processes

MDPI

Article

Multistage Production–Quality–Maintenance Optimization Under Uncertain Quality-Deviation Propagation: A K-Adaptability Approach

Honghong Jin and Jun Liu *

School of Mechanical and Electrical Engineering, Lanzhou University of Technology, Lanzhou 730050, China; jinhonghong@lut.edu.cn

* Correspondence: liuj@lut.edu.cn; Tel.: +86-138-9312-1598

Abstract

Uncertain quality-deviation propagation complicates the coordination of production, quality, and preventive maintenance in multistage manufacturing systems. This study develops a two-stage *K*-adaptability robust optimization model for integrated production–quality–maintenance decision-making under limited post-realization flexibility. The production–maintenance structure and a finite set of candidate operating policies are determined before uncertainty realization, while one policy is selected after the physical scenario is observed. Candidate policies are generated through active-scenario refinement, evaluated over the specified physical-scenario set, and compared by exhaustive subset selection within a finite common candidate-policy pool. Under the baseline uncertainty budget $\Gamma = 3$, removing zero-effect combinations yields 93 unique physical scenarios. In the four-stage benchmark, increasing the policy limit from one to two reduces the finite-pool certified worst-case total cost from 87,700.48 to 87,630.32 cost units (CU), a reduction of approximately 0.08%. Allowing a third policy does not further reduce this cost but improves $S^{\max}$, $S^{0.95}$, and S^{mean} measures by 11.21%, 19.23%, and 18.97%, respectively. A shortage-based lexicographic rule is used to distinguish cost-equivalent policy subsets in the four-stage benchmark. In the eight-stage refinement, increasing the policy limit from one to two reduces the finite-pool certified worst-case total cost by 5.46%, while a third policy provides no further improvement. These results show that the benefits of finite adaptability depend on the distinct scenario-specific responses provided by the candidate policies.

Keywords: multistage manufacturing system; production–quality–maintenance optimization; *K*-adaptability robust optimization; quality-deviation propagation; candidate operating policy

Check for updates

Academic Editor: Iqbal M. Mujtaba

Received: 21 September 2026
Revised: 17 October 2026
Accepted: 20 October 2026
Published: 22 October 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Modern manufacturing systems are increasingly characterized by high product variety, small production batches, and extensive customization, which have intensified the interactions among production capacity, product quality, equipment condition, and delivery requirements. In multistage manufacturing processes, upstream quality deviations may propagate downstream along the process chain and can be further affected by processing rates, stage loads, and equipment maintenance conditions. Consequently, when production, quality, and maintenance decisions are made independently, a locally favorable decision may propagate through subsequent stages and eventually lead to downstream quality losses or delivery risks. Coordinating production, quality, and preventive maintenance under uncertainty has therefore become an important issue in the operational optimization of multistage manufacturing systems.

Existing studies on multistage quality management have predominantly focused on quality prediction, condition monitoring, and deviation-propagation modeling. Ismail et al. [1] applied machine-learning techniques to quality monitoring in multistage manufacturing systems, Zhang et al. [2] modeled cross-stage machine dependencies using a graph-attention framework, and Moliner-Heredia et al. [3] developed a data-driven methodology for calibrating variation-propagation models using inspection data. Multi-task deep learning and neural-network architectures for continuous multistage processes have further improved the representation of accumulated quality effects and sequential dependencies across manufacturing stages [4,5]. Inspection-data-driven feedback and feedforward control, together with machining-error propagation knowledge graphs, have extended multistage quality analysis from prediction toward process control and knowledge representation [6,7]. For complex sequential manufacturing and multiple machining features, temporal deep-learning and domain-adaptation approaches have been employed to capture stage-transition dynamics and cross-feature quality relationships [8,9]. Yin et al. [10] further incorporated interstage quality inheritance into operational-parameter optimization, indicating that quality-propagation information can support not only state prediction but also process decision-making. Overall, these studies have substantially improved the identification and representation of cross-stage quality relationships and provide an important basis for embedding quality-propagation mechanisms into manufacturing decision models.

Meanwhile, the integrated optimization of production, maintenance, and quality control has received increasing attention. Ait-El-Cadi et al. [11] investigated the joint control of production, maintenance, and dynamic inspection in unreliable manufacturing systems, Hadian et al. [12] incorporated maintenance, buffer configuration, and quality control into an integrated planning framework, and Shi et al. [13] jointly optimized production-cycle, maintenance, and quality-inspection decisions. Under equipment degradation, Lv et al. [14] studied the coordinated optimization of production, maintenance, and quality control, whereas Paraschos and Koulouriotis [15] systematically reviewed learning-based approaches for production–maintenance–quality optimization in smart manufacturing. Koulinas et al. [16] further introduced machine learning into explainable integrated decision optimization. More recent studies have incorporated imperfect maintenance, dynamic inspection, and quality–reliability dependence into integrated production–maintenance–quality models [17,18], while Cao et al. [19] considered quality-based multilevel maintenance and buffer configuration in multistage small-batch production systems. These studies confirm the importance of jointly coordinating production, quality, and maintenance decisions. However, existing integrated models still provide limited explicit representation of the cross-stage propagation chain linking upstream quality deviations, stage loads, processing rates, and downstream nonconforming-product formation.

When parameter uncertainty further affects such coupled manufacturing mechanisms, designing decision structures with limited adaptability becomes another important issue. Robust optimization and its extensions provide systematic modeling tools for manufacturing decisions when accurate probability distributions are unavailable. Rahimian and Mehrotra [20] reviewed the foundations of distributionally robust optimization and its relationships with robust optimization, risk-averse optimization, and chance-constrained optimization; Tönissen et al. [21] investigated a column-and-constraint generation approach for two-stage problems; and Cohen et al. [22] applied adaptive robust optimization to parallel-machine scheduling. For large-scale or discrete uncertainty sets, scenario-reduction and adjustable robust optimization approaches have been developed to reduce computational complexity and enhance decision adjustability [23,24]. Metzker Soares et al. [25] applied adaptive robust optimization to lot-sizing under yield uncertainty, further demonstrating its applicability to manufacturing problems with stage-dependent uncertain information. The

K-adaptability approach restricts post-realization adjustment to selecting one policy from a finite set prepared before uncertainty realization [26]. More recently, machine-learning techniques have been introduced to accelerate two-stage adaptive robust optimization and *K*-adaptability search procedures [27,28]. Agra [29] investigated two-stage distributionally robust mixed-integer optimization with finite discrete support, whereas Hu et al. [30] developed a two-stage robust model and decomposition algorithm for parallel-machine selection and scheduling under uncertain release times.

Despite these advances, three research issues remain insufficiently addressed. First, existing studies generally treat quality propagation, production–maintenance coordination, and robust decision-making as related but largely separate modeling components. The joint representation of uncertain cross-stage quality-deviation propagation and its interaction with processing rates, stage loads, production capacity, preventive maintenance, nonconforming-product formation, and delivery shortage remains limited. In particular, quality deviation inherited from an upstream stage may affect downstream losses through nonlinear amplification and stage-dependent operating conditions, which creates a coupled decision structure that cannot be adequately represented by treating quality as an isolated prediction or inspection variable. Second, uncertainty in quality propagation and load-related parameters creates difficulty in determining production and maintenance decisions when reliable probability distributions are unavailable. A robust formulation can protect decisions against adverse realizations, but unrestricted post-realization adjustment may not be consistent with manufacturing practice because production–maintenance arrangements and operating plans often need to be prepared before the uncertain parameters are revealed. This motivates a finite adaptability structure in which a limited number of operating policies are prepared in advance, and one policy is selected after the physical scenario becomes known. Third, the computational treatment of such a finite-adaptability manufacturing problem requires more than solving a single robust optimization model. Candidate policies need to be generated under representative adverse scenarios, evaluated against the complete physical scenario set, and compared under a common candidate-policy pool. This is particularly important when the original model contains nonlinear quality-propagation and load-related relationships, because computational certification over a finite candidate pool should be distinguished from global optimality of the original continuous nonlinear *K*-adaptability problem.

To address these issues, this study develops a two-stage *K*-adaptability robust optimization framework for integrated production–quality–maintenance decision making in a multistage manufacturing system. The production–maintenance structure and candidate operating policies are determined before uncertainty realization; after a physical scenario is observed, one candidate policy is selected for implementation. The specific contributions are as follows. First, at the modeling level, the framework links linear quality-deviation propagation, nonlinear amplification, processing-rate deviation effects, and stage-load effects to stage capacity, preventive-maintenance-related capacity effects, nonconforming-product formation, maintenance-related costs, and delivery shortage under budgeted uncertainty. These relationships allow quality propagation and manufacturing decisions to be examined within a common optimization model. Second, at the decision level, the framework represents limited post-realization adaptability through a finite set of candidate operating policies. This structure distinguishes the production–maintenance commitments and operating policies prepared before uncertainty realization from the subsequent scenario-dependent policy selection. It therefore does not assume unrestricted adjustment of operating decisions after uncertainty is observed. Third, at the computational level, active-scenario refinement and candidate-policy augmentation are combined with complete physical-scenario evaluation and exhaustive subset selection within a finite

common candidate-policy pool. This provides a common numerical basis for comparing different policy limits. In the four-stage benchmark, an independently audited shortage-based lexicographic rule further distinguishes policy subsets with numerically equivalent primary worst-case costs.

2. Problem Description

2.1. System Configuration and Basic Assumptions

We consider a four-stage serial manufacturing system that produces a single product over a finite planning horizon. The stage set is denoted as $I = \{1, 2, 3, 4\}$.

Raw materials enter at Stage 1 and proceed sequentially through Stages 2, 3, and 4. Each stage performs its own machining operations, which may generate nonconforming products. These are removed before entering the next stage. Hence, upstream quality losses reduce the material available downstream and ultimately affect final conforming output and delivery performance.

The following assumptions are adopted throughout this work.

(1) The system is configured as a serial manufacturing line, with no parallel routings, bypasses, rework loops, or recirculation of nonconforming products;
(2) The processing capacity of each stage is determined jointly by its processing rate and the available production time;
(3) Preventive maintenance consumes production time and influences the stage-load surrogate within each maintenance interval;
(4) Quality deviations propagate from upstream to downstream stages, and are affected by both deviations of processing rates from their reference values and stage-load effects;
(5) Customer demand over the planning horizon is treated as known and deterministic, whereas the initial quality deviation, the nonlinear propagation coefficients, and the load-effect coefficients are all considered uncertain parameters;
(6) Before uncertainty is resolved, the production–maintenance structure together with a finite set of candidate operating policies must be fixed. After the physical scenario is observed, only one candidate operating policy from this pre-generated candidate-policy set may be selected for implementation.

The interactions among material flow, stage capacity, quality-deviation propagation, nonconforming-product formation, and delivery performance are illustrated in Figure 1.

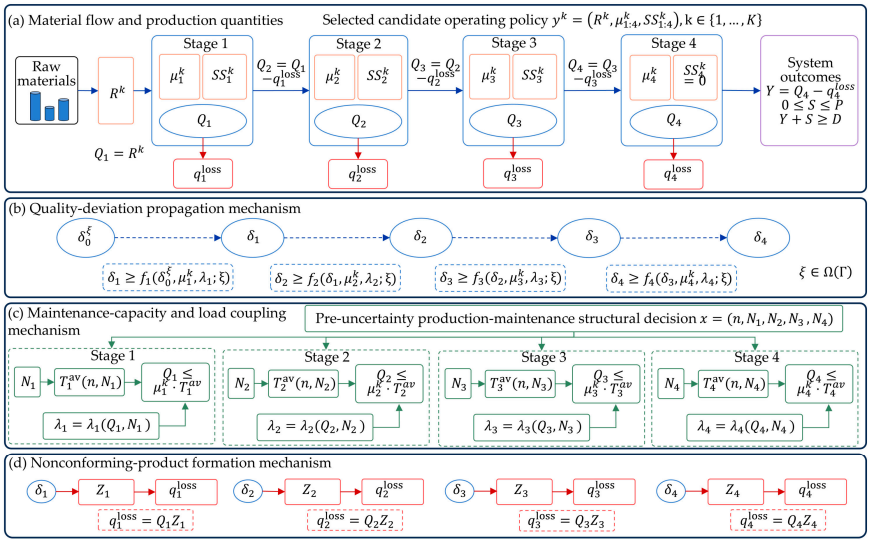


Figure 1. Material flow and production–quality–maintenance coupling in the multistage manufacturing system.

2.2. Sets and Parameters

2.2.1. Sets, Indices, and Auxiliary Notation

The sets, indices, and auxiliary notation are summarized in Table 1.

Table 1. Sets, indices, and auxiliary notation.

Symbol	Description
$I = \{1, 2, 3, 4\}$	Set of manufacturing stages
i	Manufacturing-stage index
$k = 1, \ldots, K$	Candidate operating policy index
$\Omega(\Gamma)$	Set of unique physical uncertainty scenarios under budget Γ
ξ	Physical scenario index
L	Set of uncertainty factors
P	Finite common candidate-policy pool
p	Candidate-policy index in the candidate-policy pool
η_i	Stage-specific maintenance-risk coefficient
ϖ	Weighting parameter for quality deviation and maintenance frequency
δ_i^{scale}	Normalization scale for quality deviation at Stage i

2.2.2. Production, Time, Maintenance, and Processing-Rate Parameters

The main production, planning, and maintenance parameters used in the four-stage benchmark are summarized in Table 2.

Table 2. Production, planning, and maintenance parameters.

Parameter	Value	Description
T_E	100.00	Planning horizon
D	720.00	Customer demand
Q_0	720.00	Reference production quantity in the stage-load function
t^{set}	1.50	Setup time per production cycle at each stage
t^{PM}	3.00	Preventive maintenance duration
Γ	3	Maximum number of simultaneously activated uncertainty factors

T_E and D specify the production and delivery requirements, while Q_0 provides the reference scale used in the stage-load surrogate. Setup and preventive maintenance reduce the production time available at each stage. The uncertainty budget Γ limits the number of simultaneously activated uncertainty factors and determines the range of physical scenarios considered in the robust optimization model.

The stage-specific processing-rate bounds and reference rates are given in Table 3.

Table 3. Processing-rate parameters.

Stage i	Lower Bound $\mu_i^{\min}$	Upper Bound $\mu_i^{\max}$	Reference Processing Rate μ_i^{opt}
1	8.00	13.50	10.00
2	8.00	13.00	10.00
3	7.50	12.50	9.50
4	7.50	12.00	9.50

The lower and upper bounds define the feasible processing-rate range at each stage, while μ_i^{opt} provides the reference condition used in the processing-rate-deviation term of the quality-propagation model.

2.2.3. Quality-Propagation Parameters

Quality evolution is characterized by the propagation of upstream deviations and their amplification caused by processing-rate deviation and stage-load effects. The parameters governing quality-deviation propagation are summarized in Table 4.

Table 4. Parameters governing quality-deviation propagation.

Stage i	α_i	$\overline{\beta}_i$	$\hat{\beta}_i$	γ_i	$\overline{\theta}_i$	$\hat{\theta}_i$	τ_i	ρ_i	$\delta_i^{\max}$
1	0.82	0.00	0.00	0.00035	0.0020	0.0005	1.05	2.10	0.12
2	0.96	8.00	4.00	0.00042	0.0025	0.0006	1.10	2.40	0.16
3	1.04	10.00	5.00	0.00050	0.0030	0.0007	1.15	2.70	0.21
4	1.10	12.00	6.00	0.00060	0.0035	0.0008	1.20	3.00	0.27

The coefficient α_i represents the linear propagation of upstream quality deviation, while β_i captures its nonlinear propagation. The coefficient γ_i characterizes the effect of processing-rate deviation, and θ_i represents the stage-load effect. The exponent τ_i controls the nonlinear response of the stage-load surrogate, whereas ρ_i determines the sensitivity of the nonconforming proportion to quality deviation. The parameter $\delta_i^{\max}$ specifies the upper bound of the quality deviation at Stage i. The parameter values used in the numerical benchmark are prescribed modeling inputs rather than estimates obtained from industrial production records. Their influence on the finite-pool optimization results is examined through the sensitivity analysis in Section 5.7.2.

The initial quality deviation is characterized by $\overline{\delta}_0 = 0.018$, $\hat{\delta}_0 = 0.010$.

The first stage does not receive upstream quality deviation. Therefore, the nonlinear propagation effect is disabled by setting $\overline{\beta}_1 = \hat{\beta}_1 = 0$.

For the maintenance-risk surrogate, the stage-specific coefficients and quality-deviation normalization scales are set as η = (0.18, 0.22, 0.26, 0.30), δ^{scale} = (0.12, 0.16, 0.21, 0.27), respectively, with ϖ = 0.55.

The values of τ_i, ρ_i, $\delta_i^{\max}$, η_i, δ^{scale} and ϖ are specified as benchmark parameters for numerical evaluation. They are used to represent stage-dependent load sensitivity, quality-loss sensitivity, admissible quality-deviation ranges, and maintenance-risk scaling.

2.2.4. Cost Parameters

The cost structure includes production setup, preventive maintenance, processing, nonconforming product, maintenance-risk surrogate, planned interstage buffer-demand surrogate, and shortage components. The relevant cost parameters for the four-stage benchmark are shown in Table 5.

Table 5. Cost parameters for the four-stage benchmark.

Stage i	C_i^{set}	C_i^{PM}	C_i^{p}	C_i^{q}	C_i^{CM}
1	380.00	560.00	2.00	90.00	420.00
2	420.00	620.00	2.20	105.00	500.00
3	460.00	680.00	2.50	120.00	590.00
4	500.00	760.00	2.80	135.00	680.00

The planned interstage buffer-demand surrogate cost coefficient and the shortage penalty coefficient are C^{inv} = 0.45, C^{sh} = 220.00. Here, C_i^{set} denotes the setup cost at Stage i, C_i^{PM} is the preventive-maintenance cost, C_i^{p} is the unit processing cost, and C_i^{p} is the unit cost associated with nonconforming products. The coefficient C_i^{CM} is applied to the maintenance-risk surrogate and does not represent an actual corrective-maintenance cost.

C^{inv} is the cost coefficient of the planned interstage buffer-demand surrogate, while C^{sh} penalizes delivery shortage. All costs are expressed in cost units (CU).

2.3. Decision Variables and System States

Under the two-stage decision structure, the variables fall into three categories: production–maintenance structural decisions, candidate operating policies, and system states. The production–maintenance structure and candidate operating policies are fixed before uncertainty is realized. Once a physical scenario is observed, one candidate operating policy is selected, while the production–maintenance structure stays unchanged.

The production–maintenance structural decision is written as:

$$x = (n, N_1, N_2, N_3, N_4) \tag{1}$$

where n is the number of production cycles within the planning horizon and N_i is the number of preventive-maintenance actions at Stage i.

The k-th candidate operating policy is written as $y^k = (R^k, \mu^k, SS^k)$, where $\mu^k = (\mu_1^k, \mu_2^k, \mu_3^k, \mu_4^k)$ and $SS^k = (SS_1^k, SS_2^k, SS_3^k, SS_4^k)$.

A candidate operating policy specifies the release quantity, processing rates, and planned interstage buffer-demand surrogates available for selection once the physical scenario is identified.

The information structure of the two-stage K-adaptability model is illustrated in Figure 2. The production–maintenance structure and candidate operating policies are determined before uncertainty realization. Once a physical scenario is observed, one candidate operating policy is selected, and the corresponding system states are evaluated.

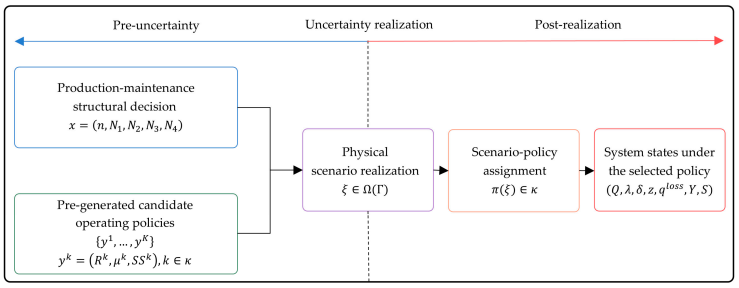


Figure 2. Information structure of the two-stage K-adaptability robust optimization model.

The candidate-policy variables and the resulting system states are summarized in Table 6.

Table 6. Definitions of decision variables and system states.

Symbol	Description
R^k	Initial release quantity under candidate operating policy k
μ_i^k	Processing rate at Stage i under candidate operating policy k
SS_i^k	Planned interstage buffer-demand surrogate after Stage i under candidate operating policy k
$Q_i^{k,\xi}$	Processing quantity at Stage i
$\lambda_i^{k,\xi}$	Stage-load surrogate at Stage i
$\delta_i^{k,\xi}$	Quality deviation at Stage i
$z_i^{k,\xi}$	Nonconforming proportion at Stage i
$q_i^{loss,k,\xi}$	Nonconforming quantity at Stage i
$Y^{k,\xi}$	Final conforming output
$S^{k,\xi}$	Shortage quantity

The system states are determined jointly by the selected candidate operating policy and the realized physical scenario. They describe the resulting production, quality, and delivery conditions and are obtained from the system relationships introduced in the following subsections.

2.4. Material Flow and Capacity Constraints

Material flow depends on the release quantity, processing capacity, and quality losses. The candidate operating policy specifies the initial release quantity and processing rates. The material transferred between stages is then affected by capacity limits and nonconforming-product losses.

At the beginning of the planning horizon, R^k is released to the first manufacturing stage:

$$Q_1^{k,\xi} = R^k \tag{2}$$

Nonconforming products may arise as quality deviations accumulate during processing. They are removed before the remaining material enters the next stage. The material balance is therefore:

$$Q_{i+1}^{k,\xi} = Q_i^{k,\xi} - q_i^{loss,k,\xi}, i = 1,2,3. \tag{3}$$

The model considers a serial production structure without rework or recirculation. Quality losses at one stage therefore reduce the quantity available to downstream stages.

The final conforming output is:

$$Y^{k,\xi} = Q_4^{k,\xi} - q_4^{loss,k,\xi} \tag{4}$$

This is the quantity available to satisfy customer demand after stage-wise quality losses have been removed.

Production setup and preventive maintenance consume part of the planning horizon. The available production time at Stage i is:

$$T_i^{av} = T_E - nt^{set} - N_i t^{PM} \tag{5}$$

For a selected processing rate, the stage processing quantity is limited by:

$$Q_i^{k,\xi} \le \mu_i^k T_i^{av}, i \in I. \tag{6}$$

A higher processing rate increases capacity, but it may also affect quality-deviation propagation. This effect is introduced in Section 2.5.

The processing rate of each candidate operating policy is restricted by the stage-specific bounds given in Table 3.

$$\mu_i^{\min} \le \mu_i^k \le \mu_i^{\max}, i \in I. \tag{7}$$

Different processing rates at adjacent stages may create an imbalance in the planned material flow. A planned interstage buffer-demand surrogate is used to represent this requirement:

$$SS_i^k \ge \frac{T_E}{2n}\left(\mu_i^k - \mu_{i+1}^k\right), i = 1,2,3. \tag{8}$$

with $SS_i^k \ge 0, SS_4^k = 0$.

The constraint gives the minimum planned buffer demand when the upstream processing rate exceeds the downstream rate. If the upstream rate is lower, the right-hand side becomes negative, and the non-negativity constraint becomes binding. No absolute-value relationship is imposed between adjacent processing rates.

These constraints define the material-flow and capacity structure of the multistage system. The resulting processing quantities depend on the selected candidate operating policy, available production time, and quality-related material losses.

2.5. Stage Load and Quality-Deviation Propagation

The production load at each stage depends on the processed quantity and the preventive maintenance schedule. A higher workload between maintenance actions may increase process sensitivity to quality variation.

To represent this relationship, the stage-load surrogate is defined as:

$$\lambda_i^{k,\xi} = \left[\frac{Q_i^{k,\xi}}{Q_0(N_i+1)} \right]^{\tau_i}, \quad i \in I. \tag{9}$$

For a fixed processing quantity, a larger N_i reduces the stage-load surrogate. Under a fixed maintenance schedule, a larger $Q_i^{k,\xi}$ increases it. Preventive maintenance therefore affects the system in two ways. It reduces available production time and changes the stage load involved in quality-deviation propagation. Equipment failure and maintenance-recovery processes are not modeled explicitly.

Previous studies have shown that second-order representations can be used to characterize nonlinear error propagation in multistage machining processes [7]. Motivated by this idea, a parsimonious second-order surrogate is adopted here to represent quality-deviation propagation while retaining upstream inheritance, nonlinear amplification, processing-rate deviation, and stage-load effects.

The quality deviation at the first stage is defined as:

$$\delta_1^{k,\xi} \geq \alpha_1 \delta_0^{\xi} + \gamma_1 (\mu_1^k - \mu_1^{opt})^2 + \theta_1^{\xi} \lambda_1^{k,\xi} \tag{10}$$

For Stages $i = 2, 3, 4$ the quality deviation propagation is formulated as:

$$\delta_i^{k,\xi} \geq \alpha_i \delta_{i-1}^{k,\xi} + \beta_i^{\xi} (\delta_{i-1}^{k,\xi})^2 + \gamma_i (\mu_i^k - \mu_i^{opt})^2 + \theta_i^{\xi} \lambda_i^{k,\xi} \tag{11}$$

The four terms represent different sources of variation in quality. The first term describes the inheritance of upstream quality deviation. The second term captures nonlinear amplification of deviations during propagation. The third term represents the quality effect caused by deviation of the processing rate from its reference value. The fourth term describes the additional quality impact associated with stage load.

The quality deviation is bounded by:

$$0 \leq \delta_i^{k,\xi} \leq \delta_i^{\max}, i \in I. \tag{12}$$

Therefore, the proposed formulation integrates upstream deviation propagation, nonlinear amplification, processing-rate deviation, and production-load effects to characterize quality evolution in the multistage manufacturing system.

2.6. Nonconforming-Product Formation and Delivery Performance

The accumulated quality deviation at each manufacturing stage determines the formation of nonconforming products. In practical multistage manufacturing processes, quality variation generated in an upstream operation may not immediately cause product rejection, but it can increase the proportion of nonconforming items generated during

subsequent processing. Therefore, a mapping relationship between quality deviation and nonconforming-product proportion is introduced:

$$z_i^{k,\xi} = 1 - \exp(-\rho_i \delta_i^{k,\xi}), i \in I. \tag{13}$$

where ρ_i represents the sensitivity of Stage i to quality deviation.

The exponential function describes the relationship between quality deviation and nonconforming-product formation. Small deviations have a limited effect on product conformity. As the deviation increases, the nonconforming proportion rises and gradually approaches its upper bound.

The quantity of nonconforming products generated at Stage i is:

$$q_i^{loss,k,\xi} = Q_i^{k,\xi} z_i^{k,\xi}, i \in I. \tag{14}$$

Nonconforming products are removed before the remaining material enters the next stage. They therefore do not return to the production flow. Only the remaining quantity proceeds to downstream operations. This setting represents a serial production process in which nonconforming products are separated after each stage.

At the end of the planning horizon, only the final conforming output is available to satisfy customer demand. The shortage quantity is constrained by:

$$Y^{k,\xi} + S^{k,\xi} \geq D, 0 \leq S^{k,\xi} \leq D. \tag{15}$$

The initial release quantity alone does not determine delivery performance. As material moves through the production stages, it is affected by capacity limitations and quality losses. Nonconforming-product removal further reduces the quantity available downstream. The remaining conforming output after the final stage determines the quantity available for delivery.

These relationships connect stage-level quality losses with the delivery performance of the manufacturing system.

2.7. Sources of Uncertainty and Scenario Construction

In practice, manufacturing quality is subject to several sources of variability. Incoming material quality may vary from batch to batch. Quality deviations may also propagate differently across manufacturing stages. In addition, the effect of production load on quality may vary with operating conditions.

Accordingly, three types of uncertain parameters are considered in the quality-deviation propagation model: the initial quality deviation, the nonlinear propagation coefficients, and the load-effect coefficients.

These uncertain parameters represent possible variations in manufacturing conditions rather than random disturbances with known probability distributions. Since accurate probability information is usually difficult to obtain for quality-propagation mechanisms, a budget-of-uncertainty approach is adopted to describe bounded parameter variations [31].

For a given physical scenario ξ the uncertain parameters are represented as:

$$\delta_0^{\xi} = \overline{\delta}_0 + \hat{\delta}_0 u_{d_0}^{\xi} \tag{16}$$

$$\beta_i^{\xi} = \overline{\beta}_i + \hat{\beta}_i u_{b_i}^{\xi}, i \in I. \tag{17}$$

$$\theta_i^{\xi} = \overline{\theta}_i + \hat{\theta}_i u_{t_i}^{\xi}, i \in I. \tag{18}$$

where $u_\ell^\xi \in \{0,1\}, \ell \in L$, indicates whether the corresponding uncertainty factor is activated in the physical scenario ξ.

The uncertainty budget is defined as $\sum_{\ell \in L} u_\ell^\xi \le \Gamma$.

Thus, at most Γ uncertainty factors can be active in a scenario. This avoids constructing scenarios in which too many adverse factors occur simultaneously.

For the four-stage benchmark, the uncertainty factor set is:

$$L = \{d_0, b_1, b_2, b_3, b_4, t_1, t_2, t_3, t_4\} \tag{19}$$

The uncertainty budget Γ limits the number of simultaneously activated adverse uncertainty factors. In the four-stage benchmark, these factors comprise one initial quality-deviation factor, four nonlinear propagation factors, and four stage-load factors.

With $\Gamma = 3$, the nine binary uncertainty factors generate $\sum_{j=0}^{3} \binom{9}{j} = 130$ initial factor-activation combinations, where j denotes the number of activated uncertainty factors. However, not all of these combinations represent distinct physical scenarios, because the nonlinear propagation term is disabled at the first manufacturing stage; the activation of uncertainty factor b_1 does not change any physical model parameter, and scenarios differing only in the activation status of b_1 produce identical system responses. After removing these physically redundant combinations, the final scenario set contains $|\Omega(\Gamma)| = \sum_{j=0}^{3} \binom{8}{j} = 93$ unique physical scenarios.

The same construction gives 9, 37, and 163 unique physical scenarios for $\Gamma = 1, 2, 4$, respectively. The uncertainty budget and perturbation magnitudes are prescribed numerical benchmark settings rather than statistically calibrated industrial uncertainty intervals. Their effects are examined through the sensitivity experiments in Sections 5.7.1 and 5.7.2.

3. Two-Stage K-Adaptability Robust Optimization Model

Section 2 described the material-flow, capacity, quality-deviation propagation, nonconforming-product formation, and delivery relationships of the multistage manufacturing system. Based on these relationships, a two-stage *K*-adaptability robust optimization model is formulated to minimize the worst-case total cost. Following the *K*-adaptability decision structure [26], a finite set of candidate operating policies is prepared before uncertainty is realized, and one policy is selected after the physical scenario is observed.

3.1. Structural and Operating Costs

The production–maintenance structural cost is:

$$C^I(x) = \sum_{i \in I} \left(n C_i^{set} + C_i^{PM} N_i \right) \tag{20}$$

It includes the production setup costs incurred over n production cycles and the preventive-maintenance costs associated with N_i maintenance actions at each stage.

To represent the combined effects of quality deviation and preventive-maintenance frequency on maintenance-related operational risk, a maintenance-risk surrogate is defined as:

$$\psi_i^{k,\xi} = \varpi \frac{\delta_i^{k,\xi}}{\delta_i^{scale}} + (1 - \varpi) \frac{\eta_i}{N_i + 1} \tag{21}$$

For candidate operating policy k under physical scenario ξ, the operating cost is:

$$C_{k,\xi}^{R} = \sum_{i \in I} C_i^p Q_i^{k,\xi} + \sum_{i \in I} C^{inv} SS_i^k + \sum_{i \in I} C_i^q q_i^{loss,k,\xi} + \sum_{i \in I} C_i^{CM} \psi_i^{k,\xi} + C^{sh} S^{k,\xi} \quad (22)$$

The cost covers stage processing, the planned interstage buffer-demand surrogate, nonconforming-product loss, maintenance-risk surrogate penalties, and shortage. Here, $C_i^{CM} \psi_i^{k,\xi}$ is not a corrective-maintenance cost.

3.2. Fixed-Policy Evaluation and K-Adaptability Formulation

For fixed x, y^k, and ξ, let $\mathcal{F}(x, y^k, \xi)$ denote the feasible set determined by the material-flow, capacity, quality-deviation propagation, nonconforming-product formation, and delivery constraints.

For this policy–scenario pair, the operating cost is:

$$\Phi(x, y^k, \xi) = \min_{v^{k,\xi}} \left\{ C_{k,\xi}^{R} : v^{k,\xi} \in \mathcal{F}(x, y^k, \xi) \right\} \quad (23)$$

where

$$v^{k,\xi} = (Q^{k,\xi}, \lambda^{k,\xi}, \delta^{k,\xi}, z^{k,\xi}, q^{loss,k,\xi}, \Upsilon^{k,\xi}, S^{k,\xi}) \quad (24)$$

If no feasible state exists for policy k under scenario ξ, set:

$$\Phi(x, y^k, \xi) = +\infty \quad (25)$$

For each physical scenario, the system selects the feasible candidate operating policy with the lowest operating cost from the pre-generated policy pool. The complete two-stage K-adaptability robust optimization model is formulated as:

$$\min_{x, y^1, \ldots, y^K} \left\{ C^I(x) + \max_{\xi \in \Omega(\Gamma)} \ \min_{k \in \mathcal{K}} \Phi(x, y^k, \xi) \right\} \quad (26)$$

To guarantee complete scenario feasibility, the following feasibility condition must hold, $\forall \xi \in \Omega(\Gamma)$, $\exists k \in \mathcal{K}$:$\mathcal{F}(x, y^k, \xi) \neq \emptyset$.

It should be noted that this requirement does not imply that every candidate operating policy must remain feasible for all physical scenarios. Instead, it only requires that each physical scenario be covered by at least one feasible candidate policy.

3.3. Model Interpretation

The proposed model captures three fundamental trade-offs among production, quality, and maintenance decisions. First, increasing the number of preventive maintenance actions reduces available production time but also simultaneously decreases the average production load carried within each maintenance interval. Second, increasing the processing rate enlarges available production capacity, whereas excessive deviations from the reference processing rate accelerate quality-deviation propagation. Third, the number of candidate operating policies determines the degree of finite adaptability. When $K = 1$, the model reduces to a conventional static robust optimization policy. When $K > 1$, different candidate operating policies are allowed to cover different physical scenarios, thereby reducing the conservatism associated with using a single policy to accommodate all uncertainty realizations.

Under an identical production–maintenance structure and policy search space, enlarging the allowable policy pool cannot theoretically worsen the optimal objective value.

However, the absence of further improvement in the primary objective does not mean additional candidate operating policies are unnecessary. Even when the worst-case total cost remains unchanged, newly introduced policies may still improve operational

performance by reducing the maximum shortage, the 95th-percentile shortage, or the equally weighted mean shortage across all enumerated physical scenarios.

4. Solution Method

The proposed model contains several nonlinear components, including power functions, quadratic quality-deviation propagation, exponential nonconforming-product formation functions, and bilinear products between production quantities and nonconforming proportions. Directly solving an integrated nonlinear optimization model over all physical scenarios with multiple candidate operating policies would result in a prohibitively large computational burden.

To address this issue, a solution framework consisting of scenario deduplication, nonlinear approximation, candidate-policy augmentation, full-scenario evaluation, and certification within a finite common candidate-policy pool is developed.

The policy-augmentation procedure is inspired by the critical-scenario identification mechanism of column-and-constraint generation for two-stage robust optimization [32]. Its purpose here is to generate candidate operating policies rather than to reproduce the upper- and lower-bound convergence mechanism of standard column-and-constraint generation (C&CG). Consequently, all reported results for different values of K are obtained through exhaustive certification within the finite common candidate-policy pool.

4.1. Treatment of Nonlinear Relationships

Once the production–maintenance structure has been fixed, the stage-load function becomes a univariate power function with respect to the production quantity. The following nonlinear relationships are approximated using piecewise linearization:

$$\lambda_i = \left[\frac{Q_i}{Q_0(N_i+1)}\right]^{\tau_i}, v_i = \delta_i^2, h_i = (\mu_i - \mu_i^{opt})^2. \tag{27}$$

Together with the exponential nonconforming-product formation function, $z_i = 1 - \exp(-\rho_i \delta_i)$, 41 breakpoints are used for both the quality-deviation and production-quantity functions, whereas 21 breakpoints are adopted for the processing-rate deviation function.

The nonconforming quantity remains represented by the bilinear relationship, $q_i^{loss} = Q_i z_i$, which is handled through the nonconvex quadratic optimization capability of the solver.

These piecewise-linear approximations are introduced solely for computational implementation and do not alter the physical relationships established in Sections 2 and 3.

4.2. Full-Scenario Evaluation of Fixed Policies

For a fixed production–maintenance structure, every candidate operating policy is evaluated under each of the 93 unique physical scenarios, yielding the policy–scenario operating cost and the corresponding shortage quantity.

$$C_{p\xi} = \Phi(x, y^{(p)}, \xi), S_{p\xi} = S(x, y^{(p)}, \xi). \tag{28}$$

If candidate operating policy p is infeasible under physical scenario ξ, its operating cost is assigned $C_{p\xi} = +\infty$.

Under a common set of model parameters, numerical approximations, and solver settings, the complete policy evaluation process produces the cost matrix and shortage matrix, $C = [C_{p\xi}]_{M\times|\Omega(\Gamma)|}, S = [S_{p\xi}]_{M\times|\Omega(\Gamma)|}$.

For an arbitrary nonempty policy subset A, the minimum operating cost under physical scenario ξ is defined as $c_A(\xi) = \min_{p \in A} C_{p\xi}$, and the corresponding cost-tie policy set is:

$$T_A(\xi) = \{p \in A : C_{p\xi} \leq c_A(\xi) + \varepsilon_{assign}\} \quad (29)$$

When several feasible policies have scenario costs within the assignment tolerance ε_{assign}, the shortage quantity is used to break the tie. If the shortage values are also tied, the policy with the smaller index is selected: $\pi_A(\xi) \in \text{lexargmin}_{p \in T_A(\xi)} (S_{p\xi}, p)$.

This rule affects only the scenario–policy assignment; it does not change the minimum scenario cost.

If no policy in A is feasible under scenario ξ, then A is infeasible for that scenario.

The worst-case total cost associated with candidate-policy subset A is:

$$Z(A) = C^I(x) + \max_{\xi \in \Omega(\Gamma)} c_A(\xi) \quad (30)$$

For each subset, the evaluation also records the shortage at the worst-cost scenario, the maximum shortage over all physical scenarios, the empirical 95th-percentile shortage, and the equally weighted mean shortage over the enumerated scenarios.

4.3. Active-Scenario Refinement and Candidate-Policy Augmentation

For $K = 1$, the initial active-scenario set contains the nominal scenario and four predefined seed scenarios: (d_0, b_2, b_3), (d_0, b_2, b_4), (d_0, b_3, b_4), (b_2, b_3, b_4).

Let N_{seed} denote the number of distinct high-cost scenarios selected to enrich the active-scenario set at each refinement iteration, and let $I_{\max}$ denote the maximum number of refinement iterations. When an additional candidate operating policy is generated, all previously obtained policies are retained. The current candidate-policy set is evaluated over the 93 unique physical scenarios, which are ranked according to their associated set costs. The current worst-cost scenario together with the next N_{seed} distinct high-cost scenarios is then used to update the active-scenario set. Cost ties between scenarios are resolved according to their canonical scenario indices.

At each refinement iteration, the candidate-policy generation problem is solved over the current active-scenario set. The resulting candidate operating policy is added to the existing candidate-policy set and subsequently evaluated over all 93 physical scenarios.

If the updated candidate-policy set yields a lower worst-case total cost over all physical scenarios, the newly generated policy is retained as the incumbent solution for the current policy-augmentation stage. If the newly identified worst-cost scenario is not already contained in the active-scenario set, it is added, and another refinement iteration is performed.

The active-scenario refinement procedure terminates when any one of the following conditions is satisfied:

(1) The current worst-cost scenario is already contained in the active-scenario set;
(2) The number of refinement iterations reaches $I_{\max}$;
(3) No feasible candidate operating policy is obtained within the prescribed time limit for single-policy optimization.

When the maximum refinement limit is reached, the candidate operating policy associated with the smallest worst-case total cost among all completed refinement iterations is retained. If no feasible new policy is generated during the current policy-augmentation stage, the procedure terminates for that stage, and the candidate-policy set from the preceding stage is preserved.

Because the initial active-scenario set, the refinement process, the computational time limits, and the nonlinear optimization procedure may all affect the candidate-policy generation trajectory, the objective value obtained from a single augmentation run is not used directly to compare different values of K. Instead, all final comparisons are performed through the finite common candidate-policy-pool certification described in Section 4.4.

4.4. Certification Within a Finite Common Candidate-Policy Pool and Secondary Tie-Breaking Rule

Using the predefined policy-deduplication tolerance ε_{policy}, numerically equivalent candidate operating policies are merged to construct the finite common candidate policy pool:

$$\mathcal{P} = \{y^{(1)}, \ldots, y^{(M)}\}, M = |\mathcal{P}|, \mathcal{I}_M = \{1, \ldots, M\}. \tag{31}$$

For a given policy limit K, the feasible collection of nonempty candidate-policy subsets is defined as:

$$\mathfrak{A}_K = \{A \subseteq \mathcal{I}_M : 1 \leq |A| \leq K\} \tag{32}$$

The corresponding optimal worst-case total cost within the finite candidate-policy pool is:

$$Z_K^{pool} = \min_{A \in \mathfrak{A}_K} Z(A) \tag{33}$$

When several candidate-policy subsets produce worst-case total costs that differ by no more than the objective tolerance ε_{obj}.

The set of cost-equivalent candidate-policy subsets is defined as follows:

$$\mathfrak{A}_K^0 = \left\{ A \in \mathfrak{A}_K : Z(A) \leq Z_K^{pool} + \varepsilon_{obj} \right\} \tag{34}$$

To prevent the selected candidate-policy subset from depending on the enumeration order, a deterministic lexicographic secondary tie-breaking rule is adopted. Among all cost-equivalent policy pools, the following criteria are minimized sequentially:

(1) The maximum shortage over all physical scenarios;
(2) The 95th-percentile shortage;
(3) The equally weighted mean shortage across all enumerated physical scenarios.

If all three service-risk indicators remain identical, the candidate-policy subset containing fewer candidate operating policies is preferred. If a tie still exists, the candidate-policy subset with the lexicographically smaller ordered policy-index vector is selected:

$$A_K^\star \in \operatorname{lexargmin}_{A \in \mathfrak{A}_K^0} \left(S_A^{max}, S_A^{0.95}, S_A^{mean}, \left|A\right|, ind(A) \right) \tag{35}$$

here, $ind(A)$ denotes the ordered vector of policy indices contained in candidate-policy subset A.

The shortage-based secondary rule is therefore applied only after the primary worst-case total cost has been made equivalent within the specified tolerance. It does not permit a material increase in the primary objective in exchange for improved shortage performance.

The shortage-based lexicographic rule in Equation (35) applies specifically to the four-stage benchmark. The archived eight-stage refinement adopts a different deterministic ranking procedure based on worst-case total cost, effective policy count, and policy-index ordering. The shortage indicators are reported as supplementary service-performance measures rather than used in the eight-stage subset-ranking procedure.

The term certification used throughout this paper has a strictly limited meaning. The certification is conducted under the following fixed conditions:

(1) A fixed production–maintenance structure;
(2) The 93 unique physical scenarios;

(3) The finite common candidate policy pool;
(4) The complete fixed-policy evaluation matrix;
(5) The numerical tolerances ε_{policy}, ε_{assign}, and ε_{obj}.

Under these conditions, all candidate-policy subsets satisfying $1 \leq |A| \leq K$ are exhaustively enumerated.

Therefore, the subset-selection problem is solved exactly in the combinatorial sense.

However, this certification does not imply that the candidate-policy generation procedure covers the continuous policy space, nor does it establish global optimality for either the original continuous *K*-adaptability model or the underlying nonlinear optimization problem. The certification results remain subject to the adopted piecewise-linear approximations and the numerical tolerances used in evaluating individual candidate operating policies. The overall workflow for candidate-policy augmentation, active-scenario refinement, and finite common candidate-policy-pool certification is illustrated in Figure 3.

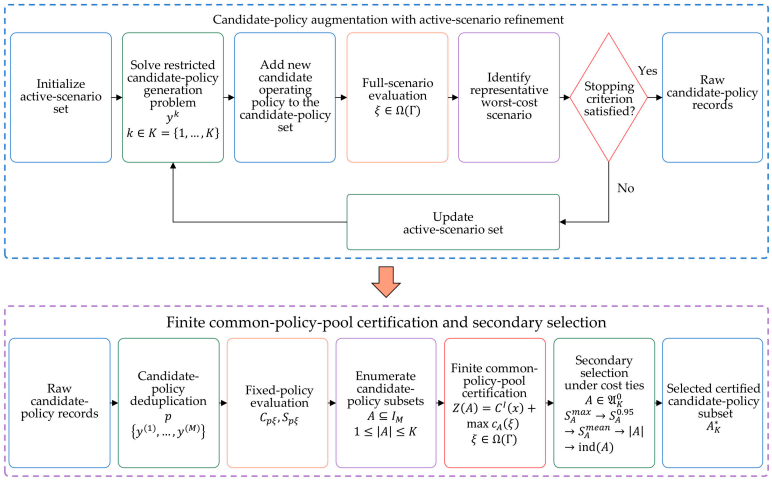


Figure 3. Algorithmic framework of candidate-policy augmentation with active-scenario refinement.

4.5. Eight-Stage Closed-Loop Computational Refinement

4.5.1. Eight-Stage Structure and Calibration Samples

The eight-stage case is constructed by dividing each of the four macro-stages into two consecutive substages:

$$I_8 = (1a, 1b, 2a, 2b, 3a, 3b, 4a, 4b) \tag{36}$$

The material-flow sequence and the main production–quality–maintenance relationships are retained. The two substages derived from the same macro-stage also share the corresponding uncertainty activation status.

The eight-stage configuration is a numerical refinement rather than an exact algebraic decomposition of the four-stage nonlinear quality-propagation model. In general, the composition of two nonlinear substage propagation relationships does not reproduce the corresponding macro-stage relationship for arbitrary operating conditions. The substage propagation and load-effect parameters are therefore calibrated against selected reference relationships and assessed through mapping and quantile-envelope validation, propagation-slack checks, and closed-loop stability diagnostics.

Let ϑ^r denote the parameter vector at refinement iteration r. It contains the propagation and load-effect parameters of the eight substages.

The calibration sample set is:

$$\mathcal{D}^r = \mathcal{D}^{base} \cup \mathcal{D}^{fb,r} \cup \mathcal{D}^{hist,r} \tag{37}$$

where D^{base} contains the samples mapped from the four-stage benchmark, $D^{fb,r}$ contains the feedback samples generated at iteration r, and $D^{hist,r}$ stores feedback samples retained from previous iterations.

Separate capacity limits and weights are used for the three sample groups. This prevents the accumulated feedback samples from dominating later calibration steps. The candidate parameter vector $\tilde{\vartheta}^{r}$ is fitted to the weighted sample set. A regularization term penalizes large departures from the parameter vector obtained in the previous refinement iteration.

4.5.2. Damped Parameter Updating and Independent Validation

The newly estimated parameter vector is not accepted directly. Instead, parameter updates are performed along a discrete damping path between the previous parameter vector and the newly calibrated one,

$$\vartheta^{r}(w) = (1-w)\vartheta^{r-1} + w\tilde{\vartheta}^{r} \tag{38}$$

$w \in \{1, 0.75, 0.50, 0.25, 0.125, 0.0625, 0\}$.

Candidate parameter vectors are examined sequentially according to the prescribed damping coefficients. A candidate update is accepted only if it simultaneously satisfies

(1) Independent mapping validation;
(2) Quantile-envelope validation;
(3) Feasibility verification over all 93 unique physical scenarios;
(4) The certification-based acceptance criterion.

Otherwise, the damping coefficient is reduced until an acceptable parameter update is obtained or the previous parameter vector is retained.

For each validation sample s, let u_s^{tar} denote the reference value obtained from either the four-stage benchmark or the envelope target, and let $\hat{u}_s$ represent the corresponding prediction generated by the eight-stage refinement model. The normalized relative error is defined as:

$$e_s = \frac{|\hat{u}_s - u_s^{tar}|}{\max(|u_s^{tar}|, \varepsilon_u)} \tag{39}$$

Both mapping validation and envelope validation report the following three performance indicators.

$$NRMSE = \sqrt{\frac{1}{n_v}\sum_{s=1}^{n_v} e_s^2} \tag{40}$$

$$Q_{0.95} = \text{quantile}_{0.95}\{e_s\} \tag{41}$$

$$e_{max} = \max_s e_s \tag{42}$$

The maximum normalized propagation slack is defined as:

$$\sigma_{max} = \max_{i,s}\left\{\frac{\max\{0, \delta_{is} - f_i(\cdot)\}}{\max(|\delta_{is}|, \varepsilon_\delta)}\right\} \tag{43}$$

To evaluate whether calibrated parameters accumulate excessively near the prescribed parameter bounds, the parameter-bound ratio is computed as:

$$r_{bd} = \frac{|\mathcal{B}_{\varepsilon_{bd}}|}{|J_\vartheta|} \tag{44}$$

where $B_{\varepsilon_{bd}}$ contains parameters located within ε_{bd} of either parameter bound, and J_ϑ contains all calibrated parameters.

4.5.3. Common Policy-Pool Updating and Decomposition Stability

Following each calibration iteration, the candidate-policy generation procedure is executed independently for $K = 1, 2, 3$, and all newly generated policies are merged into the accumulated finite common candidate-policy pool, $\mathcal{P}^{\,r} = \mathrm{dedup}\left(\mathcal{P}^{\,r-1} \cup Q_1^r \cup Q_2^r \cup Q_3^r\right)$, subject to $\mathcal{P}^{r-1} \subseteq \mathcal{P}^r$.

After deduplication, every candidate operating policy, regardless of whether it originates from the previous or current iteration, is re-evaluated under:

(1) The same calibrated parameter vector,
(2) The same fixed production–maintenance structure,
(3) The same set of 93 unique physical scenarios.

Direct comparison between objective values obtained from different calibration iterations is therefore avoided.

To distinguish the effects of parameter updating from those of policy-pool expansion, four certification cases are evaluated.

$$\begin{aligned} Z_K^{00} &= Z_K(\vartheta^{\,r-1}, \mathcal{P}^{\,r-1}), \\ Z_K^{10} &= Z_K(\vartheta^{\,r}, \mathcal{P}^{\,r-1}), \\ Z_K^{01} &= Z_K(\vartheta^{\,r-1}, \mathcal{P}^{\,r}), \\ Z_K^{11} &= Z_K(\vartheta^{\,r}, \mathcal{P}^{\,r}). \end{aligned} \tag{45}$$

The effects of parameter updating, candidate-policy-pool expansion, and their interaction are evaluated by the following normalized measures:

$$\begin{aligned} \Delta_K^{par} &= \frac{\left|Z_K^{10} - Z_K^{00}\right|}{\max\left(\left|Z_K^{00}\right|, \varepsilon_Z\right)}, \\ \Delta_K^{pool} &= \frac{Z_K^{10} - Z_K^{11}}{\max\left(Z_K^{10}, \varepsilon_Z\right)}, \\ \Delta_K^{int} &= \frac{Z_K^{11} - Z_K^{10} - Z_K^{01} + Z_K^{00}}{\max\left(Z_K^{00}, \varepsilon_Z\right)}. \end{aligned} \tag{46}$$

The closed-loop refinement process terminates only when all four indicators meet the predefined acceptance thresholds:

(1) The mapping error;
(2) The envelope error;
(3) The propagation slack;
(4) The parameter-bound ratio.

Both convergence indicators also satisfy the required thresholds, with $\max_{K=1,2,3} \Delta_K^{par} \leq 0.5\%$ and $\max_{K=1,2,3} \Delta_K^{pool} \leq 0.5\%$.

These conditions must be satisfied for two consecutive refinement iterations before the closed-loop procedure is terminated.

This stopping criterion indicates that both parameter updates and the marginal benefit of policy-pool expansion have become numerically negligible under the prescribed thresholds. It should not be interpreted as a proof of global convergence in either the continuous parameter space or the continuous policy space.

4.6. Service-Risk Indicators

The certified worst-case total cost is used as the primary criterion for evaluating candidate-policy subsets. To characterize delivery performance, the shortage outcomes associated with the selected candidate-policy subset A^* are also evaluated over all physical scenarios.

For each physical scenario ξ, let $s(\xi) = S_{\pi_{A^*}(\xi),\xi}$ denote the shortage quantity under the candidate operating policy assigned by A^*. The assignment $\pi_{A^*}(\xi)$ follows the deterministic scenario–policy assignment rule defined in Section 4.2.

The shortage performance of a selected candidate-policy subset is summarized by four indicators: the shortage under the worst-cost scenario, the maximum shortage, the empirical 95th-percentile shortage, and the equally weighted mean shortage, defined as:

$$\begin{aligned} S^{WC} &= s(\xi^{WC}), \\ S^{\max} &= \max_{\xi \in \Omega(\Gamma)} s(\xi), \\ S^{0.95} &= \text{quantile}_{0.95}\{s(\xi) : \xi \in \Omega(\Gamma)\}, \\ S^{\text{mean}} &= \frac{1}{|\Omega(\Gamma)|} \sum_{\xi \in \Omega(\Gamma)} s(\xi). \end{aligned} \tag{47}$$

These indicators describe different aspects of delivery performance. S^{WC} measures the shortage under the physical scenario that determines the certified worst-case total cost, whereas $S^{\max}$ represents the largest shortage observed over the complete physical-scenario set. These two indicators need not correspond to the same physical scenario. $S^{0.95}$ characterizes the upper tail of the shortage distribution, while S^{mean} provides an equally weighted descriptive measure across all enumerated physical scenarios.

No probability distribution is assigned to the physical scenarios. Accordingly, S^{mean} is not interpreted as an expected shortage. In the four-stage benchmark, $S^{\max}$, $S^{0.95}$, and S^{mean} are used for secondary selection among cost-equivalent policy subsets. The shortage under the worst-cost scenario is reported as an additional descriptive measure. In the eight-stage refinement, the shortage indicators are used to evaluate delivery performance rather than to rank candidate-policy subsets.

5. Numerical Experiments and Results Analysis

5.1. Experimental Setup for the Four-Stage System

The four-stage benchmark experiment used a fixed production–maintenance structure, $n = 1$, $N = (2, 3, 2, 1)$.

The model was implemented in Python3.6.5 and solved with Gurobi. Solver parameters were set to NonConvex = 2, MIPGap = 10^{-4}, and NumericFocus = 1.

Candidate operating policies were generated through three independent augmentation runs with different computational budgets. Table 7 lists the corresponding settings.

Table 7. Candidate-policy generation settings for the four-stage model.

Run	Policy-Generation Time Limit (s)	Fixed-Policy Evaluation Time Limit (s)	N_{seed}	$I_{\max}$
A	60	15	1	3
B	180	30	3	5
C	180	30	1	5

5.2. Scenario Audit and Statistics of the Finite Common Candidate-Policy Pool

After the physical deduplication described in Section 2.7, the uncertainty set contains 93 unique physical scenarios. Candidate-policy augmentation generated 18 raw candidate-policy records. After removing numerically equivalent policies, five unique candidate operating policies remained in the finite common candidate-policy pool p.

Each candidate operating policy was evaluated independently under all 93 physical scenarios. This produced 465 fixed-policy–scenario evaluations, from which the policy–scenario cost and shortage results used in the subsequent certification were obtained. The corresponding scenario and candidate-policy-pool statistics are summarized in Table 8.

Table 8. Finite common candidate-policy-pool audit and certification size.

Item	Value
Number of unique physical scenarios	93
Number of raw candidate-policy records	18
Number of unique candidate operating policies after deduplication	5
Number of fixed-policy–scenario evaluations	465
Number of one-policy subsets, $\binom{5}{1}$	5
Number of two-policy subsets, $\binom{5}{2}$	10
Number of three-policy subsets, $\binom{5}{3}$	10
Total number of nonempty policy subsets with $1 \leq \lvert A \rvert \leq 3$	25

For $K \leq 3$, a total of 25 nonempty candidate-policy subsets are included in the certification. All 93 physical scenarios are feasible under at least one policy in p. The finite-pool certification is therefore conducted on the complete policy–scenario evaluation matrix without additional feasibility repair.

5.3. Finite-Pool Certification Results for the Four-Stage Benchmark

The five candidate operating policies in p were exhaustively evaluated for policy limits $K \leq 1$, $K \leq 2$, and $K \leq 3$. For each policy limit, all candidate-policy subsets satisfying $1 \leq \lvert A \rvert \leq K$ were compared over the same 93 unique physical scenarios. The certification results are reported in Table 9.

Table 9. Four-stage certification costs and scenario assignments.

Policy Limit	Certified Policy Subset	No. of Policies Used	Certified Worst-Case Total Cost (CU)	Worst-Cost Scenario	Scenario–Policy Assignment
$K \leq 1$	(5)	1	87,700.48	(b_2, b_3, b_4)	[93]
$K \leq 2$	(2, 3)	2	87,630.32	(d_0, b_2, b_3)	[7,86]
$K \leq 3$	(2, 3, 5)	3	87,630.32	(d_0, b_2, b_3)	[7,64,22]

Increasing the policy limit from $K \leq 1$ to $K \leq 2$ reduces the certified worst-case total cost from 87,700.48 CU to 87,630.32 CU, a reduction of 70.16 CU, or approximately 0.08%. Increasing the policy limit to $K \leq 3$ produces no further reduction in the primary objective, although the selected candidate-policy subset changes from (2, 3) to (2, 3, 5) under the secondary selection rule. The third policy therefore improves delivery-performance indicators without reducing the primary worst-case cost.

The corresponding service-risk indicators are shown in Figure 4.

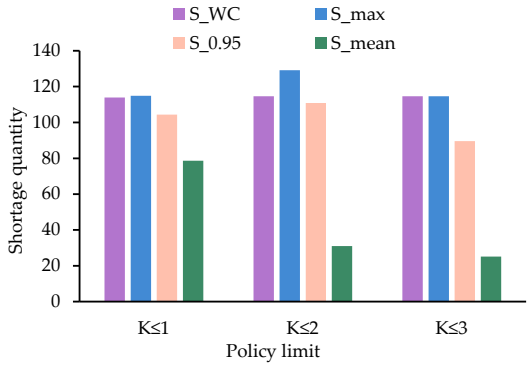


Figure 4. Shortage-related service-risk indicators under different policy limits in the four-stage benchmark.

5.4. Scenario–Policy Assignment Under the Selected Candidate-Policy Subsets

The scenario–policy assignments corresponding to the selected candidate-policy subsets in Table 9 are shown in Figure 5. Each physical scenario is assigned according to the fixed-policy evaluation and the deterministic assignment rule defined in Section 4.2.

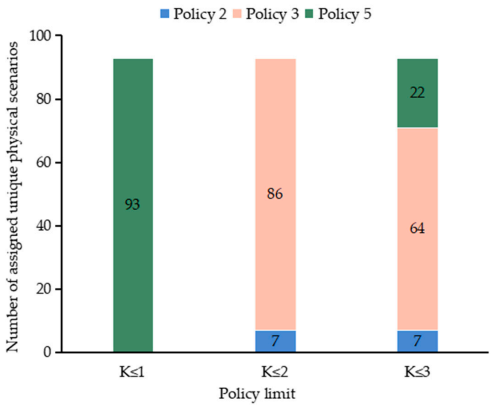


Figure 5. Scenario–policy assignment under the selected candidate-policy subsets for different policy limits.

For $K \leq 1$, Policy 5 is assigned all 93 physical scenarios. When the policy limit increases to $K \leq 2$, Policies 2 and 3 are selected, with 7 and 86 physical scenarios assigned to them, respectively. For $K \leq 3$, the selected subset becomes (2, 3, 5). The corresponding assignment changes to 7 scenarios for Policy 2, 64 for Policy 3, and 22 for Policy 5.

For $K \leq 3$, Policy 5 is assigned 22 physical scenarios, while the assignment to Policy 2 remains unchanged. The service-risk indicators in Figure 4 are lower than those for $K \leq 2$, although the certified worst-case total cost remains unchanged.

5.5. Finite-Pool Certification Results for the Eight-Stage Computational Refinement

The eight-stage computational refinement was evaluated under the fixed production–maintenance structure $n = 1$, $N = (2, 2, 3, 3, 2, 2, 1, 1)$.

After candidate-policy generation and deduplication, the finite common candidate-policy pool contained 13 unique candidate operating policies. For $K \leq 3$, the certification evaluated $\binom{13}{1} + \binom{13}{2} + \binom{13}{3} = 377$ candidate-policy subsets. The certification results are reported in Table 10.

Table 10. Eight-stage certification costs and scenario assignments.

Policy Limit	Certified Worst-Case Total Cost (CU)	Worst-Cost Scenario	No. of Active Policies	Scenario–Policy Assignment
$K \leq 1$	73,480.58	(d_0,b_2,b_3)	1	[93]
$K \leq 2$	69,470.65	(d_0,b_2,b_3)	2	[85,8]
$K \leq 3$	69,470.65	(d_0,b_2,b_3)	2	[85,8]

Increasing the policy limit from $K \leq 1$ to $K \leq 2$ reduces the certified worst-case total cost from 73,480.58 CU to 69,470.65 CU, a reduction of 4009.93 CU or 5.46%. The result for $K \leq 3$ remains 69,470.65 CU, and the selected candidate-policy subset still contains two active policies.

The corresponding shortage results are given in Table 11.

Table 11. Eight-stage shortage-risk metrics.

Policy Limit	S^{WC}	Scenario with Maximum Shortage	S^{max}
$K \leq 1$	76.34	(d_0, b_2, b_3)	76.34
$K \leq 2$	27.46	(d_0, b_3, b_4)	64.03
$K \leq 3$	27.46	(d_0, b_3, b_4)	64.03

From $K \leq 1$ to $K \leq 2$, S^{WC} decreases from 76.34 to 27.46, corresponding to a reduction of 64.03%. S^{max} decreases from 76.34 to 64.03, a reduction of 16.13%. The worst-cost scenario for $K \leq 2$ and $K \leq 3$ is (d_0, b_2, b_3), whereas the maximum shortage occurs under (d_0, b_3, b_4). The worst-cost scenario and the maximum-shortage scenario are therefore not the same.

For $K \leq 2$ and $K \leq 3$, Policies 12 and 13 are selected, with 85 and 8 physical scenarios assigned to them, respectively. Their operating settings are reported in Table 12.

Table 12. Effective candidate operating policies in the certified eight-stage pool. (**a**) Policy-level decisions and scenario assignments. (**b**) Stage-wise processing-rate and planned interstage buffer-demand surrogate profiles.

(a)

Policy	Initial Release Quantity R	No. of Assigned Scenarios
Policy 12	989.52	85
Policy 13	1087.63	8

(b)

Policy	Variable	1a	1b	2a	2b	3a	3b	4a	4b
Policy 12	μ_i	10.28	10.10	10.08	10.00	9.53	9.53	9.53	9.53
Policy 12	SS_i	9.23	0.74	4.07	23.75	0.00	0.00	0.00	0.00
Policy 13	μ_i	11.30	10.99	10.89	10.55	10.00	9.51	9.53	9.53
Policy 13	SS_i	15.32	5.41	16.91	27.54	24.28	0.00	0.00	0.00

Policy 13 has a larger initial release quantity and higher processing rates in the upstream stages than Policy 12. It is assigned to 8 of the 93 physical scenarios, while Policy 12 is assigned to the remaining 85 scenarios. Increasing the policy limit to $K \leq 3$ does not change the selected policies, the scenario–policy assignment, or the reported cost and shortage results.

5.6. Closed-Loop Refinement Stability

The eight-stage parameters were obtained through the iterative calibration procedure described in Section 4.5. To distinguish the effects of parameter updating from those of policy-pool expansion, the proposed closed-loop framework evaluates mapping accuracy, envelope approximation accuracy, propagation consistency, parameter variation, policy-pool improvement, interaction effects, and parameter-bound utilization independently. The calibration procedure terminates only when all prescribed stopping criteria are satisfied for two consecutive iterations.

The validation results for the last two calibration iterations are summarized in Table 13.

Table 13. Numerical acceptance results for two consecutive rounds of the eight-stage closed-loop refinement.

Indicator	Acceptance Criterion	Iteration 3	Iteration 4
Maximum mapping NRMSE	≤2.00%	0.04%	0.04%
Mapping $Q_{0.95}$ error	≤5.00%	0.06%	0.06%
Maximum mapping error	≤10.00%	0.09%	0.08%
Maximum envelope NRMSE	≤2.00%	0.03%	0.03%
Envelope $Q_{0.95}$ error	≤5.00%	0.06%	0.06%
Maximum envelope error	≤10.00%	0.10%	0.10%
Maximum normalized propagation slack	≤0.50%	0.05%	0.05%
Maximum parameter-update effect	≤0.50%	0.42%	0.02%
Maximum policy-pool gain	≤0.50%	<0.01%	0.48%
Maximum normalized interaction parameter-bound ratio	Diagnostic only ≤20.00%	<0.01% 12.12%	<0.01% 12.12%

The validation and stopping indicators satisfy the prescribed criteria in both Iterations 3 and 4. The maximum parameter-update effect decreases from 0.42% in Iteration 3 to 0.02% in Iteration 4. The maximum candidate-policy-pool expansion effect is 0.48% in Iteration 4, remaining below the stopping threshold of 0.50%. The normalized interaction measure remains below 0.01% and is used only as a diagnostic indicator. Since the stopping conditions are satisfied for two consecutive iterations, the closed-loop refinement terminates after Iteration 4.

To separate the effects of parameter updating and candidate-policy-pool expansion, the four-way cross-evaluation defined in Section 4.5.3 is performed. The results are reported in Table 14.

Table 14. Four-cell cross-certification results in Round 4.

K	Z_K^{00}	Z_K^{10}	Z_K^{01}	Z_K^{11}	Δ_K^{par}	Δ_K^{pool}	Δ_K^{int}
1	73,831.90	73,831.83	73,480.66	73,480.58	<0.01%	0.48%	<0.01%
2	69,484.04	69,472.20	69,482.59	69,470.65	0.02%	<0.01%	<0.01%
3	69,482.59	69,472.20	69,482.59	69,470.65	0.01%	<0.01%	<0.01%

The eight-stage closed-loop computational refinement contains five recorded iterations, indexed from 0 to 4. Under the configured paper-mode acceptance protocol, Iterations 3 and 4 satisfied the required conditions in consecutive rounds, and Iteration 4 was retained as the final numerical refinement. The validation examines several aspects of numerical consistency. Mapping and quantile-envelope errors assess agreement with the selected reference quality relationships. Propagation-slack measures quantify consistency with the prescribed quality-propagation constraints, while the four-cell cross-certification separates the effects of parameter updating and candidate-policy-pool expansion. The final physical-scenario audit confirmed 93 distinct scenarios under the shared uncertainty activation structure. The final common pool contained 13 distinct candidate operating policies, and the complete evaluation matrix contained 1209 fixed-policy–scenario records with Gurobi status OPTIMAL.

5.7. Sensitivity Analysis and Computational Performance

Under the baseline settings of the four-stage benchmark and eight-stage refinement, increasing the policy limit from one to two reduces the finite-pool worst-case total cost, whereas a third policy does not provide further primary-cost improvement. In the four-stage benchmark, the third policy nevertheless improves shortage-related indicators. The

following analyses examine the sensitivity of these findings to the uncertainty budget, quality-related parameters, and cost coefficients, together with the recorded computational performance.

The eight-stage results also show that the physical scenario governing the robust cost objective need not produce the largest delivery shortage. The worst-cost scenario and the maximum-shortage scenario are different. The certified worst-case total cost and the shortage indicators therefore characterize different aspects of system performance.

5.7.1. Sensitivity to the Uncertainty Budget and Policy Limit

The uncertainty budget was varied over $\Gamma = 1, 2, 3, 4$, and the policy limit was extended to $K = 1, 2, 3, 4$. Throughout these experiments, the original four-stage production–maintenance structure, $n = 1$ and $N = (2, 3, 2, 1)$, and the five-policy common candidate pool were retained. The four uncertainty budgets correspond to 9, 37, 93, and 163 unique physical scenarios. For $\Gamma \leq 3$, the comparisons use the verified baseline evaluation matrix restricted to the corresponding nested scenario sets. For $\Gamma = 4$, the original 93 scenarios were retained, and the five candidate policies were evaluated under 70 additional physical scenarios, requiring 350 new fixed-policy evaluations. Table 15 reports the minimum worst-case total cost obtained through exhaustive subset selection within the fixed candidate-policy pool.

Table 15. Sensitivity of finite-pool worst-case total cost to uncertainty budget and policy limit (CU).

Γ	Unique Scenarios	$K = 1$	$K = 2$	$K = 3$	$K = 4$
1	9	72,135.61	69,448.09	69,448.09	69,448.09
2	37	80,277.59	78,139.95	78,076.96	78,076.96
3	93	87,700.48	87,630.32	87,630.32	87,630.32
4	163	96,304.79	96,225.65	96,225.65	96,225.65

The cost reductions associated with increasing the policy limit from one to two are approximately 3.73%, 2.66%, 0.08%, and 0.08% for the four uncertainty budgets, respectively. The third policy provides an additional primary-cost reduction at $\Gamma = 2$, from 78,139.95 to 78,076.96 CU. In contrast, the primary objective remains unchanged when the policy limit increases from two to three at $\Gamma = 3$ and $\Gamma = 4$. Allowing a fourth policy does not further reduce the primary objective in any of the tested budget configurations. These results show that the value of additional candidate policies depends on the uncertainty support. The absence of further improvement at $K = 3$ in the baseline experiment is therefore not a general property of K-adaptability. The conclusions remain conditional on the original five-policy common candidate pool, since policies were not regenerated for each uncertainty budget.

5.7.2. Sensitivity to Quality-Propagation Parameters

Three one-factor sensitivity experiments were conducted for the nonlinear propagation coefficients $\boldsymbol{\beta_i}$, stage-load effect coefficients $\boldsymbol{\theta_i}$, and nonconforming-product sensitivity coefficients $\boldsymbol{\rho_i}$. Each coefficient family was scaled by 0.8, 1.0, and 1.2 relative to its baseline values, with the remaining parameters unchanged. For $\boldsymbol{\beta_i}$ and $\boldsymbol{\theta_i}$, the nominal and perturbation components were scaled consistently. The experiments retained the four-stage production–maintenance structure, $\Gamma = 3$, the 93 unique physical scenarios, and the original five-policy common candidate pool. The candidate policies were re-evaluated under the modified parameter settings, followed by exhaustive subset selection within the finite pool. The results of these quality-parameter sensitivity experiments are summarized in Table 16.

Table 16. Quality-parameter sensitivity of finite-pool worst-case total cost (CU).

Parameter	Scale	$K = 1$	$K = 2$	$K = 3$	$K = 1 \rightarrow 2$
β_i	0.8	73,879.47	73,485.42	73,485.42	0.53%
	1.0	87,700.48	87,630.32	87,630.32	0.08%
	1.2	104,214.25	103,924.65	103,924.65	0.28%
θ_i	0.8	86,250.29	85,208.11	85,208.11	1.21%
	1.0	87,700.48	87,630.32	87,630.32	0.08%
	1.2	90,233.39	90,163.79	90,163.79	0.08%
ρ_i	0.8	69,389.45	69,389.45	69,389.45	0.00%
	1.0	87,700.48	87,630.32	87,630.32	0.08%
	1.2	106,545.28	106,545.28	106,545.28	0.00%

The results show that the primary-cost benefit of a second policy depends on the parameter settings. Scaling $\beta_i$by 0.8 and 1.2 yields relative improvements of approximately 0.53% and 0.28%, compared with 0.08% at the baseline. For θ_i, the improvement is 1.21% at the 0.8 scale and 0.08% at the 1.2 scale. For the two nonbaseline $\rho_i$configurations, increasing the policy limit from one to two does not reduce the primary objective within the evaluated pool.

5.7.3. Sensitivity to Cost Coefficients

Three one-factor sensitivity experiments were conducted by independently scaling the shortage penalty, quality-loss coefficients, and maintenance-risk surrogate coefficients by 0.8, 1.0, and 1.2. The four-stage production–maintenance structure, baseline uncertainty budget, 93 physical scenarios, and five-policy candidate pool were retained, while fixed-policy evaluations and finite-pool subset selection were repeated under each cost configuration.

Table 17 shows that increasing the policy limit from one to two reduces the finite-pool worst-case total cost under all nine tested configurations. The relative improvement ranges from approximately 0.08% to 0.11%, corresponding to reductions of 64.91–93.82 CU. Under the original benchmark parameters, the reduction is 70.16 CU, or approximately 0.08%. The improvement depends on the assumed cost coefficients. For example, scaling the shortage penalty by 0.8 yields a relative improvement of approximately 0.11%, whereas scaling the maintenance-risk surrogate coefficients by 1.2 yields an improvement of approximately 0.08%. Allowing a third policy does not further reduce the primary objective under these tested cost configurations. However, the shortage-based secondary selection produces a reduction of approximately 18.50–19.00% in the equally weighted mean shortage compared with the selected two-policy subset. These findings indicate that the additional policy choices provide a small primary-cost benefit under the examined cost parameters, while secondary shortage indicators may provide further operational information. The corresponding cost-sensitivity results are reported in Table 17.

Table 17. Sensitivity of finite-pool worst-case total cost to cost coefficients (CU).

Cost Parameter	Scale	$K = 1$	$K = 2$	$K = 1 \rightarrow 2$
Shortage penalty	0.8	82,676.29	82,584.04	0.11%
	1.0	87,700.48	87,630.32	0.08%
	1.2	92,755.82	92,676.74	0.09%
Quality-loss coefficients	0.8	78,353.80	78,288.89	0.08%
	1.0	87,700.48	87,630.32	0.08%
	1.2	97,065.71	96,971.89	0.10%
Maintenance-risk surrogate coefficients	0.8	87,625.00	87,554.75	0.08%
	1.0	87,700.48	87,630.32	0.08%
	1.2	87,776.07	87,705.96	0.08%

5.7.4. Computational Performance, Stage Refinement, and Certification Scope

In the four-stage benchmark, the finite common candidate-policy pool contains five operating policies evaluated under 93 unique physical scenarios, producing 465 policy–scenario evaluations. Exhaustive subset selection examines 25 nonempty subsets of cardinality at most three. One archived baseline MPA-C&CG run required 148.74 s. This measurement represents an individual algorithm run rather than the total time required to generate and certify the five-policy common pool. In the supplementary $\Gamma = 4$ experiment, 350 additional fixed-policy evaluations required 14.28 s, followed by 0.03 s for finite-pool subset selection.

For the six-stage configuration, the refinement structure (1, 2, 2, 1) was constructed using previously calibrated substage propagation parameters. After reference-mapping verification, candidate policies were generated and evaluated under all 93 physical scenarios. The retained worst-case costs were 82,451.12 CU for $\boldsymbol{K} = 1$ and 82,201.11 CU for $\boldsymbol{K} = 2$. Introducing a third candidate did not further reduce the primary objective. The three candidate positions were assigned 64, 29, and 0 scenarios, respectively. The recorded incremental workflow accumulated approximately 160.95 s, including repeated evaluations during $\boldsymbol{K} = 1$ refinement and cached evaluation reuse during $\boldsymbol{K} = 3$ augmentation.

The eight-stage closed-loop computational refinement produced a final common pool of 13 distinct candidate policies and 93 unique physical scenarios. The complete evaluation matrix contains 1209 policy–scenario records, all with Gurobi status OPTIMAL. Exhaustive subset selection examines 13 single-policy, 78 two-policy, and 286 three-policy subsets, totaling 377 combinations. The certified worst-case costs within this pool were 73,480.58 CU for $\boldsymbol{K} = 1$ and 69,470.65 CU for both $\boldsymbol{K} = 2$ and $\boldsymbol{K} = 3$. In the selected eight-stage closed-loop iteration, the recorded MPA-C&CG algorithm-call times for $\boldsymbol{K} = (1, 2, 3)$ were 70.50, 87.37, and 113.92 s. Their sum, 271.79 s, excludes external parameter calibration, screening, final common-pool evaluation, and subsequent subset certification.

5.8. Discussion of Numerical Findings, Practical Implications, and Limitations

The numerical experiments provide several observations about finite adaptability in integrated production–quality–maintenance optimization under uncertain quality-deviation propagation.

First, additional candidate operating policies can reduce the worst-case total cost, but their marginal value depends on the physical scenario support and the available candidate-policy pool. In the four-stage benchmark at $\Gamma = 3$, the primary-cost improvement occurs when the policy limit increases from one to two. A third policy does not further reduce the worst-case cost, although it improves the shortage-related secondary indicators. The uncertainty-budget sensitivity analysis shows that this plateau is not general: at $\Gamma = 2$, a third policy provides an additional reduction within the same five-policy pool.

Second, the quality-parameter sensitivity results show that the magnitude of the adaptability benefit depends on the nonlinear propagation coefficients, stage-load effects, and nonconforming-product sensitivity coefficients. These parameters influence the scenario-specific manufacturing response and the relative usefulness of the available operating policies. The results therefore support the need for careful parameter selection and independent calibration before industrial application.

Third, the cost-sensitivity results indicate that the four-stage two-policy improvement remains positive but relatively small across the tested cost configurations. The approximately 0.08% reduction in the baseline case is a numerical benchmark result rather than evidence of substantial industrial monetary savings. Shortage-related indicators may improve even when the primary objective remains unchanged, but the value of maintain-

ing additional operating policies depends on implementation costs that are not explicitly included in the model.

Fourth, the four-stage, six-stage, and eight-stage configurations demonstrate computational applicability at the examined refinement scales. However, they differ in numerical representation, parameter calibration, candidate-policy construction, and execution protocol. Their cost and runtime differences therefore cannot be attributed solely to manufacturing-stage count. The available computational records characterize the work performed in the tested cases but do not establish a controlled scaling law.

Several modeling limitations also remain. The present numerical study considers serial, single-product manufacturing with deterministic demand and fixed production–maintenance structures. Nonconforming products are removed rather than reprocessed, and the formulation does not explicitly represent rework, recirculation, parallel routing, or assembly operations. Maintenance-related effects are represented through the specified capacity relationships and maintenance-risk surrogate rather than a field-validated equipment-failure model. The numerical parameters and uncertainty magnitudes have not been calibrated using independent data from a specific industrial production line. In addition, the finite-pool certification concerns the generated candidate operating policies and prescribed physical scenarios, subject to the numerical approximations and solver tolerances used in their evaluation. It does not establish global optimality for the original continuous nonlinear K-adaptability formulation. For industrial implementation, cross-stage quality-propagation coefficients would need to be estimated from production and inspection records. Processing-rate, stage-load, nonconforming-product, and maintenance-related parameters would require corresponding operational data. Enterprise-specific monetary costs and the expense of preparing multiple contingent policies would also need to be considered. The calibrated model should then be assessed using independent or out-of-sample manufacturing observations.

Future research may extend the computational experiments to jointly optimize production–maintenance structural decisions, while also considering uncertain demand, rework, recirculation, and more complex routing or assembly structures. Such extensions would require changes to the uncertainty representation, material-flow constraints, quality-state transitions, and operating-policy decisions. Data-driven parameter calibration and improved candidate-policy generation methods also remain important directions for further investigation.

6. Conclusions

This study developed a two-stage K-adaptability robust optimization framework for integrated production–quality–maintenance decisions in a multistage manufacturing system with uncertain quality-deviation propagation. The model links stage capacity, preventive maintenance, linear and nonlinear quality-deviation propagation, processing-rate deviation, stage-load effects, nonconforming-product formation, and delivery shortage. Limited post-realization adaptability is represented through a finite set of candidate operating policies prepared before uncertainty realization. The computational procedure combines active-scenario refinement, candidate-policy augmentation, complete physical-scenario evaluation, and exhaustive subset selection within a finite common candidate-policy pool.

The numerical results show that the value of finite adaptability depends on the uncertainty support, quality-related parameters, cost coefficients, and available candidate policies. In the four-stage benchmark, increasing the policy limit from one to two reduces the finite-pool worst-case total cost by approximately 0.08%. A third policy does not further reduce the primary objective but improves maximum shortage, empirical 95th-percentile shortage, and equally weighted mean shortage by approximately 11.21%, 19.23%, and

18.97%, respectively. The eight-stage refinement yields a larger two-policy cost reduction of approximately 5.46%, while the additional six-stage verification shows a reduction from 82,451.12 to 82,201.11 CU when a second policy is introduced. The supplementary sensitivity experiments demonstrate that these numerical benefits depend on the tested settings and that a third policy can provide further primary-cost improvement under a different uncertainty budget.

The study is limited to numerical manufacturing configurations with the prescribed material-flow structure and benchmark parameters. The economic costs of implementing multiple contingent policies are not explicitly modeled, and no independent industrial calibration or field validation has been performed. Furthermore, finite-pool certification does not establish global optimality over the original continuous nonlinear K-adaptability formulation, and the available runtime records do not provide a controlled stage-count scaling law. Future research should investigate industrial parameter estimation and out-of-sample validation, more complex material-flow structures, uncertain demand, endogenous maintenance-structure decisions, and computational methods for larger candidate-policy spaces.

Author Contributions: Conceptualization, H.J. and J.L.; methodology, H.J.; software, H.J.; validation, H.J.; formal analysis, H.J.; investigation, H.J.; data curation, H.J.; writing—original draft preparation, H.J.; writing—review and editing, H.J. and J.L.; visualization, H.J.; supervision, J.L.; project administration, J.L.; funding acquisition, J.L. All authors have read and agreed to the published version of the manuscript.

Funding: This research was supported by the Lanzhou Science and Technology Project (Grant No. 2025-2-107), the National Key Research and Development Plan (Grant No. 2018YFB1703105), the Key Research and Development Plan of Gansu Province (Grant No. 22YF11GA315), and the Hongliu First-Class Disciplines Development Program of Lanzhou University of Technology.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Ismail, M.; Mostafa, N.A.; El-Assal, A. Quality monitoring in multistage manufacturing systems by using machine learning techniques. *J. Intell. Manuf.* **2022**, *33*, 2471–2486. [CrossRef]
2. Zhang, D.; Liu, Z.; Jia, W.; Liu, H.; Tan, J. Path enhanced bidirectional graph attention network for quality prediction in multistage manufacturing process. *IEEE Trans. Ind. Inform.* **2022**, *18*, 1018–1027. [CrossRef]
3. Moliner-Heredia, R.; Peñarrocha-Alós, I.; Abellán-Nebot, J.V. A methodology for data-driven adjustment of variation propagation models in multistage manufacturing processes. *J. Manuf. Syst.* **2023**, *67*, 281–295. [CrossRef]
4. Wang, P.; Qu, H.; Zhang, Q.; Xu, X.; Yang, S. Production quality prediction of multistage manufacturing systems using multi-task joint deep learning. *J. Manuf. Syst.* **2023**, *70*, 48–68. [CrossRef]
5. Cho, H.; Kim, K.; Yoon, K.; Chun, J.; Kim, J.; Lee, K.; Lee, J.; Lim, C. MMP Net: A feedforward neural network model with sequential inputs for representing continuous multistage manufacturing processes without intermediate outputs. *IISE Trans.* **2024**, *56*, 1058–1069. [CrossRef]
6. Magnanini, M.C.; Demir, O.; Colledani, M.; Tolio, T. Performance evaluation of multi-stage manufacturing systems operating under feedback and feedforward quality control loops. *CIRP Ann.* **2024**, *73*, 349–352. [CrossRef]
7. Shi, H.L.; Jiang, P.Y. Quality control in multistage machining processes based on a machining error propagation event-knowledge graph. *Adv. Manuf.* **2024**, *12*, 679–697. [CrossRef]
8. Sun, X.; Beghi, A.; Susto, G.A.; Lv, Z. Deep learning-based quality prediction for multi-stage sequential hot rolling processes in heavy rail manufacturing. *Comput. Ind. Eng.* **2024**, *196*, 110466. [CrossRef]
9. Wang, P.; Tao, H.; Qi, J.; Li, P. Machining quality prediction of complex thin-walled parts using multi-task dual domain adaptive deep transfer learning. *Adv. Eng. Inform.* **2024**, *62*, 102640. [CrossRef]
10. Yin, X.; Chen, Y.; Mu, Z.; Cui, Q. A quality propagation path and implementation resolution-aware bidirectional operational parameter optimization framework for multistage manufacturing processes. *Comput. Ind. Eng.* **2025**, *209*, 111445. [CrossRef]

11. Ait-El-Cadi, A.; Gharbi, A.; Dhouib, K.; Artiba, A. Integrated production, maintenance and quality control policy for unreliable manufacturing systems under dynamic inspection. *Int. J. Prod. Econ.* **2021**, *236*, 108140. [CrossRef]
12. Hadian, S.M.; Farughi, H.; Rasay, H. Joint planning of maintenance, buffer stock and quality control for unreliable, imperfect manufacturing systems. *Comput. Ind. Eng.* **2021**, *157*, 107304. [CrossRef]
13. Shi, L.; Lv, X.; He, Y.; He, Z. Optimising production, maintenance, and quality control for imperfect manufacturing systems considering timely replenishment. *Int. J. Prod. Res.* **2024**, *62*, 3504–3525. [CrossRef]
14. Lv, X.; Shi, L.; He, Y.; He, Z.; Lin, D.K.J. Joint optimization of production, maintenance, and quality control considering the product quality variance of a degraded system. *Front. Eng. Manag.* **2024**, *11*, 413–429. [CrossRef]
15. Paraschos, P.D.; Koulouriotis, D.E. Learning-based production, maintenance, and quality optimization in smart manufacturing systems: A literature review and trends. *Comput. Ind. Eng.* **2024**, *198*, 110656. [CrossRef]
16. Koulinas, G.K.; Paraschos, P.D.; Koulouriotis, D.E. A machine learning framework for explainable knowledge mining and production, maintenance, and quality control optimization in flexible circular manufacturing systems. *Flex. Serv. Manuf. J.* **2024**, *36*, 737–759. [CrossRef]
17. Wei, X.; Wang, Y.; He, Y.; Liu, Z.; He, Z. Integrated production, maintenance and quality control for complex manufacturing systems considering imperfect maintenance and dynamic inspection. *Reliab. Eng. Syst. Saf.* **2025**, *259*, 110896. [CrossRef]
18. Zhang, J.; Chen, X.; An, Y.; Zhang, L.; Shi, H.; Xu, W. Joint optimization of preventive maintenance and product quality improvement policies for deteriorating manufacturing systems with quality-reliability dependency. *Reliab. Eng. Syst. Saf.* **2025**, *255*, 110661. [CrossRef]
19. Cao, Y.; Wang, P.; Xv, W.; Dong, W. Joint optimization of quality-based multi-level maintenance and buffer stock within multi-specification and small-batch production. *Front. Eng. Manag.* **2025**, *12*, 754–773. [CrossRef]
20. Rahimian, H.; Mehrotra, S. Frameworks and results in distributionally robust optimization. *Open J. Math. Optim.* **2022**, *3*, 85. [CrossRef]
21. Tönissen, D.D.; Arts, J.J.; Shen, Z.J.M. A column-and-constraint generation algorithm for two-stage stochastic programming problems. *TOP* **2021**, *29*, 781–798. [CrossRef]
22. Cohen, I.; Postek, K.; Shtern, S. An adaptive robust optimization model for parallel machine scheduling. *Eur. J. Oper. Res.* **2023**, *306*, 83–104. [CrossRef]
23. Goerigk, M.; Khosravi, M. Optimal scenario reduction for one- and two-stage robust optimization with discrete uncertainty in the objective. *Eur. J. Oper. Res.* **2023**, *310*, 529–551. [CrossRef]
24. Detienne, B.; Lefebvre, H.; Malaguti, E.; Monaci, M. Adjustable robust optimization with objective uncertainty. *Eur. J. Oper. Res.* **2024**, *312*, 373–384. [CrossRef]
25. Metzker Soares, P.; Thevenin, S.; Adulyasak, Y.; Dolgui, A. Adaptive robust optimization for lot-sizing under yield uncertainty. *Eur. J. Oper. Res.* **2024**, *313*, 513–526. [CrossRef]
26. Hanasusanto, G.A.; Kuhn, D.; Wiesemann, W. K-adaptability in two-stage robust binary programming. *Oper. Res.* **2015**, *63*, 877–891. [CrossRef]
27. Bertsimas, D.; Kim, C.W. A machine learning approach to two-stage adaptive robust optimization. *Eur. J. Oper. Res.* **2024**, *319*, 16–30. [CrossRef]
28. Julien, E.; Postek, K.; Birbil, Ş.İ. Machine learning for *K*-adaptability in two-stage robust optimization. *Inf. J. Comput.* **2025**, *37*, 644–665. [CrossRef]
29. Agra, A. Two-stage distributionally robust optimization with a finite support. *Comput. Oper. Res.* **2025**, *183*, 107142. [CrossRef]
30. Hu, L.; Zhang, Y.; Wen, M.; Leus, R.; Zhang, N. Robust parallel machine selection and scheduling with uncertain release times. *Eur. J. Oper. Res.* **2025**, *327*, 838–856. [CrossRef]
31. Bertsimas, D.; Sim, M. The price of robustness. *Oper. Res.* **2004**, *52*, 35–53. [CrossRef]
32. Zeng, B.; Zhao, L. Solving two-stage robust optimization problems using a column-and-constraint generation method. *Oper. Res. Lett.* **2013**, *41*, 457–461. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.