

[8] International Organization for Standardization. (2022). Information security, cybersecurity and privacy protection - Guidance on managing information security risks (ISO/IEC 27005:2022). ISO.

[9] Joint Task Force Transformation Initiative. (2012). Guide for conducting risk assessments (NIST Special Publication 800-30 Rev. 1). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-30r1

[10] The Open Group. (2025). Risk taxonomy (O-RT), Version 3.1 (Standard C251). https://publications.opengroup.org/c251

[11] Directive (EU) 2022/2555 of the European Parliament and of the Council of 14 December 2022 on measures for a high common level of cybersecurity across the Union (NIS 2 Directive). (2022). Official Journal of the European Union, L 333, 80-152.

[12] Regulation (EU) 2022/2554 of the European Parliament and of the Council of 14 December 2022 on digital operational resilience for the financial sector (DORA). (2022). Official Journal of the European Union, L 333, 1-79.

[13] Organisation for Economic Co-operation and Development, & Joint Research Centre. (2008). Handbook on constructing composite indicators: Methodology and user guide. OECD Publishing. https://doi.org/10.1787/9789264043466-en

[14] Methodology for the assessment and calculation of financial and material losses and operational consequences of cybersecurity breaches of infrastructure and digital systems. (2026). Unpublished draft, Appendix 2.6. Supplied file: Приложение 2.6 Методика оценки ущерба ИБ.docx. Sections 5–7 and Appendices B–G (Russian Б–Г). Astana, Kazakhstan (in Russian).

[15] Methodology for the assessment of damage from cybersecurity incidents. (2026). Unpublished working revision 2.2 with highlighted changes. Supplied file: Методика_оценки_ущерба_ред_2.2_с_выделением_правок.docx. Sections 6.1–6.3 and Appendices B, C, G (Russian Б, В, Г), pp. 73–77, 97–110. Kazakhstan (in Russian).

[16] National Institute of Standards and Technology. (2024). The NIST Cybersecurity Framework (CSF) 2.0 (NIST Cybersecurity White Paper 29). https://doi.org/10.6028/NIST.CSWP.29

[17] The Open Group. (2025). Risk analysis (O-RA), Version 2.1 (Standard C250). https://publications.opengroup.org/c250

[18] Freund, J., & Jones, J. (2014). Measuring and managing information risk: A FAIR approach. Butterworth-Heinemann.

[19] European Network and Information Security Agency. (2012). Introduction to return on security investment. ENISA.

[20] World Economic Forum. (2015). Partnering for cyber resilience: Towards the quantification of cyber threats. World Economic Forum.

[21] Orlando, A. (2021). Cyber risk quantification: Investigating the role of cyber value at risk. Risks, 9(10), 184. https://doi.org/10.3390/risks9100184

[22] Aldasoro, I., Gambacorta, L., Giudici, P., & Leach, T. (2022). The drivers of cyber risk. Journal of Financial Stability, 60, 100989. https://doi.org/10.1016/j.jfs.2022.100989

[23] Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data (General Data Protection Regulation). (2016). Official Journal of the European Union, L 119, 1-88.

[24] European Commission. (2024). Commission Delegated Regulation (EU) 2024/1772 of 13 March 2024 supplementing Regulation (EU) 2022/2554 with regard to regulatory technical standards specifying the criteria for the classification of ICT-related incidents and cyber threats, setting out materiality thresholds and specifying the details of reports of major incidents. Official Journal of the European Union, L, 2024/1772. https://eur-lex.europa.eu/eli/reg_del/2024/1772/oj

[25] European Commission. (2024). Commission Implementing Regulation (EU) 2024/2690 of 17 October 2024 laying down rules for the application of Directive (EU) 2022/2555 as regards technical and methodological requirements of cybersecurity risk-management measures and further specification of significant incidents for specified digital service providers. Official Journal of the European Union, L, 2024/2690. https://eur-lex.europa.eu/eli/reg_impl/2024/2690/oj

[26] International Organization for Standardization. (2019). Security and resilience - Business continuity management systems - Requirements (ISO 22301:2019). ISO.

[27] International Organization for Standardization. (2021). Security and resilience - Business continuity management systems - Guidelines for business impact analysis (ISO/TS 22317:2021). ISO.

[28] International Electrotechnical Commission. (2020). Security for industrial automation and control systems - Part 3-2: Security risk assessment for system design (IEC 62443-3-2:2020). IEC.

[29] Stouffer, K., Pease, M., Tang, C., Zimmerman, T., Pillitteri, V., Lightman, S., Hahn, A., Saravia, S., Sherule, A., & Thompson, M. (2023). Guide to operational technology (OT) security (NIST SP 800-82 Rev. 3). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-82r3

[30] Leontief, W. (1986). Input-output economics (2nd ed.). Oxford University Press.

[31] Rinaldi, S. M., Peerenboom, J. P., & Kelly, T. K. (2001). Identifying, understanding, and analyzing critical infrastructure interdependencies. IEEE Control Systems Magazine, 21(6), 11-25. https://doi.org/10.1109/37.969131

[32] European Union Agency for Network and Information Security. (2014). Methodologies for the identification of critical information infrastructure assets and services. ENISA.

[33] Saaty, T. L. (2008). Decision making with the analytic hierarchy process. International Journal of Services Sciences, 1(1), 83-98. https://doi.org/10.1504/IJSSCI.2008.017590

[34] International Organization for Standardization. (2018). Risk management - Guidelines (ISO 31000:2018). ISO.

[35] International Electrotechnical Commission. (2019). Risk management - Risk assessment techniques (IEC 31010:2019). IEC.

[36] FAIR Institute. (2025). Factor Analysis of Information Risk (FAIR) Standard, Version 3.0. January 2025. https://www.fairinstitute.org/hubfs/Standards%20Artifacts/Factor%20Analysis%20of%20Information%20Risk%20(FAIR)%20Standard%20v3.0%20(January%202025).pdf

[37] Anderson, R., Barton, C., Böhme, R., Clayton, R., Ganán, C., Grasso, T., … Vasek, M. (2019). Measuring the changing cost of cybercrime. In Proceedings of the 18th Workshop on the Economics of Information Security (WEIS 2019). Boston, MA.

[38] Verizon. (2025). 2025 data breach investigations report. Verizon Business.

[39] European Network and Information Security Agency. (2013). Recommendations for a methodology of the assessment of severity of personal data breaches. ENISA.

[40] The White House. (2016). Presidential Policy Directive 41: United States cyber incident coordination (Annex: Cyber incident severity schema). Washington, DC.

[41] European Supervisory Authorities. (2024). Joint Guidelines on the estimation of aggregated annual costs and losses caused by major ICT-related incidents under Regulation (EU) 2022/2554 (JC 2024 34). Version with completed regulatory references published by ESMA in March 2025. https://www.esma.europa.eu/sites/default/files/2025-03/JC_2024_34_Guidelines_on_the_estimation_of_aggregated_annual_costs_and_losses_caused_by_major_ICT-related_incidents.pdf

[42] Nelson, A., Rekhi, S., Souppaya, M., & Scarfone, K. (2025). Incident response recommendations and considerations for cybersecurity risk management: A CSF 2.0 community profile (NIST Special Publication 800-61r3). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-61r3

[43] International Electrotechnical Commission. (2016). Functional safety - Safety instrumented systems for the process industry sector - Part 3: Guidance for the determination of the required safety integrity levels (IEC 61511-3:2016). IEC.

[44] Young, W., & Leveson, N. G. (2014). An integrated approach to safety and security based on systems theory. Communications of the ACM, 57(2), 31-35. https://doi.org/10.1145/2556938

[45] Organisation for Economic Co-operation and Development. (2019). Good governance for critical infrastructure resilience. OECD Publishing. https://doi.org/10.1787/02f0e5a0-en

[46] Boyens, J., Smith, A., Bartol, N., Winkler, K., Holbrook, A., & Fallon, M. (2022, updated 2024). Cybersecurity supply chain risk management practices for systems and organizations (NIST SP 800-161 Rev. 1, Update 1). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-161r1-upd1

[47] Ghadge, A., Weiß, M., Caldwell, N. D., & Wilding, R. (2020). Managing cyber risk in supply chains: A review and research agenda. Supply Chain Management, 25(2), 223-240. https://doi.org/10.1108/SCM-10-2018-0357

[48] International Organization for Standardization. (2016). Information technology - Security techniques - Information security management - Monitoring, measurement, analysis and evaluation (ISO/IEC 27004:2016). ISO/IEC.

[49] Republic of Kazakhstan. (2026). Tsifrovoi kodeks Respubliki Kazakhstan [Digital Code of the Republic of Kazakhstan], No. 255-VIII, 9 January 2026. Adilet Legal Information System. https://adilet.zan.kz/rus/docs/K2600000255 (in Russian).

[50] Republic of Kazakhstan. (2013). O gosudarstvennykh uslugakh [On state services], Law No. 88-V, 15 April 2013. Adilet Legal Information System. https://adilet.zan.kz/rus/docs/Z1300000088 (in Russian).

[51] Dambra, S., Bilge, L., & Balzarotti, D. (2020). SoK: Cyber insurance - Technical challenges and a system security roadmap. In 2020 IEEE Symposium on Security and Privacy (SP) (pp. 1367-1383). IEEE. https://doi.org/10.1109/SP40000.2020.00019

[52] Kianpour, M., Kowalski, S. J., & Øverby, H. (2021). Systematically understanding cybersecurity economics: A survey. Sustainability, 13(24), 13677. https://doi.org/10.3390/su132413677