Bayesian Business Continuity Analytics for Cyber-Disrupted Industrial Process Workflows

Yisong Chen[1, *], Chuanhao Nie[1], Yexin Tian[1]

[1] College of Computing, Georgia Institute of Technology, Atlanta, GA, USA 30332

* Email: ychen841@gatech.edu (Corresponding Author)

Received 20 December 2025 · Revised 3 February 2026 · Accepted 20 March 2026 · Published 30 June 2026

Abstract

Industrial process workflows depend on interconnected operational technology, information systems, people, and external services. When a cyber incident interrupts or undermines confidence in one dependency, continuity decisions must be made before the state of the whole workflow is known. This article develops a Bayesian decision framework that connects uncertain cyber and physical states to production, quality release, contractual obligations, and recovery choices. A dynamic Bayesian network represents evolving asset and workflow states; evidence from security alerts, process observations, and verified recovery checks updates the posterior distribution. A decision layer evaluates feasible actions in terms of expected business loss, service shortfall, and, where appropriate, downside risk, subject to independently assessed safety and operational constraints. The framework makes explicit the difference between inferring a state, predicting an outcome under an action, and authorizing an intervention. A fully specified hypothetical example illustrates how an imperfect diagnostic signal changes a choice between controlled shutdown and a conditionally available isolated operating mode. With an initial compromise probability of 0.20 and the stated loss assumptions, the minimum expected loss falls from 28 to 20 thousand US dollars before a diagnostic cost; the expected value of the signal is therefore 8 thousand US dollars. This is an analytical illustration, not an empirical estimate or a claim of tested plant performance. The article discusses dependence, sparse data, calibration, model governance, and a prospective evaluation design. Its contribution is a transparent bridge from evolving evidence to safety-constrained continuity decisions, with a reproducible example and clearly stated limits on operational use.

Keywords: Bayesian networks; business continuity; industrial cyber-physical systems; operational technology; cyber resilience; workflow dependencies; decision analytics; value of information

Bayesian Business Continuity Analytics for Cyber-Disrupted Industrial Process Workflows

1. Introduction

Industrial production is sustained by sequences of interdependent tasks rather than by isolated machines. A typical production order passes through material confirmation, scheduling, process control, inspection, release, and dispatch. Each task relies on particular personnel, equipment, digital identities, configuration records, suppliers, and communications. A loss of confidence in a scheduling application may delay a process even when the physical equipment is undamaged; conversely, a compromised controller may invalidate product quality while enterprise dashboards continue to report output. Industrial control systems research therefore treats the combination of cyber, physical, and organizational dependencies as an essential feature of risk analysis (Ding et al., 2018; Jbair et al., 2022; Knowles et al., 2015).

Business continuity teams must make decisions under considerable ambiguity. An alarm may indicate malicious activity, a faulty sensor, an authorized maintenance action, or an interaction between these causes. A response that rapidly restores a server can be ineffective if the server's configuration or the identity infrastructure it depends on is untrustworthy. A manual operating mode may preserve a limited service, but only when trained staff, independent monitoring, required approvals, and plant-specific engineering limits are available. In a high-consequence process, production continuity cannot be treated as permission to relax physical safety requirements. Research on integrated process safety and security has similarly shown why cyber threats and physical safeguards should be examined together (Yuan, Reniers & Yang, 2024; Yuan, Yang & Reniers, 2024).

The decision problem changes over time. Initial information may be fragmented, and a decision maker may need to choose between isolating an asset, continuing a restricted operation, diverting work, or waiting for a verification result. Later evidence can alter the estimated compromise probability and reveal whether a backup, production batch, or dependent service can be trusted. A static risk register is useful for preparation but does not by itself explain how a decision should change when new evidence arrives. Dynamic Bayesian networks provide one way to represent temporal dependencies and revise beliefs, while recognizing that their validity depends on model structure and data quality (Tong et al., 2020; Zarei et al., 2017; Zhang et al., 2018).

Several research traditions contribute pieces of this problem. Studies of industrial cybersecurity model attack paths, detection, or process effects; supply chain studies analyze disruption propagation and resilience; and restoration models compare alternative recovery sequences (Heath et al., 2020; Hosseini & Ivanov, 2020; Ojha et al., 2018). However, connecting these perspectives for a specific industrial workflow requires decisions about the state of dependencies, the meaning of usable output, the cost of time, and the conditions under which an action is permissible. It also requires care not to interpret a statistical association as proof that a proposed intervention will work.

This article develops a decision-analytic framework for that connection. Its objectives are to represent cyber and process uncertainty in a dynamic Bayesian model; translate posterior beliefs into workflow service and business loss measures; compare feasible recovery actions while treating safety gates as independent requirements; and show, with a complete numerical example, how imperfect evidence can change a decision. The paper is a methodological and analytical contribution. It does not report plant observations, a deployed software product, field validation, or estimates of the frequency of real attacks. Its intended users are researchers and multidisciplinary continuity teams seeking a transparent basis for a later, site-specific implementation.

The remainder of the paper reviews the related literature, defines the workflow and decision problem, presents the Bayesian and decision models, works through a hypothetical case, and discusses implementation, validation, and limitations. The illustrative calculations are deliberately small enough to check by hand; the

dynamic model describes the general extension and is not presented as empirically validated by the small example.

2. Literature Review

2.1. Cyber risk and industrial process dependencies

Industrial cyber-physical systems couple digital communications and control logic with physical processes. Security assessments must consequently consider the timing of control actions, the validity of sensor and actuator data, and the potential interaction of cyber compromises with conventional equipment failures. Surveys of industrial control security and control-oriented detection provide foundations for differentiating these systems from ordinary enterprise information systems (Ding et al., 2018; Knowles et al., 2015). Threat-modeling research extends the analysis across system life cycles and dependencies, supporting the identification of attack surfaces before a disruption occurs (Jbair et al., 2022).

A further difficulty is that an observable deviation does not identify its cause with certainty. Physics-based monitors may detect inconsistency between a process model and measurements, but their effectiveness depends on model accuracy, disturbances, operating regimes, and possible manipulation of measurements (Giraldo et al., 2018). Fault-diagnosis literature also emphasizes that different diagnostic approaches make different assumptions about available data and system behavior (Dowdeswell et al., 2020). For continuity planning, a useful assessment therefore retains multiple hypotheses and does not equate an anomaly score with proof of an attack.

Bayesian approaches have been applied to industrial cybersecurity to represent uncertain dependencies and combine available evidence. A fuzzy-probability approach addresses dynamic cybersecurity risk in industrial control systems, while system-level dependency modeling links interactions to operational risk identification (Rotibi et al., 2025; Zhang et al., 2018). A recent review also considers context-based modeling, digital twins, and machine learning for industrial cybersecurity (Abraham et al., 2025). Work combining behavioral anomaly detection with Bayesian modeling and work distinguishing intentional attacks from technical faults are particularly relevant to a continuity team deciding what an alert actually implies (Chockalingam et al., 2021; Faramondi et al., 2023). These studies motivate probabilistic reasoning, but a posterior estimate still has to be connected to feasible actions and business outcomes.

2.2. Safety, process resilience, and recovery

Process safety research has developed dynamic Bayesian methods for changing operating conditions and accident scenarios. Examples include dynamic safety risk modeling, data-driven failure analysis, and adaptive networks incorporating process loss functions (Adedigba et al., 2017; Amin & Khan, 2022; Zarei et al., 2017). These methods support the use of process evidence, but practical deployment must preserve safety-engineering judgment, conservative treatment of uncertainty, and independence of protective systems. Combining safety and security analysis is important because the same intervention can change both exposure to an adversary and the availability of a safeguard (Yuan, Reniers & Yang, 2024; Yuan, Yang & Reniers, 2024).

Resilience is often defined in relation to the ability to maintain, adapt, and recover a valued function after disruption; its measurement depends on the selected performance variable and time horizon. Reviews of cyber-physical and process-industry resilience show that definitions, boundaries, and methods vary substantially (Cassottana et al., 2023; Ghaljahi et al., 2025). Dynamic Bayesian models can represent changes in system functionality and recovery over time and capture dependencies among process-industry facilities (Tong & Gernay, 2023; Tong et al., 2020); context-based production research emphasizes operator participation and heterogeneous dependencies (Bagozi et al., 2021). A resilience score alone, however, does

not say whether a batch is acceptable, whether a customer's critical order is fulfilled, or whether a proposed restoration action is safe.

Prior manufacturing work considers resilience of cyber-physical control systems and quantitative measures for production recovery (Alexopoulos et al., 2022; Moghaddam & Deshmukh, 2019). Research on active defenses in industrial control systems further illustrates the need to evaluate protective measures in relation to physical process operation (Chaves et al., 2017). These studies support recovery-oriented analysis, yet the appropriate objectives and permissions must still be specified for each process. This article therefore uses service and loss measures alongside independently established engineering constraints instead of treating any one resilience index as a universal decision rule.

2.3. Supply chain disruption and business continuity

Supply chain resilience research examines how disturbances propagate through connected organizations and how redundancy, substitution, and recovery can mitigate their effects. The ripple-effect literature explains why a localized disruption can alter capacity, inventory, and downstream delivery obligations (Dolgui et al., 2018; Hosseini et al., 2019). Bayesian network studies model dependencies between suppliers and risk propagation, including temporal recovery and multiple network layers (Garvey et al., 2015; Hosseini & Ivanov, 2020; Hosseini & Ivanov, 2022a; Hosseini et al., 2020; Ojha et al., 2018). Complementary research has developed resilience measures that account for propagation and recovery and considered cyber-physical interdependencies within supply networks (Chen et al., 2020; Hosseini & Ivanov, 2022b; Zhang et al., 2024).

The continuity problem addressed here differs in its unit of analysis: the production workflow and the business service it supports. Supplier and enterprise dependencies remain relevant, but the analysis must also distinguish control availability from trusted process execution and approved quality release. A restored machine is not sufficient if the records needed to verify a batch are unavailable. A supply network can be modeled in the same Bayesian architecture, but aggregating every dependency into a single "available" state would hide these distinctions. Prior restoration and network-risk optimization studies provide useful decision formulations, including sequencing under uncertainty and protection of interdependent facilities (Heath et al., 2020; Liu et al., 2024; Ruskey & Rosenberg, 2022).

A separate business continuity literature emphasizes the need to identify critical activities, assess consequences, and select controls suited to organizational contexts (Torabi et al., 2016). The present framework connects those concerns to time-indexed beliefs and explicit decision alternatives. It does not assume that resilience investment will always reduce loss: a redundant route can share a compromised identity provider, and an emergency procedure can impose its own delay, cost, or safety limitations. Such dependencies need to be represented and tested rather than assumed away.

2.4. Research gap and position of this article

The literature already contains Bayesian models for industrial security, safety, and supply-chain resilience, as well as optimization models for cyber-physical mitigation (Bhosale et al., 2024; Yuan, Yang & Reniers, 2024; Żebrowski et al., 2022). Consequently, the contribution of this article is not the invention of Bayesian networks or the observation that security affects business continuity. Its proposed contribution is the explicit assembly of four components: a workflow-level dependency map, posterior updating with evidence provenance, a feasibility screen for safety and authorization, and a loss-based comparison of recovery alternatives that can include the value of information.

This synthesis also clarifies several common analytical errors. First, alert probabilities cannot be multiplied as though sources were independent when they share a sensor, asset, or data pipeline. Second, a graph that accurately predicts observed outcomes does not automatically identify the effect of an intervention. Third, a low expected cost is not evidence that a proposal satisfies a process safety case. Fourth, simulated improvements must not be reported as measured operational gains. The framework addresses these issues

through explicit assumptions and governance requirements, while leaving site-specific causal identification and field validation as future work.

3. Problem Definition and Workflow Representation

3.1. System boundary and decision setting

Consider a manufacturer responsible for delivering a defined set of products or services over a planning horizon divided into periods t = 0, 1, …, T. The system boundary contains a set of production tasks, digital assets, physical assets, staff capabilities, and external dependencies. A task can be executable, restricted, unavailable, or unverified. Dependencies may include a manufacturing execution system, controller, engineering workstation, historian, quality record, trained operator, supplier, and logistics service. The boundary must be set around the business service of interest; an enterprise-wide total is not meaningful unless the relevant workflows and constraints are included.

The illustrative workflow in Figure 1 follows an order from authorization to production, quality release, and dispatch. It also shows shared information and security dependencies. The drawing is a conceptual representation, not a diagram of a particular plant. It highlights three kinds of edge: a prerequisite for executing a task, a source of evidence about a state, and a potential effect of a recovery action. These edges should be recorded separately because a monitoring relationship does not necessarily imply a physical production dependency, and an observed association does not automatically establish an intervention effect.

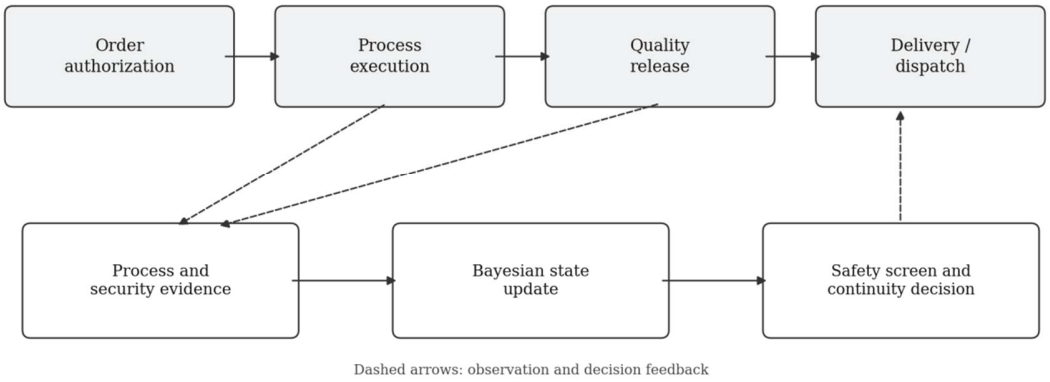


Figure 1. Conceptual Relationship between the Workflow, Evidence, and Continuity Decisions

Industrial workflows need more than a binary up-or-down description. For example, the production line may be technically operable while its recipe version is uncertain, and an order may remain blocked if batch records cannot be authenticated. A replacement controller can restore capacity but may still depend on a common engineering workstation. A paper-based workaround may be possible for a narrow product family while a regulated or high-risk product requires a validated digital record. The workflow map should therefore distinguish capacity, integrity, authorization, quality release, and downstream delivery.

A task may have AND prerequisites, where all required dependencies must be available, or OR alternatives, where approved substitutes are possible. A shared asset feeding multiple tasks creates a common cause of disruption. Recovery can be constrained by the availability of skilled personnel, specialized equipment, safe access, backup integrity, or vendor support. These are not generic probabilities: the owner of the model must define them with process engineers, cybersecurity personnel, quality staff, and business continuity managers. Research on interdependent infrastructure and manufacturing provides precedents for representing such shared functions (Hossain, Jaradat, et al., 2019; Hosseini et al., 2016; Moghaddam & Deshmukh, 2019).

3.2. States, observations, actions, and timing

Let X_t be the collection of latent states at the start of period t. Components can describe compromise status, process integrity, quality evidence, resource availability, and supplier condition. Let Y_t represent the evidence received during the period, and let a_t denote an action selected after reviewing that evidence. Examples of observations are a security alert, an independent physical reading, an authenticated backup check, and an operator's report. Examples of actions are controlled shutdown, a validated isolated operating mode, restoration from a verified image, expedited inspection, and reassignment of an order to an approved facility.

States and observations should not be conflated. "Alert present" is an observation; "controller compromised" is a latent hypothesis. "Backup verified" may be an observation derived from a defined verification procedure, while "backup usable for safe restoration" is a state requiring additional prerequisites. A stale or missing signal should be represented explicitly if its absence is informative. An evidence item should carry a source, timestamp, applicable operating regime, integrity status, and known dependence on other observations. In cases of possible evidence manipulation, the model should include a credibility or compromise state for the sensing pipeline where adequate information exists.

A decision epoch is defined by a practical opportunity to change the response, such as receipt of a diagnostic result or completion of an authorization check. Within the model, the sequence is: observe available evidence, update beliefs, determine feasible actions, select an action, and allow the state to evolve. Real operations may require immediate protection before full inference; a predefined emergency trip or other safety action takes precedence over this sequence. No Bayesian calculation should delay an established protective response or modify independent control logic without authorization.

The model distinguishes an action being theoretically listed from being available at a particular time. If qualified operators are absent, a workaround is unavailable even if it is attractive in expected cost. If independence of a protective system cannot be verified, the site's safety rules may prohibit a restart. Because these gates are determined outside the analytical optimizer, a missing approval is not converted into a low probability of harm and traded against revenue. The optimization is performed only over actions that pass the applicable gate.

3.3. Notation and assumptions

Table 1 summarizes the main notation. The general framework uses discrete time for clarity, although event-driven or continuous-time variants can be built where process dynamics require them. It assumes that the analyst can define a bounded set of operational states and a set of candidate actions, and can either estimate or elicit the distributions needed for inference. These assumptions are substantive limitations, not claims that industrial incidents are fully observable or that rare attacks follow stable historical frequencies.

Table 1. Principal Notation

Symbol	Definition
X_t	Unobserved cyber, physical, quality, and resource states at time t
Y_t	Observations available at time t, with provenance and timing
h_t	History of observations and actions available at the decision time
$b_t(x)$	Posterior probability that $X_t = x$ given the history
a_t and π	A candidate action and a multi-period decision policy
$A_{safe}(h_t)$	Actions independently approved under the applicable safety and authorization gates
d_{jt} and q_{jt}	Required and eligible, quality-released service for workflow j in period t

Symbol	Definition
L and R	Modeled loss and normalized priority-weighted service measure
EVSI	Expected value of sample information before diagnostic cost

The model treats the observation mechanism and transition mechanism as potentially dependent on operating conditions and past actions. The likelihood for a sensor during startup may differ from that during steady production, and a shutdown action can change the transition probabilities of an equipment state. When historical records do not support a reliable estimate, expert elicitation can provide a structured prior, preferably documented as ranges or distributions rather than a single confident number. The provenance, calibration period, confidence, and dependencies of every influential assumption should be recorded. Work on dynamic assessment and dependency-based operational risk provides a methodological basis for this transparent representation (Kanes et al., 2017; Rotibi et al., 2025; Zhu et al., 2018).

The framework does not presume that cyber events are independent of ordinary faults, that historical observations are representative of future adversarial behavior, or that all losses can be monetized accurately. Those are hypotheses to be examined. Where uncertainty cannot be bounded credibly, a conservative rule, human review, or additional information may be more appropriate than ranking actions on a fragile expected-value calculation. The small example later in this article imposes stronger assumptions so that its arithmetic is transparent.

4. Bayesian Continuity Analytics

4.1. Dynamic dependency model

A dynamic Bayesian network describes a joint probability distribution over system states and observations in successive periods. Let $b_t(x) = P(X_t = x \mid h_t)$ denote the belief in state x given the available history h_t, including earlier actions and observations. For a chosen action a_t, a transition model $P(x' \mid x, a_t)$ describes possible states in the next period, while an observation model $P(y' \mid x', a_t)$ describes the evidence expected from those states. The resulting Bayesian update is:

$$b_{t+1}(x') = [P(y_{t+1} \mid x', a_t) \times \Sigma_x P(x' \mid x, a_t) b_t(x)] / P(y_{t+1} \mid h_t, a_t). \quad (1)$$

The denominator normalizes the probabilities and is the sum of the numerator over all possible x′. This expression assumes the specified state is sufficient for the chosen Markov representation and that the observation model is valid for the action and operating regime. If a diagnostic action produces an observation before any meaningful state transition, a same-period Bayes update is the appropriate simplification. If the physical process has long memory, a state variable can be enlarged to retain relevant history or a different temporal model can be used.

A dynamic graph may contain separate nodes for an enterprise application, an identity service, a control function, production capability, quality evidence, and a customer order. Within each time slice, directed edges encode the conditional dependencies assumed by the model; between slices, they encode persistence and recovery. Bayesian network construction for industrial control and risk assessment shows how domain dependencies can be translated into probabilistic structures (Bhosale et al., 2024; Zhang et al., 2018; Zhu et al., 2018). Nevertheless, a graph is a model of assumptions: an arrow in a diagram is not sufficient evidence of physical causation, and a missing arrow can suppress an important common cause.

For example, two alerts produced by the same historian can be correlated even conditional on the true plant state. Treating them as independent may exaggerate certainty by counting the same underlying information twice. A practical graph can include their shared source and use a joint or conditional likelihood; if the data do not support this, an analyst can compare dependence assumptions and report a range of posteriors.

Likewise, a restored server and a restored workstation may fail together if they use the same authentication service. Shared dependencies should be represented as nodes or other explicit common factors rather than hidden in independent component reliabilities.

4.2. Learning from sparse and potentially unreliable evidence

In well-observed settings, conditional probability tables or parametric transition models can be estimated from logged incidents, exercises, testing, or maintenance records. In sparse settings, priors may combine engineering judgment with limited observations. For a binary event with a Beta(α, β) prior and n comparable trials containing k events, the posterior is Beta($\alpha + k$, $\beta + n - k$), provided the trials are conditionally independent and exchangeable under the model. The posterior mean is $(\alpha + k)/(\alpha + \beta + n)$. This simple update is illustrative; adversarial incidents are not automatically exchangeable, and evidence collected after a policy change may need a different model.

Expert judgments should be elicited in terms of clearly defined events, horizons, and conditioning situations. An estimate of "the probability of attack" is ambiguous unless the attack, exposed system, observation period, and existing controls are specified. Multiple experts can be asked for low, central, and high estimates, then their disagreements and areas of shared evidence can be recorded. Pooling should avoid treating repeated opinions based on the same source as independent observations. Sensitivity analysis across plausible priors is often more informative than claiming precision from one elicited table.

Evidence may be missing because a device is unavailable, a logging pipeline has been disabled, or communication has been intentionally interrupted. If the reason for missingness depends on the system state, simply ignoring the missing observation may bias the posterior. Conversely, the same indicator can reflect an authorized intervention or a hostile action. Distinguishing these possibilities requires maintenance records, independent process measurements, or further diagnostics when available (Chockalingam et al., 2021; Dowdeswell et al., 2020; Giraldo et al., 2018). Conflicting evidence should remain visible to reviewers rather than being hidden in a single score.

4.3. Mapping beliefs to service performance

Let d_{jt} be the quantity of service or output required for workflow j in period t, and let $q_{jt}(x, a)$ be the quantity that can be produced, verified, and released in that period under state x and the modeled action. The eligible quantity excludes output that cannot satisfy mandatory safety and quality conditions. For a single product family, a simplified representation is:

$q_{jt} = \min\{\text{production capacity, quality-release capacity, dispatch capacity, } d_{jt}\}$, subject to all required gates. (2)

This expression is a bottleneck approximation, not a full scheduling model. In a real plant, buffers, work in process, product substitutions, material yields, sequence-dependent setup times, and shared resources may require a flow or scheduling model. Capacity in the formula must refer to comparable units in the same period, and output cannot be counted as delivered merely because it was produced. A separate inventory balance is needed if an undelivered or unverified batch carries forward to a later period.

A normalized, horizon-specific continuity measure can be constructed as total priority-weighted eligible output divided by total priority-weighted demand. Let w_j be a documented nonnegative priority weight. Then $R = \Sigma_t \Sigma_j w_j q_{jt} / \Sigma_t \Sigma_j w_j d_{jt}$, provided the denominator is positive. If q_{jt} is bounded by d_{jt}, R lies between zero and one. Posterior scenarios yield a distribution of R rather than only a point value. The weights express an organization's chosen priorities; they are not an objective measure of the value of human safety, environmental protection, or legal compliance. Those matters remain independent constraints.

Other outputs should be reported alongside R: the duration below a minimum service level, the probability of missing a contractual delivery date, the time until a quality-approved product can be released, and the

probability of a permitted recovery by a specified time. These measures are defined with reference to a particular workflow, baseline, and horizon. Resilience measures and supply network analyses in prior research similarly depend on the performance function and disruption boundary selected (Chen et al., 2020; Hosseini & Ivanov, 2022b; Hosseini et al., 2019). Comparisons across facilities are meaningful only if definitions and operating contexts are comparable.

4.4. Decision objectives and safety gates

For an admissible action or policy π, a financial loss L_π can include incremental response expenditure, contribution margin lost on unserved demand, contractual penalties, additional logistics costs, and verified rework. The model should count each consequence once: for example, lost contribution and a contract penalty are distinct only if the accounting definition supports both, while revenue and contribution margin should not both be subtracted for the same undelivered unit. Deferred orders should be distinguished from permanently lost orders, and any terminal backlog must be valued consistently with the horizon.

A risk-neutral one-period decision chooses the action that minimizes posterior expected loss:

$$a^* = \arg\min_{a \in A_{safe}(h_t)} E[L(a, X_t) \mid h_t]. \quad (3)$$

Here $A_{safe}(h_t)$ is the set of actions that have passed applicable site safety, quality, legal, and authorization checks at the decision time. It may be empty if no modeled action is approved; in that case the optimizer cannot provide an operational instruction and the established incident procedure applies. A policy for multiple periods can instead minimize expected cumulative loss over feasible actions, including transitions and future observations. The one-period numerical example below does not solve this larger sequential problem.

If management is concerned about severe losses, the expected value can be complemented by a tail-risk measure such as conditional value at risk (CVaR). For confidence level α and loss L, CVaR can be calculated as $\min_\eta \{\eta + E[(L - \eta)_+]/(1 - \alpha)\}$. A weighted objective combining expected loss and CVaR is one possible choice, but the weight and α must be disclosed and reviewed. For discrete distributions, this optimization form correctly handles probability mass at the quantile. Financial tail-risk preferences must not replace non-negotiable safety restrictions; a small modeled probability of a catastrophic event is not by itself an authorization to proceed.

A site may also use an independently approved probabilistic threshold for a specified risk as one additional screening condition, but a modeled chance constraint cannot substitute for a safety case or mandatory rule. Its credibility depends on the validity of the probability model and the coverage of possible failure modes. When estimates are poorly supported, decision makers should examine conservative parameter ranges and compare the resulting recommendations. Robust Bayesian network and optimization research shows why data scarcity and model uncertainty deserve explicit treatment (Liu et al., 2021; Liu et al., 2024).

4.5. Value of information and sequential decisions

A diagnostic can be valuable because it changes an action, not simply because it produces a more precise probability. Before paying for a test Z, let C_0 be the minimum expected loss using current information. Let C_Z be the expected minimum loss after observing the test result, before the cost of the test. The expected value of sample information is $EVSI = C_0 - C_Z$. In this definition both optimizations use the same feasible action set and credible decision model. A diagnostic is financially attractive in the simplified risk-neutral setting when EVSI exceeds its total incremental cost, including delay and any extra operational exposure.

The theoretical expected value of perfect information provides an upper bound when the perfect and imperfect observations concern the same state, the action set is unchanged, and observation is costless and does not alter the system. Actual tests may delay recovery or expose an asset, so this bound does not imply that testing is safe or worthwhile. In a sequential setting, observations can change later actions and future evidence, making the timing of information important. The general framework can be viewed as a partially

observed decision problem, but solving one in practice may require approximations and careful checking of computational error.

Association and intervention must be separated at this stage. A model trained on observed restarts may reflect selection: teams may restart only relatively healthy systems. Applying that association to predict outcomes under a forced restart would be misleading. Action-specific transition distributions should be supported by engineering mechanisms, controlled exercises, appropriate observational identification, or carefully qualified expert judgment. If the data cannot identify an action's effect, the result is a scenario assumption, not an empirically verified causal effect. Mitigation frameworks and restoration studies illustrate the importance of tying decision alternatives to explicit assumptions (Heath et al., 2020; Żebrowski et al., 2022).

4.6. Governance and decision trace

Each recommendation should be accompanied by its decision time, workflow boundary, available actions, evidence provenance, posterior distribution, assumed action effects, material cost parameters, and safety-gate status. The record should identify the model version and the person or role authorized to approve an action. Operators must be able to inspect the main drivers, reject a recommendation, and record reasons for overriding it. An explanation should identify the thresholds and missing evidence that matter, rather than merely displaying a numerical probability with excessive decimal places.

During an incident, security and operations teams may disagree about whether evidence is trustworthy or whether an asset is ready. The model should preserve alternative scenarios and their consequences rather than suppressing disagreement in a single consensus estimate. A protected decision log supports later reconstruction, but storage and access controls must be appropriate for sensitive industrial information. Model outputs should be advisory until the organization has completed site-specific engineering, security, and governance reviews. The framework is not a safety instrumented system and is not intended to issue control commands.

5. Illustrative Numerical Analysis

5.1. Hypothetical workflow and assumptions

Consider a fictional process workflow whose scheduling and controller support environment has generated a security alert. A decision team must choose, for a fixed decision horizon, between a controlled shutdown S and an isolated operating mode M. The isolated mode is assumed to have passed all relevant engineering, staffing, authorization, and quality gates before it enters the choice set. If any gate fails, M is unavailable and the comparison below does not apply. The example uses no data from a real manufacturer and does not represent an actual incident.

Let C mean that the relevant support environment is compromised and N mean it is not compromised. The prior probability is $P(C) = 0.20$. A diagnostic produces either a positive indication A or a negative indication $\neg A$. Its sensitivity is $P(A \mid C) = 0.80$ and its false-positive probability is $P(A \mid N) = 0.10$. The diagnostic is assumed to leave the underlying state unchanged, generate its result before the action, and be available at an incremental cost of 2 thousand US dollars with no additional delay or safety impact. Those assumptions would require verification for a real diagnostic.

Losses for the fixed horizon are 32 thousand dollars for S in either state, 12 thousand dollars for M in state N, and 92 thousand dollars for M in state C. These are total modeled losses, not separate terms to be added to other costs. The higher loss for M under C represents a hypothetical combination of business disruption and unsuccessful containment or rework; it is not an estimate of physical harm. Table 2 provides all parameters needed to reproduce the example. All amounts in Tables 2–4 are expressed in thousands of US dollars unless a probability is shown.

Table 2. Hypothetical Diagnostic and Loss Assumptions

Quantity	Value	Interpretation
Prior	P(C) = 0.20	Compromise probability before the diagnostic
Positive if compromised	P(A | C) = 0.80	Diagnostic sensitivity
Positive if not compromised	P(A | N) = 0.10	False-positive probability
Shutdown S	32 in C or N	Total loss over the fixed horizon
Isolated mode M	92 in C; 12 in N	Total loss; available only after all required gates pass
Diagnostic cost	2	Incremental cost; no delay assumed
Timing	Result before action	Diagnostic does not alter the underlying state

5.2. Decisions without and with the diagnostic

Without the diagnostic, the expected loss of M is $0.20 \times 92 + 0.80 \times 12 = 28$, while the expected loss of S is 32. Because both are assumed feasible and the decision maker is risk neutral, M minimizes expected loss. In general, if p is the probability of C, the expected loss of M is 12 + 80p, so M and S are equal at p = 0.25. M is preferred for $p < 0.25$ and S for $p > 0.25$; at the threshold the two are tied. This is an economic decision threshold under the stipulated losses, not a safety threshold.

The probability of a positive diagnostic result is $P(A) = 0.20 \times 0.80 + 0.80 \times 0.10 = 0.24$. Bayes' rule gives $P(C \mid A) = 0.16/0.24 = 0.6667$. For a negative result, $P(\neg A) = 0.76$ and $P(C \mid \neg A) = 0.04/0.76 = 0.0526$. Consequently, the preferred feasible action after a positive result is S, and after a negative result it is M. Table 3 shows the posterior expected losses of both actions. The example contains one diagnostic; it does not assume that multiple security alerts are conditionally independent.

Table 3. Posterior Decisions and Expected Losses

Information	P(C | information)	Loss S	Loss M	Preferred
No diagnostic	0.2000	32.00	28.00	M
Positive A	0.6667	32.00	65.33	S
Negative ¬A	0.0526	32.00	16.21	M

The expected loss under this diagnostic-contingent policy, excluding the diagnostic cost, is $0.24 \times 32 + 0.76 \times 16.2105 = 20.00$. The same result follows by summing joint probabilities and associated losses: $0.24 \times 32 + 0.04 \times 92 + 0.72 \times 12 = 20.00$. The EVSI is therefore $28 - 20 = 8.00$, and the net expected benefit after the diagnostic cost is 6.00. The total expected loss with the diagnostic and its cost is 22.00. These values follow from the assumed loss matrix and probabilities; they do not establish the value of a real industrial monitoring tool.

For comparison, perfect knowledge of the state would lead to S in C and M in N, giving an expected loss of $0.20 \times 32 + 0.80 \times 12 = 16.00$. The expected value of perfect information is $28 - 16 = 12.00$, so the computed EVSI of 8.00 is consistent with the perfect-information upper bound under the stated assumptions. This arithmetic check is useful, but it is not a validation of the prior, likelihoods, losses, or action effects.

5.3. Sensitivity to the prior and loss assumptions

The recommendation and the value of the diagnostic depend on the prior probability. Table 4 varies P(C) with other inputs fixed; loss after testing excludes its cost, while net benefit subtracts the cost from EVSI. At a prior of 0.05, a positive test only just raises the posterior above the 0.25 decision threshold; the resulting EVSI is 0.50, below the assumed cost of 2. At a prior of 0.10, EVSI is 3.00 and the net benefit is 1.00. At a prior of 0.40, the no-test decision is S, but a negative result supports switching to M. At 0.60, the negative-result posterior equals the decision threshold, so the diagnostic has no decision value in this stylized setting.

Table 4. Sensitivity of Information Value to the Prior

Prior P(C)	Loss without test	Loss with test	EVSI	Net benefit
0.05	16.00	15.50	0.50	-1.50
0.10	20.00	17.00	3.00	1.00
0.20	28.00	20.00	8.00	6.00
0.40	32.00	26.00	6.00	4.00
0.60	32.00	32.00	0.00	-2.00

These results are non-monotonic. A more severe prior belief does not automatically imply a greater value of testing: at very high prior probabilities, both outcomes can lead to the same decision. Equally, a very reliable test may be of limited economic value when it cannot change any feasible choice. The conclusion depends on the assumed action costs and information quality. If the compromised-state loss of M were 72 rather than 92, the threshold would rise to 20/60 = 0.3333 and, at the baseline prior, EVSI would be 4.80. If that loss were 112, the threshold would be 0.20 and EVSI would be 11.20; before testing the two actions would be tied. These variants hold all other inputs fixed and ignore diagnostic cost when reporting EVSI.

The example does not model the diagnostic's effect on restoration time, sensor compromise, correlated observations, dynamic attacker adaptation, or multiple products. It also assumes the same feasible action set before and after testing. If a test confirms a mandatory condition for a workaround, then the action set itself changes and a simple comparison with the same-set EVSI definition is insufficient. In that situation, the diagnostic should be evaluated as part of a larger intervention pathway, including the cost and time of verification, and the safety gate remains independently binding.

6. Managerial and Operational Implications

6.1. From asset inventory to critical workflow mapping

A continuity team can start with one high-value or safety-relevant workflow instead of attempting to model the entire enterprise. The first task is to document the service, customer commitments, acceptable interruption, and actual prerequisites for production and release. Process engineers can map physical constraints; security specialists can identify digital assets and trusted paths; quality staff can identify records needed for release; and business owners can specify order priorities and financial consequences. The map should include off-site services and people, because a nominally redundant asset can remain dependent on the same personnel or shared identity service.

A useful dependency register records both a normal route and any approved substitute, the conditions for switching, the expected restoration sequence, and evidence required to declare the task trustworthy. For instance, a substitute communications route may be available only if a separate credential and a verified configuration are accessible. A quality inspection may continue manually but require later reconciliation

before release. These distinctions support exercise design and prevent an availability dashboard from presenting technically running equipment as a completed customer service.

The appropriate unit of business impact is a fulfilled obligation, not simply machine uptime. A line that runs at full speed while producing unreleasable material can increase loss through scrap, rework, and delayed deliveries. Conversely, a lower throughput route may protect a critical customer commitment if it maintains quality and traceability. Studies of supply network resilience and restoration provide useful analogies, but local workflows determine which dependencies and consequences are material (Heath et al., 2020; Hosseini et al., 2020; Zhang et al., 2024). Priority weights and acceptable service levels should be approved before an incident to reduce arbitrary decisions under pressure.

6.2. Evidence selection and recovery sequencing

The model can help identify evidence that is decision-relevant. A test of backup integrity is potentially useful when it determines whether a rapid restoration action is trustworthy; a further alert from the same compromised logging source may add little. Evidence should be evaluated for accuracy, timeliness, independence, and the potential for tampering. Before an incident, teams can test whether their supposedly independent diagnostic is actually dependent on a shared network, power supply, clock, authentication service, or configuration database.

Recovery sequencing should consider enabling relationships. Restoring an identity service, verified engineering environment, or validated record store may unlock several downstream tasks, but it can also create exposure if compromised trust relationships are reintroduced. A state-transition model should represent both benefits and residual risk, including the possibility that a restored dependency fails again. If two recovery tasks compete for the same specialist, a schedule that assumes simultaneous work is infeasible. Existing restoration and Bayesian supply-network methods provide useful starting points for the necessary resource and dependency constraints (Heath et al., 2020; Hossain, Nur, et al., 2019; Ruskey & Rosenberg, 2022).

The value of an observation is sensitive to the decisions available after it arrives. If no safe alternative exists, an additional diagnostic might improve situational awareness without changing the immediate action; its other benefits should be evaluated separately. If a test is slow enough to miss a contractual deadline, its delay cost can exceed the expected benefit of a better choice. Conversely, a quick, independently verified result can prevent an unnecessarily long shutdown or an ill-advised restart. None of these outcomes should be assumed in advance; their likelihood and operational cost require exercises and validation.

6.3. Resource allocation, tail risk, and shared dependencies

Before an incident, the same framework can support comparisons of measures such as an independently administered backup, additional cross-trained staff, protected quality records, and a qualified alternate site. Each measure changes specific model inputs: an independent backup may change recovery time and common-cause exposure; a trained team may change the set of feasible actions; and an alternate site may change the service capacity available during disruption. Modeling the mechanism explicitly helps prevent adding a generic “resilience improvement” percentage to every scenario.

Budget allocation should also reflect correlated failures. Two backups stored under the same administrative control may offer less protection against a common compromise than two truly separated recovery paths. A small number of shared services can dominate multiple workflows, but concentrating all resources on one high-centrality asset can be misleading if its loss is easy to bypass or its recovery requires unavailable expertise. Criticality depends on consequences, substitutability, time, and decision alternatives as well as graph position. Integrated cyber-physical network research and resilience measurement motivate attention to such dependencies (Chen et al., 2020; Hossain, Jaradat, et al., 2019; Zhang et al., 2024).

Expected loss provides a useful baseline for business planning, while tail measures and scenario analysis expose outcomes that the average can hide. However, a precise tail metric estimated from a very small sample can be misleading. Teams should report the assumptions that drive tail losses, specify exclusions, and consider stress scenarios outside the estimated distribution. A review of quantitative resilience methods and work on risk under incomplete data both suggest caution about relying on a single numerical summary (Hosseini et al., 2019; Liu et al., 2024). A result can be useful for comparison even when the correct decision is to gather more evidence or preserve a conservative operating state.

6.4. Coordination and accountability

The framework assigns different types of expertise to different parts of the decision. Control and process engineers define physical limits and approved procedures; cybersecurity personnel assess compromise hypotheses and evidence integrity; quality staff decide whether records and output satisfy release requirements; and business owners estimate service consequences. The decision authority then considers the model output within the organization's established incident governance. A model developer should not independently authorize physical operation merely because an optimization selects it.

Exercises should test the entire information-to-decision chain, including incomplete records, misleading alerts, conflicting reports, loss of a trusted clock, and unavailability of a key employee. After an exercise, the team should compare its initial dependencies and estimated action times with observed exercise outcomes, then record corrections. A model that cannot be maintained, understood, and challenged may increase cognitive burden during an incident. Context-based production resilience work emphasizes the role of the operator and the surrounding operating context (Bagozi et al., 2021); the decision trace proposed here gives that principle a practical expression.

7. Validation Strategy and Limitations

7.1. Proposed validation design

The present study contains an analytical example and does not validate a deployed model. A future evaluation should begin with a preregistered set of workflow boundaries, assumptions, decision alternatives, and outcome measures for a collaborating facility. Retrospective incident records can be separated chronologically into development and evaluation periods, with maintenance windows and known configuration changes recorded. A prospective tabletop or testbed exercise can then assess whether required evidence arrives in time and whether proposed actions are operationally feasible. Such an evaluation would require appropriate site approvals and must not disturb protective functions.

Probability forecasts should be tested for calibration and discrimination, not only whether the largest posterior state matches the final label. Calibration plots and proper scoring rules such as the Brier score can compare forecast probabilities against verified outcomes where enough independent cases exist. Confidence intervals should account for clustering by incident, asset, or facility, and rare outcomes should not be used to claim finely calibrated probabilities without sufficient data. Where ground truth is uncertain, the evaluation must describe how labels were established and report sensitivity to ambiguous cases.

Decision evaluation should compare pre-specified alternatives using the same information available at each decision time. Suitable comparators include a documented fixed response, a static risk assessment, and the Bayesian policy with and without diagnostic evidence. The outcomes may include realized or simulated service shortfall, time to a verified recovery, response cost, and the number of recommendations rejected by safety or quality gates. Counterfactual outcomes for unchosen actions are not directly observed; causal estimates require an appropriate design, and simulated counterfactuals should be reported as model-dependent.

Ablation tests can remove the workflow dependency map, common-cause nodes, or evidence-integrity model one at a time to examine whether they change decisions. Sensitivity studies should vary priors, observation likelihoods, recovery times, loss coefficients, dependence structures, and assumptions about attacker adaptation. An action ranking that reverses under small plausible changes should be labeled unstable. External assessment across different industries would help determine whether the representation transfers; a model calibrated for one plant should not be assumed to work in another without review.

7.2. Data, modeling, and causal limitations

Real industrial cyber incidents are relatively sparse in any one facility, and the cases most visible in logs may not represent undetected or unreported events. A Bayesian prior makes assumptions explicit but does not create missing evidence. An attacker can adapt to defenses and change the observation process, so stable historical transition and detection rates may not carry forward. Generic public incident data can inform hypotheses but may not match a site's architecture, configurations, staffing, or protective procedures.

The simplified graph may omit common dependencies or contain poorly specified conditional probability tables. Some variables are continuous and evolve on timescales much shorter than the chosen decision period; discretization can conceal transient excursions and unsafe physical dynamics. The network should therefore be integrated with, rather than substituted for, a validated process model when time-sensitive dynamics are material. Dynamic failure analysis and adaptive process-safety modeling show the value of using operational information while also illustrating the modeling choices required (Adedigba et al., 2017; Amin & Khan, 2022).

The quality of financial estimates is also limited. Losses can depend on demand substitution, capacity recovery, customer behavior, insurance terms, legal obligations, and the passage of time. Some consequences cannot be expressed responsibly as a monetary amount, and uncertain values may differ across stakeholders. The hypothetical example fixes these values to make the decision logic inspectable, but it does not demonstrate that its costs are realistic. Neither its dollar amounts nor its posterior probabilities should be transferred to an actual facility.

A final limitation is causal identification. The probability of a successful recovery among systems selected for restoration is not necessarily the probability that restoration would succeed for any system. Similarly, the presence of a contingency plan may correlate with other capabilities that affect outcomes. Action effects should be represented as assumptions unless supported by suitable evidence. Expert review, mechanistic engineering models, and exercises can reduce uncertainty, but they do not automatically establish general causal effects or prove the safety of a proposed action.

8. Conclusion

This article presents a Bayesian business continuity framework for industrial process workflows disrupted by cyber events. It links a time-indexed representation of uncertain cyber, physical, and organizational states to evidence updates, workflow service measures, admissible recovery actions, and business loss. Safety, quality, and authorization requirements define the action set before the economic comparison is performed. The framework also distinguishes observed associations from the action effects required for decision analysis.

In a hypothetical two-state example, an imperfect diagnostic changes the preferred action and has an expected value of 8 thousand US dollars before its 2-thousand-dollar cost. Sensitivity analysis shows that the value varies with the prior and can disappear when test outcomes do not alter the decision. These are results of the specified assumptions, not measurements from an operational facility. The example makes the calculation reproducible while illustrating why posterior beliefs, decision thresholds, and the cost of information should be considered together.

Future work should implement the model for a defined workflow, evaluate forecast calibration and decision performance against documented alternatives, examine common-cause failures and action-dependent observations, and verify proposed procedures with process-safety and operations specialists. A useful deployment would require credible local data, maintained dependency maps, transparent evidence provenance, and accountable human authorization. Until such validation is complete, the framework is a research and planning aid rather than a basis for automatic plant control.

Acknowledgement

We acknowledge the researchers whose published work informed this study. Any remaining errors are our responsibility.

Author Contributions

Conceptualization, Y.S.C.; Methodology, Y.S.C.; Formal Analysis, Y.S.C.; Investigation, Y.S.C. and C.H.N. and Y.X.T.; Data Curation, Y.S.C. and C.H.N. and Y.X.T.; Software, Y.S.C. and C.H.N. and Y.X.T.; Validation, Y.S.C. and C.H.N. and Y.X.T.; Resources, C.H.N. and Y.X.T.; Visualization, Y.S.C. and C.H.N. and Y.X.T.; Writing – Original Draft, Y.S.C.; Writing – Review & Editing, Y.S.C. and C.H.N. and Y.X.T.; Supervision, Y.S.C.; Project Administration, Y.S.C. We have read and approved the final manuscript and take responsibility for the work.

Funding

This research received no external funding.

Conflicts of Interest

We declare no conflicts of interest.

Ethical Approval and Consent

This analytical study reports no human-participant, personal-data, or animal research. Any future facility evaluation should obtain the approvals applicable to that study.

Data Availability Statement

The numerical analysis uses only the hypothetical values and formulas reported in Section 5 and Tables 2–4; no confidential operational dataset was used.

Declaration of Generative AI Use

Generative AI assisted us with drafting and editorial preparation. We take full responsibility for the accuracy of the content and references and for the final manuscript.

References

Abraham, D., Erceylan, G., Gkioulos, V., & Houmb, S. H. (2025). Towards enhanced cybersecurity in industrial control systems: A systematic review of context-based modeling, digital twins, and machine learning approaches. International Journal of Information Security, 24(6), 243. https://doi.org/10.1007/s10207-025-01158-1

Adedigba, S. A., Khan, F., & Yang, M. (2017). Dynamic failure analysis of process systems using principal component analysis and Bayesian network. Industrial & Engineering Chemistry Research, 56(8), 2094–2106. https://doi.org/10.1021/acs.iecr.6b03356

Alexopoulos, K., Anagiannis, I., Nikolakis, N., & Chryssolouris, G. (2022). A quantitative approach to resilience in manufacturing systems. International Journal of Production Research, 60(24), 7178–7193. https://doi.org/10.1080/00207543.2021.2018519

Amin, M. T., & Khan, F. (2022). Dynamic process safety assessment using adaptive Bayesian network with loss function. Industrial & Engineering Chemistry Research, 61(45), 16799–16814. https://doi.org/10.1021/acs.iecr.2c03080

Bagozi, A., Bianchini, D., & Antonellis, V. D. (2021). Context-based resilience in cyber-physical production system. Data Science and Engineering, 6(4), 434–454. https://doi.org/10.1007/s41019-021-00172-2

Bhosale, P., Kastner, W., & Sauter, T. (2024). AutomationML meets Bayesian networks: A comprehensive safety-security risk assessment in industrial control systems. IEEE Open Journal of the Industrial Electronics Society, 5, 823–835. https://doi.org/10.1109/ojies.2024.3439388

Cassottana, B., Roomi, M. M., Mashima, D., & Sansavini, G. (2023). Resilience analysis of cyber-physical systems: A review of models and methods. Risk Analysis, 43(11), 2359–2379. https://doi.org/10.1111/risa.14089

Chaves, A., Rice, M., Dunlap, S., & Pecarina, J. (2017). Improving the cyber resilience of industrial control systems. International Journal of Critical Infrastructure Protection, 17, 30–48. https://doi.org/10.1016/j.ijcip.2017.03.005

Chen, L., Dui, H., & Zhang, C. (2020). A resilience measure for supply chain systems considering the interruption with the cyber-physical systems. Reliability Engineering & System Safety, 199, 106869. https://doi.org/10.1016/j.ress.2020.106869

Chockalingam, S., Pieters, W., Teixeira, A., & van Gelder, P. (2021). Bayesian network model to distinguish between intentional attacks and accidental technical failures: A case study of floodgates. Cybersecurity, 4(1), 29. https://doi.org/10.1186/s42400-021-00086-6

Ding, D., Han, Q. L., Xiang, Y., Ge, X., & Zhang, X. M. (2018). A survey on security control and attack detection for industrial cyber-physical systems. Neurocomputing, 275, 1674–1683. https://doi.org/10.1016/j.neucom.2017.10.009

Dolgui, A., Ivanov, D., & Sokolov, B. (2018). Ripple effect in the supply chain: An analysis and recent literature. International Journal of Production Research, 56(1-2), 414–430. https://doi.org/10.1080/00207543.2017.1387680

Dowdeswell, B., Sinha, R., & MacDonell, S. G. (2020). Finding faults: A scoping study of fault diagnostics for industrial cyber–physical systems. Journal of Systems and Software, 168, 110638. https://doi.org/10.1016/j.jss.2020.110638

Faramondi, L., Flammini, F., Guarino, S., & Setola, R. (2023). A hybrid behavior- and Bayesian network-based framework for cyber–physical anomaly detection. Computers and Electrical Engineering, 112, 108988. https://doi.org/10.1016/j.compeleceng.2023.108988

Garvey, M. D., Carnovale, S., & Yeniyurt, S. (2015). An analytical framework for supply network risk propagation: A Bayesian network approach. European Journal of Operational Research, 243(2), 618–627. https://doi.org/10.1016/j.ejor.2014.10.034

Ghaljahi, M., Omidi, L., & Karimi, A. (2025). Resilience assessment in process industries: A review of literature. Heliyon, 11(4), e42498. https://doi.org/10.1016/j.heliyon.2025.e42498

Giraldo, J., Urbina, D., Cárdenas, Á., Valente, J., Faisal, M., Ruths, J., Tippenhauer, N. O., Sandberg, H., & Candell, R. (2018). A survey of physics-based attack detection in cyber-physical systems. ACM Computing Surveys, 51(4), Article 76. https://doi.org/10.1145/3203245

Heath, E. A., Mitchell, J. E., & Sharkey, T. C. (2020). Models for restoration decision making for a supply chain network after a cyber attack. The Journal of Defense Modeling and Simulation: Applications, Methodology, Technology, 17(1), 5–19. https://doi.org/10.1177/1548512918808410

Hossain, N. U. I., Jaradat, R., Hosseini, S., Marufuzzaman, M., & Buchanan, R. K. (2019). A framework for modeling and assessing system resilience using a Bayesian network: A case study of an interdependent electrical infrastructure system. International Journal of Critical Infrastructure Protection, 25, 62–83. https://doi.org/10.1016/j.ijcip.2019.02.002

Hossain, N. U. I., Nur, F., Hosseini, S., Jaradat, R., Marufuzzaman, M., & Puryear, S. M. (2019). A Bayesian network based approach for modeling and assessing resilience: A case study of a full service deep water port. Reliability Engineering & System Safety, 189, 378–396. https://doi.org/10.1016/j.ress.2019.04.037

Hosseini, S., Al Khaled, A., & Sarder, M. (2016). A general framework for assessing system resilience using Bayesian networks: A case study of sulfuric acid manufacturer. Journal of Manufacturing Systems, 41, 211–227. https://doi.org/10.1016/j.jmsy.2016.09.006

Hosseini, S., & Ivanov, D. (2020). Bayesian networks for supply chain risk, resilience and ripple effect analysis: A literature review. Expert Systems with Applications, 161, 113649. https://doi.org/10.1016/j.eswa.2020.113649

Hosseini, S., & Ivanov, D. (2022a). A multi-layer Bayesian network method for supply chain disruption modelling in the wake of the COVID-19 pandemic. International Journal of Production Research, 60(17), 5258–5276. https://doi.org/10.1080/00207543.2021.1953180

Hosseini, S., & Ivanov, D. (2022b). A new resilience measure for supply networks with the ripple effect considerations: A Bayesian network approach. Annals of Operations Research, 319(1), 581–607. https://doi.org/10.1007/s10479-019-03350-8

Hosseini, S., Ivanov, D., & Dolgui, A. (2019). Review of quantitative methods for supply chain resilience analysis. Transportation Research Part E: Logistics and Transportation Review, 125, 285–307. https://doi.org/10.1016/j.tre.2019.03.001

Hosseini, S., Ivanov, D., & Dolgui, A. (2020). Ripple effect modelling of supplier disruption: Integrated Markov chain and dynamic Bayesian network approach. International Journal of Production Research, 58(11), 3284–3303. https://doi.org/10.1080/00207543.2019.1661538

Jbair, M., Ahmad, B., Maple, C., & Harrison, R. (2022). Threat modelling for industrial cyber physical systems in the era of smart manufacturing. Computers in Industry, 137, 103611. https://doi.org/10.1016/j.compind.2022.103611

Kanes, R., Ramirez Marengo, M. C., Abdel-Moati, H., Cranefield, J., & Véchot, L. (2017). Developing a framework for dynamic risk assessment using Bayesian networks and reliability data. Journal of Loss Prevention in the Process Industries, 50, 142–153. https://doi.org/10.1016/j.jlp.2017.09.011

Knowles, W., Prince, D., Hutchison, D., Disso, J. F. P., & Jones, K. (2015). A survey of cyber security management in industrial control systems. International Journal of Critical Infrastructure Protection, 9, 52–80. https://doi.org/10.1016/j.ijcip.2015.02.002

Liu, M., Liu, Z., Chu, F., Zheng, F., & Chu, C. (2021). A new robust dynamic Bayesian network approach for disruption risk assessment under the supply chain ripple effect. International Journal of Production Research, 59(1), 265–285. https://doi.org/10.1080/00207543.2020.1841318

Liu, M., Tang, H., Chu, F., Ding, Y., Zheng, F., & Chu, C. (2024). A signomial programming-based approach for multi-echelon supply chain disruption risk assessment with robust dynamic Bayesian network. Computers & Operations Research, 161, 106422. https://doi.org/10.1016/j.cor.2023.106422

Moghaddam, M., & Deshmukh, A. (2019). Resilience of cyber-physical manufacturing control systems. Manufacturing Letters, 20, 40–44. https://doi.org/10.1016/j.mfglet.2019.05.002

Ojha, R., Ghadge, A., Tiwari, M. K., & Bititci, U. S. (2018). Bayesian network modelling for supply chain risk propagation. International Journal of Production Research, 56(17), 5795–5819. https://doi.org/10.1080/00207543.2018.1467059

Rotibi, A. O., Saxena, N., Burnap, P., & Read, C. (2025). System-level operational cyber risks identification in industrial control systems. Cyber-Physical Systems, 11(3), 265–296. https://doi.org/10.1080/23335777.2024.2373388

Ruskey, B., & Rosenberg, E. (2022). Minimizing risk in Bayesian supply chain networks. Computers & Industrial Engineering, 169, 108134. https://doi.org/10.1016/j.cie.2022.108134

Tong, Q., & Gernay, T. (2023). Resilience assessment of process industry facilities using dynamic Bayesian networks. Process Safety and Environmental Protection, 169, 547–563. https://doi.org/10.1016/j.psep.2022.11.048

Tong, Q., Yang, M., & Zinetullina, A. (2020). A dynamic Bayesian network-based approach to resilience assessment of engineered systems. Journal of Loss Prevention in the Process Industries, 65, 104152. https://doi.org/10.1016/j.jlp.2020.104152

Torabi, S. A., Giahi, R., & Sahebjamnia, N. (2016). An enhanced risk assessment framework for business continuity management systems. Safety Science, 89, 201–218. https://doi.org/10.1016/j.ssci.2016.06.015

Yuan, S., Reniers, G., & Yang, M. (2024). Integrated management of safety and security barriers in chemical plants to cope with emerging cyber-physical attack risks under uncertainties. Reliability Engineering & System Safety, 250, 110320. https://doi.org/10.1016/j.ress.2024.110320

Yuan, S., Yang, M., & Reniers, G. (2024). Integrated process safety and process security risk assessment of industrial cyber-physical systems in chemical plants. Computers in Industry, 155, 104056. https://doi.org/10.1016/j.compind.2023.104056

Zarei, E., Azadeh, A., Aliabadi, M. M., & Mohammadfam, I. (2017). Dynamic safety risk modeling of process systems using Bayesian network. Process Safety Progress, 36(4), 399–407. https://doi.org/10.1002/prs.11889

Żebrowski, P., Couce-Vieira, A., & Mancuso, A. (2022). A Bayesian framework for the analysis and optimal mitigation of cyber threats to cyber-physical systems. Risk Analysis, 42(10), 2275–2290. https://doi.org/10.1111/risa.13900

Zhang, L., Zhou, J., Ma, Y., Wang, X., & Zhang, F. (2024). Resilience improvement of cyber-physical supply chain networks considering cascading failures with mixed failure modes. Computers & Industrial Engineering, 187, 109812. https://doi.org/10.1016/j.cie.2023.109812

Zhang, Q., Zhou, C., Tian, Y. C., Xiong, N., Qin, Y., & Hu, B. (2018). A fuzzy probability Bayesian network approach for dynamic cybersecurity risk assessment in industrial control systems. IEEE Transactions on Industrial Informatics, 14(6), 2497–2506. https://doi.org/10.1109/tii.2017.2768998

Zhu, Q., Qin, Y., Zhou, C., & Gao, W. (2018). Extended multilevel flow model-based dynamic risk assessment for cybersecurity protection in industrial production systems. International Journal of Distributed Sensor Networks, 14(6), 1550147718779564. https://doi.org/10.1177/1550147718779564