From Composition to Cost: Modelling Material Variability in Grey Cast Iron Machining

Chenghan Qian

2026

LUND UNIVERSITY

LTH
FACULTY OF ENGINEERING

MASTER THESIS
DIVISION OF PRODUCTION AND MATERIALS ENGINEERING
LUND UNIVERSITY

Supervisor: Christina Windmark, Senior Lecturer

Examiner: Mats Andersson, Associate Professor

Author: Chenghan Qian
Lund, Sweden 2026

Avdelningen för Industriell Produktion
Lunds Tekniska Högskola
Lunds universitet
Box 118
221 00 Lund
Sverige

Division of Production and Materials Engineering
LTH, Faculty of Engineering
Lund University
Box 118
SE-221 00 Lund
Sweden

Printed in Sweden
Media-Tryck
Lund University

Foreword

This thesis is the result of my master's project at the Division of Production and Materials Engineering, LTH, Faculty of Engineering, Lund University. The work was carried out during the spring of 2026 and concludes my master's degree in Production and Materials Engineering.

The motivation for this study came from a question that often appears at the intersection of metallurgy and production engineering: when material composition varies within an accepted specification window, what happens further down the line — at the cutting tool, at the bottleneck operation, and ultimately on the unit cost of the finished part? Grey cast iron is a material in which this question is particularly relevant, because its production increasingly relies on recycled feedstock with varying residual-element content. Connecting the metallurgical reality of incoming castings with the economic reality of the machining shop became the central interest of this thesis.

I would like to thank my supervisor, Christina Windmark, for her guidance, patience and clear feedback throughout the project, and for encouraging me to keep questioning my own assumptions when the simulation results turned out differently from what I initially expected. I would also like to thank my examiner, Mats Andersson, for valuable input during the planning stage of the work.

I am grateful to the industrial practitioner who took the time to share their experience during the consultation described in Section 3.2; the conversation shaped how I came to interpret the difference between average production loss and event-based disturbance in real machining operations.

Finally, I would like to thank my family and friends for their support during the long evenings of coding, writing and rewriting that produced this document.

Lund 2026-05-25

Chenghan Qian

Sammanfattning

Materialvariation i gråjärn diskuteras vanligtvis utifrån metallurgi eller bearbetbarhet, medan dess konsekvenser för produktionsflödet och styckkostnaden är mindre systematiskt karaktäriserade. Detta examensarbete undersöker hur sammansättningsvariation i gråjärn propagerar genom bearbetningsrespons till produktionssystemets beteende, och kvantifierar de resulterande ekonomiska konsekvenserna per detalj som underlag för försörjningskedjebeslut.

Ett förenklat modelleringsramverk i fyra lager utvecklas. Det första lagret översätter representativa sammansättningsscenarier (referensfall S0, CE/matrisförskjutning S1, lokal heterogenitet S2, restelementshärdning S3 och kombinerat ogynnsamt fall S4) till materialresponsdeskriptorer. Det andra lagret översätter dessa deskriptorer till operationsspecifika bearbetningsmultiplikatorer för bearbetningstid, verktygsslitage och störningsintensitet. Det tredje lagret propagerar dessa multiplikatorer genom en diskret händelsesimulering implementerad i SimPy, som representerar en seriell bearbetningslinje bestående av OP10 planfräsning, OP20 borrning och OP30 tung svarvning, sammankopplade med ändliga buffertar. Det fjärde lagret aggregerar produktionsflödets utdata till en styckkostnadsindikator genom ett transparent kostnadslager som omfattar maskintidskostnad, verktygsbyteskostnad, råmaterialkostnad och probabilistisk kassationskostnad.

Tre simuleringskonfigurationer analyseras. Fall A karaktäriserar produktionsflödets respons vid nominella skärparametrar. Fall B introducerar skärhastigheten vid OP30 som beslutsvariabel genom ett normaliserat hastighetsförhållande kopplat till verktygslivslängd via Taylors ekvation, och identifierar den produktionsoptimala skärhastigheten för varje scenario. Fall C utvidgar skärhastighetssvepningen med kostnadslagret och identifierar den kostnadsoptimala skärhastigheten samt den oundvikliga kostnadsökningen för varje scenario, definierad som den styckkostnadsökning som kvarstår efter scenariospecifik omoptimering av skärhastigheten.

Det huvudsakliga resultatet är att materialvariation ger upphov till en scenariospecifik oundviklig styckkostnadsökning vars storlek skiljer sig markant mellan mekanismerna. Vid en grundkostnad på 80,76 € per detalj ger restelementshärdning en oundviklig ökning på +10,46 % (8,45 € per detalj), kombinerat ogynnsamt fall +11,98 % (9,68 € per detalj), CE/matrisförskjutning +5,27 % (4,26 € per detalj) och lokal heterogenitet

+1,28 % (1,03 € per detalj). Det ungefär åttafaldiga förhållandet mellan den mest och minst allvarliga enskilda mekanismen är robust mot stor variation i kostnadsparameterantaganden, vilket bekräftas av en känslighetsanalys med 27 parameterkombinationer, där restelementshärdningens kostnadsökning genomgående hamnar i intervallet +9,05 % till +12,42 %. Styckkostnadsuppdelningen visar dessutom att maskintidskostnad och råmaterialkostnad tillsammans utgör mer än 90 % av styckkostnaden, vilket indikerar att materialvariation når den ekonomiska målfunktionen främst genom maskintidsförstärkning vid flaskhalsoperationen snarare än genom direkt verktygsförbrukning eller kvalitetsförlust. Dessa värden per detalj fastställer kvantitativa tak för den ekonomiska motiveringen av kvalitetsförbättringar för ingående material som riktar sig mot varje materialmekanism, och tillhandahåller ett strukturerat beslutsramverk för försörjningskedjan som kopplar materialmekanismer till rekommenderade process- och försörjningskedjeåtgärder.

Nyckelord: gråjärn; materialvariation; bearbetning; verktygslivslängd; Taylors skärhastighetsoptimering; diskret händelsesimulering; styckkostnadsanalys; beslutsramverk för försörjningskedjan; produktionsflödesmodellering.

Abstract

Material variability in grey cast iron is usually discussed in terms of metallurgy or machinability, while its consequences for production-flow behaviour and per-part economics are less systematically characterised. This thesis investigates how compositional variability in grey cast iron propagates through machining response into production-system behaviour, and quantifies the resulting per-part economic consequences as a basis for supply-chain decision-making.

A reduced-order modelling framework with four layers is developed. The first layer translates representative compositional scenarios (baseline S0, CE/matrix drift S1, local heterogeneity S2, residual hardening S3 and combined adverse S4) into material-response descriptors. The second layer translates these descriptors into operation-specific machining multipliers for processing time, tool-life consumption and incident intensity. The third layer propagates these multipliers through a SimPy-based discrete-event production-flow model representing a serial machining line of OP10 face milling, OP20 drilling and OP30 heavy turning, connected by finite buffers. The fourth layer aggregates the production-flow outputs into a per-part unit-cost indicator through a transparent cost layer covering machine-hour cost, tool-change cost, raw-material cost and probabilistic scrap cost.

Three simulation configurations are studied. Case A characterises the production-flow response under nominal cutting parameters. Case B introduces the OP30 cutting speed as a decision variable through a normalised speed ratio coupled to tool life via Taylor's equation, and identifies the production-optimal cutting speed for each scenario. Case C extends the cutting-speed sweep with the cost layer and identifies the cost-optimal cutting speed and the unavoidable cost penalty for each scenario, defined as the per-part cost increase that persists after per-scenario re-optimization of the cutting speed.

The principal finding is that material variability produces a scenario-dependent unavoidable unit-cost penalty whose magnitude differs substantially across mechanisms. At a baseline unit cost of €80.76 per part, residual hardening produces an unavoidable penalty of +10.46% (€8.45 per part), combined adverse conditions +11.98% (€9.68 per part), CE/matrix drift +5.27% (€4.26 per part) and local heterogeneity +1.28% (€1.03 per part). The approximately eight-fold ratio between the most and least severe single mechanisms is robust to wide variation in cost-parameter assumptions, as confirmed by a 27-combination sensitivity sweep that places the residual-hardening penalty consistently in the range +9.05% to +12.42%. The unit-cost decomposition further shows that machine-hour cost and raw-material cost together account for more than 90% of unit cost, indicating that material variability reaches the cost objective primarily through machine-time amplification at the bottleneck operation rather than through direct tool consumption or quality loss. These per-part values establish quantitative ceilings for the economic justification of incoming-material quality improvements targeting each material mechanism, and provide a structured supply-chain decision framework that links material mechanisms to recommended process-side and supply-chain responses.

Keywords: grey cast iron; material variability; machining; tool life; Taylor cutting-speed optimization; discrete-event simulation; unit-cost analysis; supply-chain decision framework; production-flow modelling

Table of Content

1. Introduction

1.1.Background

In contemporary grey cast iron production, the increasing reliance on steel scrap as a primary charge material has fundamentally altered the nature of material variability entering the production system. While traditional cast iron production was dominated by pig iron with relatively stable composition, modern industrial practice increasingly adopts synthetic cast iron routes, in which a significant proportion of the charge consists of steel scrap and recycled returns. This transition is driven by both economic and environmental considerations, but it also introduces a more complex and less predictable compositional structure at the material input stage[1], [2].

Unlike pig iron, which is produced under relatively controlled conditions, steel scrap originates from a wide range of sources, including structural steels, coated materials, alloyed components, and industrial waste streams. As a result, scrap-based charge materials inherently carry residual and trace elements that are not always explicitly controlled in the melting process. Elements such as phosphorus and sulphur are often introduced through pig iron or low-quality scrap, while manganese is frequently present in combination with sulphur in the form of MnS inclusions. In addition, residual elements such as copper, tin, chromium, and nickel may accumulate through repeated recycling cycles, forming a background composition that is difficult to eliminate and only partially adjustable through charge design [3], [4], [5].

Importantly, these elements do not enter the melt as isolated variables, but as part of correlated compositional patterns associated with specific scrap streams or charge practices. For example, increasing the proportion of steel scrap may reduce sulphur and phosphorus levels when high-quality scrap is used, but it may simultaneously introduce residual alloying elements or reduce nucleation potential, thereby requiring additional metallurgical treatment such as inoculation [2], [6]. Consequently, the resulting material variability is structured rather than random, reflecting the combined influence of charge composition, recycling history, and melting practice.

From a metallurgical perspective, the influence of such element's manifests through several recurring mechanisms. Carbon and silicon primarily govern graphite morphology and the ferrite–pearlite balance, thereby affecting overall machinability. Manganese and sulphur contribute to the formation of inclusions, often in the form of MnS, which can introduce localized heterogeneity. Residual elements such as copper, tin, chromium, and nickel tend to stabilize pearlitic or hard phases, leading to cumulative hardening effects over time [1], [4], [7]. These mechanisms are well documented at the microstructural level, but their implications for machining behaviour and production flow are less systematically understood.

As a result, material-related uncertainty in modern grey cast iron production should not be interpreted as isolated deviations in individual elements, but as the system-level consequence of evolving charge-material structures. Understanding how these compositional patterns propagate through microstructure, machining behaviour, and ultimately production flow is therefore essential for both process design and operational control.

1.2. Material Input Structure and Metallurgical Mechanisms

While the increasing use of steel scrap has introduced a higher degree of compositional variability in grey cast iron production, the resulting material uncertainty cannot be understood solely in terms of individual elemental deviations. Instead, it reflects a structured input condition arising from the combined effects of charge composition, recycling history, and metallurgical processing [1], [2].

In industrial practice, grey cast iron is typically produced from a mixture of pig iron, steel scrap, internal returns, and ferroalloy additions. Each of these input streams carries its own compositional signature. Pig iron generally provides a relatively stable base in terms of carbon and silicon content, whereas steel scrap introduces a wider spectrum of residual and trace elements depending on its origin. Internal returns tend to reinforce existing compositional patterns through repeated recycling cycles, leading to cumulative effects over time [2], [3].

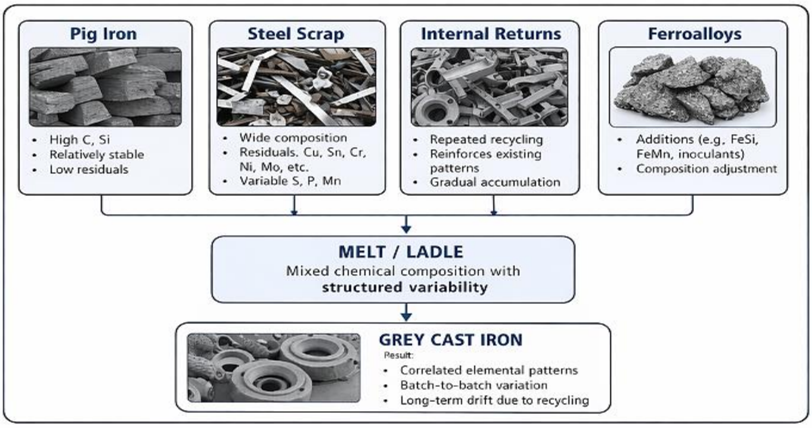

Figure 1. Flow of Material inputs and their characteristics, leading to structured compositional variability in the melt and final cast iron

Within this context, several groups of elements can be distinguished based on their dominant metallurgical roles.

Carbon and silicon form the fundamental basis of grey cast iron, governing graphite formation and the balance between ferrite and pearlite. Their influence is typically continuous and manifests at the scale of the entire batch. Variations in these elements, even within specification limits, can shift the overall microstructure and thereby affect machinability in a systematic manner [1], [8].

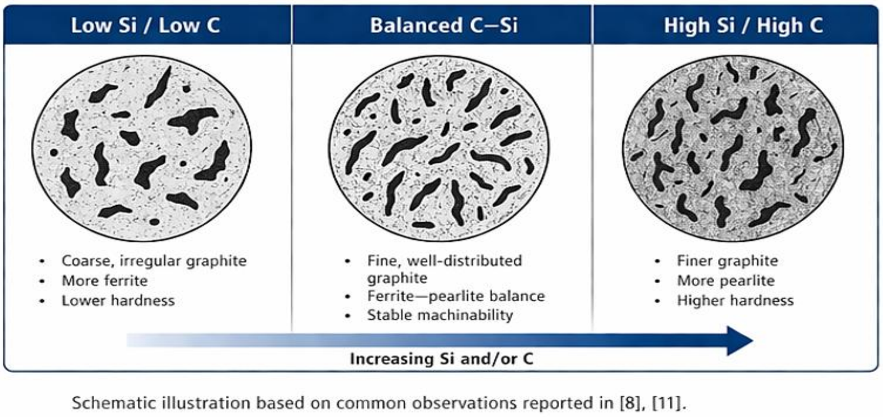

Figure 2.Effect of C and Si on graphite morphology and ferrite-pearlite balance

In contrast, elements such as manganese, sulphur, and phosphorus contribute to the formation of inclusions or secondary phases that are often unevenly distributed. Manganese and sulphur commonly combine to form MnS inclusions, which may act as nucleation sites but can also introduce local heterogeneity when present in non-uniform distributions. Phosphorus tends to form steadite, a hard and brittle eutectic constituent segregated along cell boundaries. The influence of these elements is therefore localized and not fully captured by average composition values [4], [7].

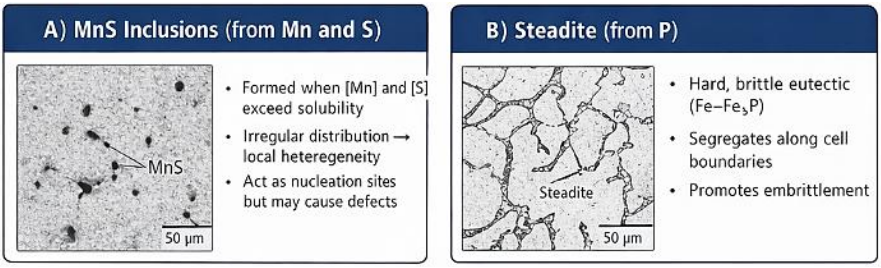

Figure 3.Localized mechanisms:(A) MnS inclusions from Mn and S; (B) Steadite formation from P. Adapted from the original sources [4],[7]

A third group consists of residual elements such as copper, tin, chromium, and nickel, which tend to accumulate through repeated recycling of scrap. These elements stabilize pearlitic or carbide-forming phases and exert a cumulative effect on the microstructure. Their influence is gradual but persistent, leading to increased hardness and a narrowing of the process window. Because these elements are difficult to remove once introduced, their concentrations reflect long-term trends in raw material sourcing rather than short-term batch variation [3], [5].

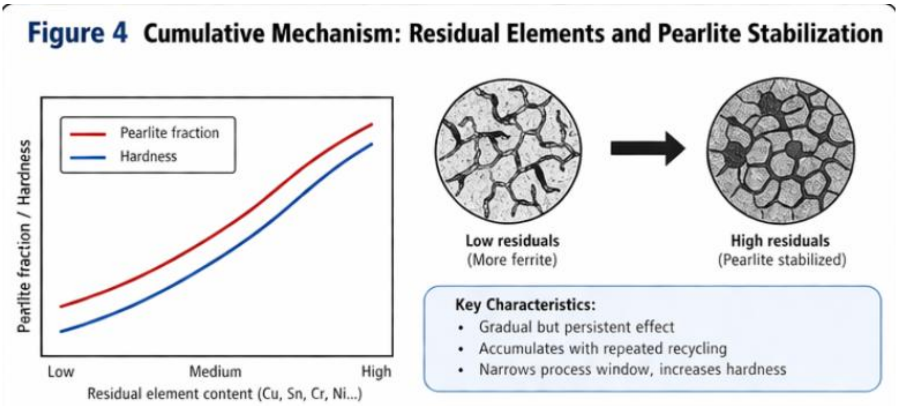

Figure 4.Cumulative effect of residual elements (Cu, Sn, Cr, Ni, etc.) on pearlite stabilization and hardness.

An important characteristic of these metallurgical mechanisms is that they rarely occur in isolation. Instead, elements co-vary as part of recurring compositional patterns associated with specific charge practices or scrap streams. For instance, increasing the proportion of steel scrap may simultaneously reduce sulphur and phosphorus levels while introducing residual alloying elements or altering nucleation conditions, thereby requiring additional metallurgical control such as inoculation [2], [3].

From a production-oriented perspective, this structured variability can be interpreted in terms of how metallurgical mechanisms translate into observable material behaviour. Some mechanisms act continuously at the batch level, producing gradual shifts in machinability. Others introduce localized heterogeneity, leading to sporadic disturbances. Still others accumulate over time, progressively altering the baseline material condition.

This observation motivates a functional classification of material-related uncertainty, in which elements are grouped according to their dominant mode of influence rather than their nominal chemical identity. Such a classification provides a suitable basis for linking material behaviour to machining response and production-system performance.

1.3.Aim, Research Questions and Scope

The aim of this thesis is to develop a production-oriented modelling framework that links structured compositional variability in grey cast iron to machining response and production-flow consequences. The objective is not to predict the exact metallurgical state of a specific industrial batch, but to construct a transparent and simulation-supported propagation model from compositional tendencies to tool-life consumption, incident risk, bottleneck behaviour, throughput and OEE-related losses.

The first research question is, how can structured compositional variability in grey cast iron be represented through reduced material-response indicators that remain meaningful for machining and production-flow analysis?

The second research question will focus on, how can these material-response indicators be translated into operation-specific machining-response parameters, including processing-time multipliers, tool-life consumption and stochastic incident intensity?

The third and the last research question will be, how do these machining-response parameters affect production-flow behaviour and the resulting per-part unit cost in a serial machining line, including throughput, tool-life consumption, incident burden and the unavoidable cost penalty that persists after process-side cutting-speed re-optimization?

The model is scenario-based and reduced-order. The compositional scenarios are literature-bounded and feedstock-informed. The simulation results should therefore be interpreted as structured scenario outputs rather than direct predictions of a specific factory.

In addition to literature-based modelling and discrete-event simulation, the thesis uses an exploratory industrial practitioner consultation to support the practical interpretation of machining disturbances and production-control responses. This consultation is not treated as a formal qualitative interview study or as empirical validation of the simulation model. Instead, it is used to ground the modelling assumptions and to compare the simulated disturbance mechanisms with industrially recognizable phenomena.

2. Literature Review/Problem Formulation

Grey cast iron has been extensively studied from a metallurgical perspective, particularly with respect to the formation of graphite, the evolution of the matrix, and the role of alloying and trace elements during solidification. Directional solidification studies have shown that cooling rate and inoculation strongly affect graphite type, graphite size, and the spatial location of graphite precipitation, thereby altering the final microstructure in a systematic way [9]. More recent work has further demonstrated that variations in sulphur and manganese, even when other parameters are held constant, can significantly modify graphite flake morphology and distribution. This indicates that the response of grey cast iron to compositional variation is not limited to average chemistry, but is expressed through changes in local graphite geometry and microstructural heterogeneity [10].

A second group of studies has focused on the machining consequences of this microstructural complexity. In finish milling of grey cast iron, the anisotropic distribution of graphite flakes has been shown to inhibit the formation of a high-quality machined surface. Experimental and simulation work indicates that the interaction between the primary shear zone and graphite flakes can generate microcavities on the machined surface, providing a direct mechanism by which material heterogeneity translates into surface roughness or irregular surface features[11]. A related finite element study similarly emphasized that grey cast iron is difficult to machine to a high-quality surface finish precisely because it is a multiphase and anisotropic material, rather than a homogeneous metallic matrix [12]. These studies are particularly relevant to the present work because they show that material-related variability is not merely reflected in bulk hardness or strength, but can appear directly in the generated surface.

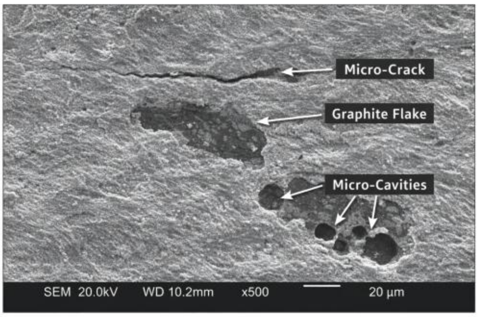

Adapted from: [10] De Santiago-Méndez et al., Materials, 2025; [11] Odum et al., Int. J. Adv. Manuf. Technol.,

Figure 5.Manifestation of microstructural heterogeneity in the machined surface of grey cast iron

Machining studies have also shown that unstable cutting of grey cast iron cannot be understood only in terms of classical abrasive wear. In high-speed machining with pcBN tools, protective transfer layers containing MnS, SiO_2, and Al2O3 have been observed on worn tool surfaces, indicating that the cutting interface evolves dynamically as a function of both cutting conditions and workpiece chemistry [13]. More recent work on ageing effects in grey cast iron reported that tool performance depends strongly on whether a stable built-up layer can form; insufficient protection accelerates wear, while an appropriate cutting speed can promote the formation of a stable Al2O3 and MnS layer[14]. In practical terms, this means that process instability may manifest not only as faster wear, but also as a changing and sometimes unstable interfacial condition, including material accumulation near the cutting edge or on the tool surface. The machining response is therefore partly governed by whether the material–tool system enters a stable or unstable tribochemical regime.

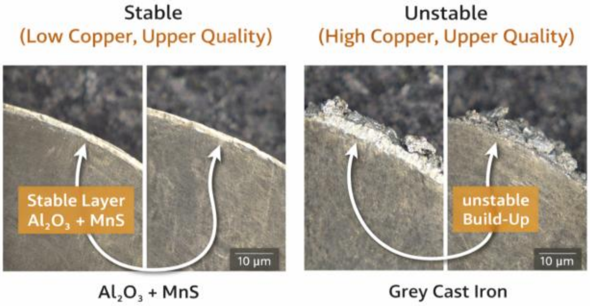

Figure 6.Schematic illustration based on machining studies of grey cast iron

A further body of literature addresses grey cast iron machining from a process-optimization perspective. Studies on turning have quantified how cutting speed, feed, and depth of cut influence surface roughness, cutting force, temperature, and productivity, and have proposed statistical optimization strategies for selecting better cutting conditions [15]. Such work is valuable because it makes the machining consequences measurable and comparable. At the same time, it generally treats the workpiece material as a given condition rather than as a structured source of uncertainty. As a result, the parameter-optimization literature explains how to improve performance for a defined material state, but offers less guidance on how to represent shifts in material behaviour that originate upstream in casting and charge composition.

At the production-system level, the literature is again substantial but differently focused. OEE-based studies treat availability, performance, and quality losses as key dimensions of manufacturing effectiveness and provide a well-established language for analysing productivity losses [16]. Bottleneck research has likewise emphasized that throughput is constrained by limiting stations whose identity may shift dynamically over time, particularly in systems with variability, buffering, and changing operating conditions [17], [18]. These studies are essential for understanding line-level behaviour, but they usually begin from machine states, stop histories, or system events rather than from material-related causes. In other words, production-system models are generally capable of describing disturbances once they have appeared in the system, but less able to describe how they originate from structured variability in the material input itself.

Taken together, the literature reveals a clear but unresolved discontinuity. Metallurgical studies explain how grey cast iron develops different graphite morphologies and local heterogeneities. Machining studies explain how such heterogeneity can influence surface generation, tool wear, and interfacial layer formation. Production studies explain how disturbances reduce OEE and create or shift bottlenecks. What remains insufficiently integrated is the propagation path between these levels. The existing literature does not yet provide a sufficiently explicit framework for describing how structured material-related variability is translated into machining disturbances and then absorbed, amplified, or concealed within production flow.

The problem addressed in this thesis therefore lies not in the absence of knowledge at the individual levels, but in the lack of a connecting representation between them. The central research problem can be formulated as follows: how can structured material-related variability in grey cast iron be represented in a way that remains meaningful at the metallurgical level, yet is also operationally relevant for machining analysis and production-system evaluation? More specifically, how can different forms of material-induced disturbance—continuous microstructural drift, localized heterogeneity, and cumulative hardening tendencies—be translated into analytical variables that explain not only machining response but also throughput loss, bottleneck behaviour, and the use of conservative safety margins in production planning? The following chapter addresses this problem by proposing a functional and production-oriented analytical framework.

Table 1.Literature gap motivating the present modelling framework

Literature stream	Main contribution	Limitation for this thesis	Role in this study
Cast iron metallurgy	Explains graphite, matrix, inclusions, residual elements	Often remains at material/microstructure level	Provides composition-to-material mechanisms
Machining of grey cast iron	Explains tool wear, surface generation, cutting response	Often treats workpiece material as fixed	Provides machining-response mechanisms
Production system / OEE / bottleneck	Explains throughput, downtime, buffering, bottleneck behaviour	Usually starts from machine events, not material-originated causes	Provides production-flow interpretation
Present study	Connects composition patterns to machining parameters and DES outputs	Reduced-order and scenario-based	Provides a transparent propagation framework

3. Methodology and Modelling Framework

3.1. Overall modelling architecture

The methodology is organized as a reduced-order propagation model. The model does not attempt to predict the full metallurgical state of a specific grey cast iron batch. Instead, it constructs a transparent pathway from compositional tendencies to production-flow consequences. The pathway consists of four modelling layers: first, literature-bounded compositional scenarios are defined; second, these scenarios are converted into material-response indicators; third, the material-response indicators are translated into operation-specific machining parameters; fourth, the resulting machining parameters are used in a SimPy-based discrete-event simulation of a serial machining line.

The three layers of the modelling framework correspond to the three research questions: material-response representation, machining-response translation and production-flow evaluation.

The purpose of this structure is to preserve traceability between material mechanisms and production-system outcomes. Each model layer reduces complexity, but also makes explicit where assumptions enter the analysis. The model is therefore intended as an explanatory and simulation-supported framework, rather than as an empirically calibrated prediction model.

3.2. Industrial practitioner consultation

An exploratory industrial practitioner consultation was used to support the practical grounding of the modelling framework. The purpose of the consultation was not to generate statistically generalizable empirical data, but to identify whether the disturbance mechanisms represented in the model are recognizable in industrial machining practice.

The consultation was conducted with a practitioner, who has experience in machining process development and works with process trials, industrial partners and production-related machining problems. The discussion focused on material-related machining stability, tool-life behaviour, sudden tool or quality incidents, monitoring practices and how disturbances are usually detected and handled in production.

The consultation was semi-structured. The main themes included: whether material-related variation is usually gradual or sudden; how tool damage, quality loss or abnormal machining behaviour is detected; how quality gates and monitoring systems are used; whether tool changes are mainly reactive or preventive; and whether predictive models are used in everyday production practice.

The information obtained from the consultation was used qualitatively to guide the interpretation of the simulation model. Several observations were particularly relevant. First, machining processes may appear stable under normal material specifications, but sudden incidents such as tool damage or quality deviations can still occur. Second, such deviations are often detected through quality gates, vibration-related monitoring or process checks. Third, tools are frequently changed preventively to reduce risk. Fourth, predictive models are not always central to everyday shop-floor decision-making, even when monitoring and process-control routines are present.

These observations support the modelling distinction between gradual machining drift, wear-related tool-life burden and stochastic incident mechanisms. They also support the decision to include tool-change states, incident states and event-oriented interpretation in the simulation model. However, the consultation is not used as a direct source for numerical parameter calibration. The numerical scenarios, multipliers and simulation settings remain literature-bounded and assumption-based, as described in the following sections.

Table 2.Industrial consultation themes and modelling implications

Consultation theme	Industrial observation	Modelling implication
Material and process stability	Processes may appear stable under normal specifications, but sudden tool or quality incidents can still occur	Supports separating stable drift from stochastic incident mechanisms
Detection of deviations	Deviations may be detected through quality gates, vibration-related monitoring or process checks	Supports event-oriented indicators in addition to mean throughput
Tool-life management	Tools may be changed preventively to reduce risk	Supports explicit tool-change states and tool-life burden modelling

Consultation theme	Industrial observation	Modelling implication
Practical use of prediction	Predictive models may not be central in everyday production control	Supports treating the model as an explanatory framework rather than a direct industrial predictor
Production response	Incidents may lead to process checks, adjustments or temporary interruptions	Supports modelling incidents as production-flow disturbances

3.3. Functional classification of material uncertainty in cast iron

Material-related uncertainty in cast iron does not arise solely from individual alloying elements, but from the functional roles these elements play within the production system and from the structured compositional variability inherent in industrial feedstock.

Table 3. Dominant material-response modes used for scenario construction

Functional Category	Microstructural Influence	Machining / Process Effects	Production Flow Implications
A. Graphite & Matrix Modifiers，C, Si, Al (low content)	Alters graphite morphology and ferrite–pearlite balance; sensitive to cooling conditions.	Changes in cutting forces, surface and vibration behaviour.	Batch variability; reduced predictability without immediate quality failure.
B. Inclusion & Local Heterogeneity Drivers，Mn (via MnS), S, P	Formation of inclusions, eutectics, or brittle phases with uneven spatial distribution.	Local hard spots; sporadic tool chipping; surface defects; wider tool-life variation.	More tool-changes; unplanned disturbances; harder scheduling and root-cause analysis.

C. Residual / Scrap-Driven Hardeners， Cu, Sn, Cr, Ni and Mo, V, Ti	Stabilization of pearlite or hard phases; cumulative effects from recycling.	Narrower process window; increased sensitivity to parameter deviation.	Tighter safety margins; higher takt-time mismatch risk; gradual loss of flow robustness.

The categories in Table 3 are not mutually exclusive elemental groups. They are reduced functional response modes. In real grey cast iron metallurgy, C, Si, Mn, S, P and residual elements interact through carbon equivalent, Mn/S balance, inoculation, cooling rate, segregation and recycling history. The purpose of the classification is therefore not to isolate individual elements, but to define dominant response channels that can be propagated into machining and production-system variables.

The first category, referred to as graphite and matrix modifiers, includes elements such as carbon, silicon, and aluminium at low concentrations. These elements primarily affect graphite morphology and the balance between ferrite and pearlite. Their influence tends to be continuous rather than localized, leading to systematic shifts in machining behaviour across batches. While such variations rarely cause immediate quality rejection, they reduce predictability and often require implicit or explicit retuning of process parameters. Importantly, these effects are not necessarily driven by changes in a single element, but may result from combined shifts in composition that preserve nominal specifications while altering machinability at the batch level.

The second category, inclusion and local heterogeneity drivers, comprises elements such as manganese (through MnS inclusions), sulphur, and phosphorus. These elements promote the formation of inclusions or brittle phases with uneven spatial distribution. Their impact is therefore highly localized, manifesting as sporadic tool damage, surface defects, or abnormal wear events. From a production perspective, these effects are particularly disruptive, as they introduce instability without clear or repeatable patterns. Such disturbances are consistent with the presence of localized compositional features that are not captured by average material descriptors, but nonetheless dominate operational risk.

The third category, residual or scrap-driven hardeners, includes elements such as copper, tin, chromium, and nickel, as well as other carbide-forming residuals that may accumulate through recycling. These elements exert a gradual but persistent influence by stabilizing harder microstructural constituents. Rather than causing abrupt failures, they progressively narrow the available process window, making production more sensitive to parameter variations and reducing system-level robustness. In this context, individual elemental concentrations are less informative than their cumulative hardening tendency, which reflects long-term compositional drift rather than discrete batch changes.

By distinguishing these three functional categories, the analysis shifts focus from individual elemental effects to their systemic implications for production flow. This classification provides a structured basis for subsequent empirical investigation, including interviews with industrial practitioners and simplified production scenario modelling, while remaining consistent with observed patterns of industrial compositional variability. Beyond their direct impact on machining behaviour, such structured material-related uncertainty also shapes production planning practices, particularly through the use of conservative safety margins. When material variability cannot be explicitly resolved or attributed to specific causes, production systems tend to compensate by underestimating effective tool life and narrowing planned process windows in order to maintain robustness. This suggests that safety margins are not merely a secondary effect, but a central mechanism through which material uncertainty propagates into production flow and cost.

3.4. Construction of literature-bounded compositional scenarios

The compositional scenarios used in this thesis are constructed as literature-bounded and feedstock-informed scenario inputs. Their purpose is to represent plausible compositional tendencies that may arise from charge-material structure, pig iron quality, scrap input and residual-element accumulation.

The scenario construction is based on the assumption that material-related uncertainty in grey cast iron is not introduced as random variation in isolated elements, but as recurring compositional tendencies associated with charge material, pig iron quality, scrap input, recycling history and melt treatment. This is consistent with the feedstock-oriented perspective adopted in this thesis. Recent XRF-based work on pig iron samples from three industrial foundries showed clear differences in sulphur content, total impurity level, and concentrations of carbide-forming and flake-promoting elements. Such observations support the use of structured compositional scenarios rather than a single nominal material state. At the same time, the present scenarios are not intended to reproduce those pig iron measurements directly. They are used as controlled model inputs for testing how different compositional tendencies may propagate through material response, machining behaviour and production flow[19].

Table 4.Compositional scenarios used in the model

Scenario	**Interpretation**	**Main compositional tendency**	**Intended material response**
S0	Baseline	Controlled nominal grey cast iron	Reference
S1	CE/matrix drift	C, Si, P / CE shift	Continuous machinability drift
S2	Local heterogeneity	S, P, Mn/S deviation	Higher local incident and defect risk
S3	Residual hardening	Cu, Sn, Cr, Ni, Mo, V, Ti	Higher hardness / tool wear
S4	Combined adverse	Combined tendencies	Stress-test case

Table 4 summarizes the role of the compositional scenarios used in the model. S0 is the reference condition. S1 represents a continuous CE- and matrix-related drift, where the main expected effect is a gradual change in machinability rather than a discrete failure event. S2 represents a local heterogeneity-sensitive condition associated with sulphur, phosphorus and Mn/S imbalance, where the main expected effect is an increased probability of sporadic incidents or quality disturbances. S3 represents a residual hardening condition associated with accumulated pearlite- and carbide-promoting elements, where the main expected effect is increased tool-life consumption. S4 is retained as an optional stress-test case and should not be interpreted as a typical industrial state.

3.5. Composition descriptors and material-response indicators

The first step in the modelling framework is to convert compositional scenarios into reduced material-response descriptors. The purpose is not to predict the complete metallurgical state of the cast iron, but to represent the compositional tendencies that are most relevant for machining and production-flow analysis. Therefore, the model uses normalized descriptors rather than absolute microstructural predictions. These descriptors preserve the distinction between batch-level drift, local heterogeneity risk and residual hardening tendency, while remaining simple enough to be used as inputs to the subsequent machining and production simulation model.

$$CE = C + \frac{Si+P}{3} \quad \textit{Equation 1}$$

Equation 1 defines the carbon equivalent used in this thesis. CE is used as a reduced descriptor for the combined influence of carbon, silicon and phosphorus. This avoids treating carbon and silicon as fully independent variables and provides a compact way to represent shifts in graphite and matrix tendency at the batch level. In the present model, CE is not used as a full solidification predictor. It is used as a production-oriented descriptor for continuous material-response drift.

$$I_{drift} = \left|\frac{CE-CE_0}{s_{CE}}\right| \quad \textit{Equation 2}$$

Equation 2 defines the drift index used to represent batch-level CE-related deviation from the baseline material state. Here, CE_0 is the carbon equivalent of the baseline scenario, and s_{CE} is a scaling factor used to normalize the magnitude of the CE deviation. The absolute value is used because both upward and downward shifts in CE may indicate a departure from the reference material state and may therefore affect machining response.

$$I_{Mn/S} = ln\left(\frac{(Mn/S)}{(Mn/S)_0}\right) \quad \textit{Equation 3}$$

Equation 3 introduces the Mn/S-related descriptor. This term is included because manganese and sulphur should not be interpreted only as isolated elemental concentrations. Their balance influences sulphide formation and the way sulphur is controlled in the melt. A deviation from the baseline Mn/S condition is therefore interpreted as a potential contributor to local heterogeneity, especially when combined with elevated sulphur or phosphorus levels.

$$E_j^+ = max\left(0, \frac{x_j-x_{j,0}}{s_j}\right) \quad \textit{Equation 4}$$

Equation 4 defines the positive excess term for element j. In this expression, x_j is the concentration of element j, $x_{j,0}$ is the baseline concentration, and s_j is a scaling factor used to normalize the deviation. Only positive deviations are retained because the purpose of this term is to represent the additional contribution of elevated residual or trace elements relative to the baseline condition. The scaling factors are modelling parameters and should be interpreted as normalization constants rather than experimentally fitted values.

$$I_{res} = \sum_{j \in \mathcal{R}} w_j E_j^+ , \quad \mathcal{R} = \{Cu, Sn, Cr, Ni, Mo, V, Ti\} \qquad \textit{Equation 5}$$

Equation 5 defines the residual hardening index. This index combines the positive excess of residual or carbide-forming elements such as Cu, Sn, Cr, Ni, Mo, V and Ti. The index represents a cumulative tendency toward pearlite stabilization, hard-phase formation and increased relative hardness. It is not a thermodynamic phase-fraction calculation. Instead, it is a reduced-order indicator that allows residual-element accumulation to be propagated into tool-life and production-flow consequences.

$$I_{local} = aI_{Mn/S} + bS^+ + +cP^+ \qquad \textit{Equation 6}$$

Equation 6 defines the local heterogeneity index. This index combines Mn/S imbalance, sulphur excess and phosphorus-related risk. It is intended to represent the tendency toward sulphide-related heterogeneity, steadite/phosphide formation and local segregation effects. Unlike the residual hardening index, this term is not expected to primarily increase average cycle time. Its main role in the model is to affect stochastic incident intensity and output variability.

$$\begin{cases} H^* = 1 + \alpha_1 I_{drift} + \alpha_2 I_{res} + \alpha_3 I_{local} \\ M^* = 1 + \beta_1 I_{drift} + \beta_2 I_{res} + \beta_3 I_{local} \end{cases} \qquad \textit{Equation 7}$$

The resulting material-response indicators are expressed as relative factors rather than measured mechanical properties. H^* represents a relative hardness or hard-phase tendency, while M^* represents relative machining difficulty. A value above 1 indicates a more demanding material state compared with the baseline. These indicators should not be interpreted as direct predictions of Brinell hardness, tensile strength or elastic modulus. Their purpose is to provide a transparent intermediate layer between compositional scenarios and machining-response parameters.

The coefficients and scaling factors introduced in this section are modelling parameters. They are used to normalize and weight different compositional tendencies in the reduced-order framework, and they should be calibrated in future work using measured chemistry, hardness data, tool-life records and machining trials.

3.6. Translation from material response to machining response

The second step is to translate the material-response indicators into machining-response parameters. In the simulation model, each operation is affected by three intermediate parameters: the processing-time multiplier $k_{time,i}$, the wear multiplier $k_{wear,i}$, and the incident intensity λ_i. These parameters are operation-specific because different machining processes respond differently to the same material state. For example, rough machining may be more sensitive to general machinability drift, drilling may be more sensitive to localized heterogeneity and sudden tool damage, while heavy finishing may be more sensitive to wear and residual hardening.

$$\begin{cases} k_{time,i} = 1 + a_i I_{drift} + b_i I_{local} + c_i I_{res} \\ k_{wear,i} = 1 + d_i I_{drift} + e_i I_{local} + f_i I_{res} \\ \lambda_i = \lambda_{0,i}(1 + g_i I_{drift} + h_i I_{local} + l_i I_{res}) \end{cases} \quad \textit{Equation 8}$$

In these expressions, i denotes the operation, namely OP10, OP20 or OP30. The coefficients a_i to l_i are sensitivity coefficients. They do not represent universal material constants. Instead, they express the assumed sensitivity of each operation to drift, local heterogeneity and residual hardening. This structure prevents the model from assigning a material mechanism exclusively to one operation. All material-response modes may influence all operations, but their relative influence differs according to the machining process.

Table 5. Operation sensitivity logic

Operation	**Physical meaning**	**Most sensitive to**	**Model expression**
OP10	Rough machining / face milling	General machinability drift	higher $\boldsymbol{time}_{drift}$
OP20	Drilling / constrained tool	Local heterogeneity and incident risk	higher $\boldsymbol{\lambda}_{local}$
OP30	Heavy finishing / turning	Wear and hardening	higher $\boldsymbol{wear}_{res}$

The sensitivity logic in Table 5 therefore replaces a one-to-one mechanism mapping with an operation-dependent response model. OP10 is treated as more sensitive to continuous machinability drift because it represents a rough machining operation with high material-removal interaction. OP20 is treated as more sensitive to localized heterogeneity because drilling involves a constrained cutting environment, limited chip evacuation and a mechanically slender tool. OP30 is treated as more sensitive to residual hardening and wear because it represents a long-duration heavy finishing operation where tool-life consumption and conservative process windows have a direct effect on line capacity.

3.7. Tool-life and incident modelling

Tool life is the main link between material response and production loss in the present model. The classical Taylor tool-life equation is introduced as the conceptual basis for this link.

$$VT^n = C_T \qquad \textit{Equation 9}$$

In this equation, V is cutting speed, T is tool life, n is the Taylor exponent, and C_T is a tool-workpiece constant. Since the present study does not include cutting trials, n and C_T are not fitted. The equation is used to justify the modelling principle that a change in workpiece material condition can change effective tool life even when the nominal process setting remains unchanged.

In Case B and Case C of the simulation study, the Taylor equation is used not only as a conceptual link between material response and tool life but also as the operational coupling between OP30 cutting speed and OP30 tool life. The OP30 cutting speed is varied around its baseline through a normalized ratio $v_{ratio} = v_c / v_{baseline}$. The corresponding nominal cycle time and nominal tool life at any v_{ratio} are computed from the baseline values according to the inverse-proportional relationship between cutting time and cutting speed for a fixed material removal volume per part, and according to Taylor's equation for tool life:

$$t_{cut}(v_{ratio}) = \frac{t_{cut,base}}{v_{ratio}} \qquad \textit{Equation 10}$$

$$T(v_{ratio}) = T_{base} \cdot v_{ratio}^{-1/n} \qquad \textit{Equation 11}$$

The Taylor exponent is taken as n = 0.25, a value in the range commonly reported for cemented-carbide machining of grey cast iron [20], [21]. The material-derived wear multiplier k_{wear} is then applied on top of the Taylor-derived nominal tool life to obtain the effective tool life in the simulation: T_{eff} = T(v_{ratio}) / k_{wear}. This formulation separates two effects on tool life that act through different mechanisms. The first is the cutting-speed effect, which is governed by the Taylor equation and is independent of material scenario. The second is the material-response effect, which is governed by k_{wear} and is independent of cutting speed. The decoupled formulation allows each effect to be examined independently and supports the use of v_{ratio} as a process-engineering decision variable in Case B and Case C.

The Taylor exponent n is treated as a fixed parameter rather than as a function of material scenario. In practice, the wear sensitivity to cutting speed may also vary with material condition, with harder material conditions typically producing a slightly smaller n and therefore a steeper speed–life relationship. Treating n as fixed therefore tends to underestimate the speed sensitivity of tool life under the more adverse material scenarios (S3 and S4) and correspondingly underestimates the wear-driven cost penalty at higher cutting speeds. This simplification is consistent with the role of the model as a structured framework rather than a calibrated predictive tool, and the direction of the simplification is conservative for the principal conclusions of the cost analysis in Section 4.4.

$$T_{eff,i} = \frac{T_{nom,i}}{k_{wear,i}} \qquad \textit{Equation 12}$$

The effective tool life in the simulation is therefore calculated by reducing the nominal tool life according to the wear multiplier. This formulation allows residual hardening, local hard phases or broader machinability shifts to enter the production model through increased tool-change frequency and reduced effective availability.

In addition to continuous wear, some material-related disturbances appear as discrete events. Examples include sudden tool chipping, tool breakage, abnormal surface damage, burr formation or short-term quality deviations. These events are represented using a Poisson incident formulation. The incident intensity λ_i is expressed per cutting hour and is modified by the material-response model.

$$P(N_i(t) = k) = \frac{(\lambda_i t)^k e^{-\lambda_i t}}{k!} \qquad \textit{Equation 13}$$

$$\lambda_i = \lambda_{0,i}\left(1 + g_i I_{drift} + h_i I_{local} + l_i I_{res}\right) \quad \textit{Equation 14}$$

$$P(N_i(t) \geq 1) = 1 - e^{-\lambda_i t} \quad \textit{Equation 15}$$

$$p_{inc,i} = 1 - e^{-\lambda_i t_{cut,i}} \quad \textit{Equation 16}$$

This formulation makes incident risk dependent on both the cutting time and the material scenario. It also keeps the disturbance model compatible with future empirical calibration, since λ_i could later be estimated from tool-breakage logs, quality records, maintenance events or operator reports.

3.8. SimPy-based discrete-event simulation model

The production model is implemented as a Python-based discrete-event simulation using SimPy. The simulated system represents a serial machining line in which parts move through OP10, Buffer1, OP20, Buffer2 and OP30. Each operation has its own processing time, tool-life accumulation, tool-change behaviour and stochastic incident risk. The simulation is used to observe how material-derived changes in k_{time}, k_{wear} and λ propagate into throughput, machine states, buffer levels and bottleneck behaviour.

The simulation is intended to serve as a transparent experimental model for comparing material scenarios under controlled production configurations. It is not intended to reproduce a specific industrial production line. Its value lies in representing finite buffers, blocking, starvation, tool-changes and stochastic incidents over time. These dynamic behaviours cannot be fully captured by a static operation-level calculation.

3.8.1. Production-line structure and machine states

The simulated line follows the structure:

Raw parts → OP10 → Buffer1 → OP20 → Buffer2 → OP30 → Finished parts.

The intermediate buffers have finite capacity. When a downstream buffer is full, the upstream operation becomes blocked. When an operation has no available input part, it becomes starved. Each operation may also enter tool-change or incident states. The machine states recorded in the simulation are therefore busy, blocked, starved, tool-change and incident. These states are used as production-flow indicators, because they show how local machining disturbances are absorbed, amplified or redistributed by the line structure.

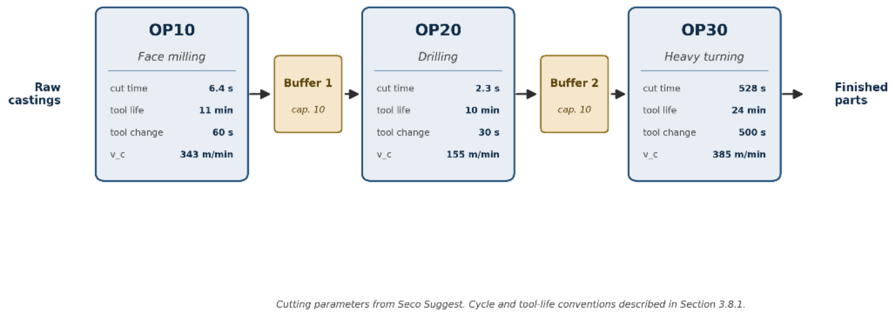

Figure 7.Production-line schematic and per-operation baseline parameters. The line consists of three sequential operations (OP10 face milling, OP20 drilling, OP30 heavy turning) connected by two finite-capacity buffers. Tool-life values include an engineering safety margin applied as part of this study.

The complete process diagram is as follows:

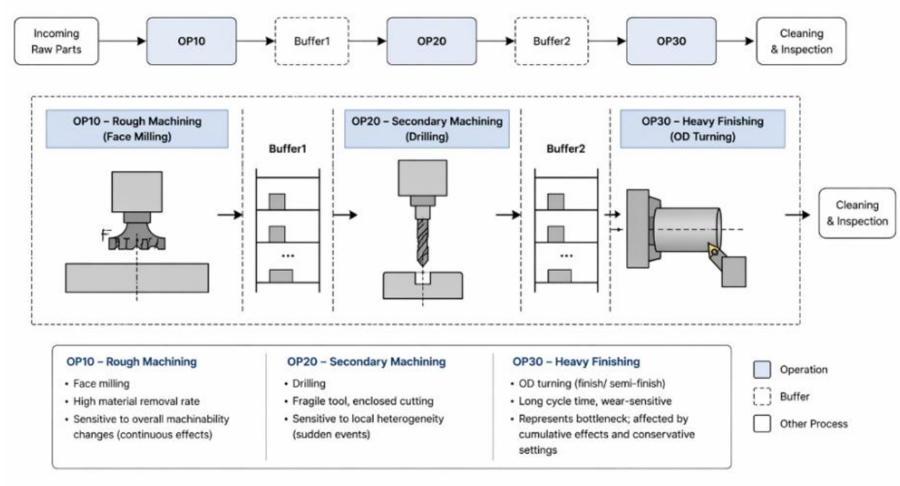

Figure 8. Overall illustration of production simulation

The production system considered in this study is represented as a simplified serial machining line consisting of three consecutive operations (OP10, OP20, OP30) connected by two intermediate buffers (Buffer1 and Buffer2).

The line is preceded by an incoming raw-part stage and followed by a conceptual cleaning and inspection stage.

This representation provides a structured production context in which the propagation of material-related disturbances can be examined at system level. By arranging the machining stages in series, the model makes it possible to trace how local variations in machining behaviour influence downstream flow performance.

Within this line structure, the two buffers play an essential mediating role. They allow certain disturbances to be absorbed locally, while also creating the conditions under which starvation and blocking may emerge when variability exceeds the system's buffering capacity. The buffers therefore act not only as logistical elements, but also as analytical components that reveal how local instability is transformed into flow-level consequences.

Overall, the simplified line serves as an intermediate analytical layer between material-related mechanisms and production-system indicators such as cycle time, throughput, bottleneck behaviour, and cost. In this way, the analysis moves beyond isolated machining responses and examines how uncertainty is redistributed, amplified, or temporarily absorbed within a serial production system.

The baseline cutting parameters used for the three operations are selected using the Seco Suggest cutting-data tool [Seco Tools 2026] as a process-planning aid, and are then verified against the cutting-speed ranges reported in the academic literature on grey cast iron machining. The Seco Suggest tool is used here as an engineering parameter-selection utility consistent with industrial process-planning practice; the academic references cited below establish the reasonableness of the resulting parameter values rather than serving as their primary source. The OP30 turning speed of 385 m/min lies in the upper range of cemented-carbide turning of EN-GJL-250 grey cast iron experimentally examined at 216, 314 and 433 m/min by Twardowski et al. [20]. The OP10 face-milling speed of 343 m/min is consistent with the dry face-milling cutting speeds of 230 and 350 m/min experimentally examined on high-strength cast iron alloys with cemented-carbide tools by Dos Reis et al.[22]. The general use of cemented-carbide tooling at these cutting-speed levels for grey cast iron is documented in standard machining references[23]. The OP30 nominal tool life of 24 minutes incorporates an engineering safety marginal estimated tool life of 38.9 minutes, reflecting the conservative tool-change practice commonly applied in industrial production environments rather than allowing the tool to operate close to its theoretical end-of-life limit.

3.8.2. Operation event logic

For each operation, the simulation follows the same event sequence. First, the operation waits for an available input part. If no part is available, the operation is recorded as starved. Once a part is available, the operation enters the busy state and processes the part using a processing time modified by $k_{time,i}$. During processing, cutting time contributes to accumulated tool-life consumption. When the accumulated tool consumption exceeds the effective tool-life threshold, the operation enters the tool-change state for a defined tool-change duration.

After processing and possible tool-change, the operation may experience a stochastic incident. The incident probability is calculated from the operation-specific incident intensity and the relevant cutting time. If an incident occurs, the operation enters the incident state for a randomly sampled downtime. Finally, the part is transferred to the downstream buffer or to the finished-part sink. If the downstream buffer is full, the operation remains blocked until space becomes available.

Event sequence implemented for each operation is summarized in Box 1.

Box 1. Event logic for each simulated operation

For each operation i:

1. Wait for an available input part.

   If no input part is available, record the operation state as starved.

2. When a part is available, enter the busy state.

   Process the part using the material-dependent processing-time multiplier k_time,i.

3. Accumulate tool-life consumption during cutting.

   If accumulated tool consumption exceeds the effective tool-life threshold,

   enter the tool_change state for the defined tool-change duration.

4. Sample the occurrence of a stochastic incident using the incident intensity λ_i.

   If an incident occurs, enter the incident state for a sampled downtime.

5. Transfer the part to the downstream buffer or finished-part sink.

   If the downstream buffer is full, record the operation state as blocked

   until buffer space becomes available.

3.8.3. Simulation configurations and experimental design

Three simulation configurations are used in this thesis, sharing a single physical line structure (OP10 → Buffer1 → OP20 → Buffer2 → OP30) and identical material-derived machining multipliers. They differ only in which decision variable is examined and which production objective is evaluated.

The first configuration, Case A: configured bottleneck line, treats the line as a fixed structural setup in which OP30 is a single long-cycle, wear-sensitive finishing operation. Each station operates with one machine. Cutting parameters at all stations are held at their nominal values. The purpose of Case A is not to discover the bottleneck, but to examine how material-related uncertainty is amplified when it acts on a structurally constrained operation. Case A therefore establishes the throughput and event-burden baseline against which subsequent configurations are compared.

The second configuration, Case B: cutting-speed sweep with production objective, retains the single-machine line structure of Case A but treats the OP30 cutting speed as an operational decision variable. The OP30 cutting speed is varied around its baseline through a normalized ratio $v_{ratio} = v_c / v_{baseline}$, swept from 0.70 to 1.30 in steps of 0.05. The coupling between cutting speed, cycle time, and tool life is governed by Taylor's tool-life equation (Section 3.7). For each material scenario, the configuration produces a production-rate response curve over v_ratio, from which a scenario-specific production-optimal cutting speed is identified. The purpose of Case B is to test whether different material scenarios shift the production-optimal operating point of the line, rather than merely shifting its peak performance level.

The third configuration, Case C: cutting-speed sweep with cost objective, uses the same v_{ratio} sweep as Case B but evaluates each operating point through a unit-cost layer rather than through throughput alone. The cost layer combines machine-hour cost, tool-change cost, raw-material cost, and a probabilistic scrap cost driven by stochastic incidents. For each material scenario, the configuration produces a unit-cost response curve over v_{ratio}, from which a scenario-specific cost-optimal cutting speed and minimum achievable unit cost are identified. The purpose of Case C is twofold: (i) to translate the production-level effect of material variability into an economic quantity expressed in EUR per finished part, and (ii) to quantify the unavoidable cost penalty that persists for each material scenario even after per-scenario re-optimization of the cutting speed. The unavoidable cost penalty is defined as the difference between the cost-optimal unit cost of a given scenario and that of the baseline scenario, and it represents an upper bound on the per-part value of any incoming-material quality improvement targeting the corresponding material mechanism.

The same material-derived machining multipliers k_{time}, k_{wear} and λ are used across the three configurations. Differences between the configurations therefore reflect the interaction between material response and production-system decision variables (line structure for Case A, cutting speed for Case B and C), rather than changes in the material model itself. A separate cost sensitivity analysis (Section 4.5) further examines the robustness of the Case C cost results to variation in the cost-parameter assumptions.

3.8.4. OEE-oriented interpretation of production losses

Availability-related losses are represented by tool-changes and stochastic incidents. Performance-related losses are represented by material-dependent changes in effective processing time. Flow-level losses are represented by blocking and starvation caused by finite buffers and station-capacity imbalance. Quality losses such as scrap, rework and additional inspection are conceptually recognized but not yet parameterized in the current numerical simulation.

In this thesis, OEE is therefore used as an interpretive framework rather than as a fully calculated indicator. The numerical simulation focuses on the underlying loss states that would contribute to OEE degradation, especially tool changes, incidents, blocking, starvation and reduced processing performance.

The OEE-oriented framing described above is the primary interpretive lens for Case A, in which OP30 is operated at its nominal cutting speed and the question of interest is how material variability degrades the underlying availability, performance and flow components of the line. For Case B and Case C, in which the OP30 cutting speed is varied through the v_{ratio} sweep, the interpretive emphasis shifts from loss states toward operating-point selection. In Case B, the same underlying loss states are present, but they are evaluated against a production-rate objective that depends on v_{ratio}. In Case C, the cost layer described in Section 3.9 aggregates the loss-state contributions into a single per-part economic indicator, and the operating-point selection is examined under a unit-cost objective. The OEE framing and the cost framing are therefore complementary: OEE clarifies how material variability degrades the line, while the cost framing quantifies the economic consequence of that degradation and the value of any intervention against it.

3.8.5. Model inputs, assumptions, verification and limitations

The model combines four types of quantities: baseline machining inputs, constructed compositional scenario inputs, derived intermediate parameters and simulated production outputs. Baseline machining inputs include nominal cutting time, auxiliary time, nominal tool life and tool-change time for OP10, OP20 and OP30. These values define the reference machining line and are treated as planning-level inputs rather than measured production data.

The compositional scenarios are literature-bounded and feedstock-informed inputs. Their role is to represent controlled compositional tendencies that can be propagated through the reduced-order material-response model.

The intermediate machining parameters k_{time}, k_{wear} and λ are derived from the material-response indicators through operation-specific sensitivity coefficients. They are therefore not selected directly as arbitrary scenario multipliers. Throughput, finished parts, machine-state shares, buffer levels, tool-changes, incident counts and dynamic bottleneck indicators are simulation outputs.

In addition to the four types of quantities listed above, the unit-cost layer described in Section 3.9 introduces a fifth type: cost-parameter inputs. These include the machine-hour rates assigned to each operation, the tool-change cost assigned to each operation, the raw-material cost per part and the scrap conversion rate. Cost-parameter inputs are external to the SimPy simulation itself and do not influence the simulation dynamics; they are applied as a post-processing accounting layer over the simulation outputs. The cost-parameter inputs are therefore treated separately from the machining and material inputs of the simulation model. Their assigned values, sources and the sensitivity of the principal cost conclusions to their variation are documented in Sections 3.9 and 4.5.

Table 6.Status of model inputs and method

Model quantity	Source or basis	Role in the model	Status
Baseline cutting and auxiliary times	Planning-level machining data	Define nominal processing behaviour of OP10–OP30	Baseline input
Nominal tool life and tool-change time	Planning-level tooling assumptions	Define reference tool-life accumulation and tool-change behaviour	Baseline input
Compositional scenario values	Literature-bounded and feedstock-informed construction	Represent controlled material input states	Scenario input
CE, Mn/S imbalance, residual hardening index, local heterogeneity index	Calculated from compositional scenarios	Convert chemistry into reduced material-response descriptors	Derived descriptor
H^*, M^*, material-response indicators	Reduced-order material model	Represent relative hardness / hard-phase tendency and machining difficulty	Derived intermediate
k_{time}, k_{wear}, λ	Operation sensitivity model	Translate material response into machining response	Derived simulation input
Finished parts, throughput, machine states, buffer levels, bottleneck indicator	SimPy DES	Quantify production-flow consequences	Simulation output
Machine-hour rates, tool-change costs, raw-material cost, scrap conversion rate	European CNC market rates, grey cast iron market price	Translate simulation outputs into a per-part economic indicator	Cost-layer input
Unit cost per finished part, unit-cost decomposition, unavoidable cost penalty	Cost-layer post-processing of simulation outputs	Quantify economic consequence of material variability and provide upper bound on incoming-material quality improvement value	Cost-layer output

Verification is performed by checking whether the simulation logic behaves consistently with the model assumptions. This includes confirming that parts move through the line in the intended order, that finite buffers generate blocking and starvation, that tool-changes are triggered after accumulated tool-life consumption, and that incident frequency increases when the incident intensity is increased. Validation is limited to plausibility validation based on literature mechanisms, planning-level machining data and industrial interpretation. Full empirical validation would require measured batch chemistry, hardness data, tool-life records, incident logs, quality-gate data and production-flow records.

3.8.6. Simulation outputs

The main simulation outputs are finished parts per shift, throughput, machine-state shares, tool-change counts, incident counts, buffer-level trajectories and dynamic bottleneck indicators. Finished parts and throughput measure production capacity. Machine-state shares distinguish productive time from blocked, starved, tool-change and incident states. Buffer-level trajectories show whether disturbances are absorbed, accumulated or transmitted downstream. The dynamic bottleneck indicator is based on the relative active share of each operation or station, where active time includes busy, tool-change and incident states. This indicator is used to compare whether the system remains dominated by OP30 or whether the limiting operation changes under different material scenarios and capacity configurations.

In addition to these production-flow outputs, the unit-cost layer of Section 3.9 produces a complementary set of cost-related outputs that are derived from the simulation results as a post-processing step. These include the unit cost per finished part, the decomposition of unit cost into machine-hour, tool-change, raw-material and scrap components, the cost-optimal v_{ratio} for each material scenario and the unavoidable cost penalty of each scenario relative to the S0 baseline. The cost-related outputs are reported alongside the production-flow outputs in the results chapter and are interpreted jointly with them in the discussion. Together, the production-flow outputs and the cost-related outputs allow the simulation framework to be examined from two complementary angles: a production-system angle that emphasises throughput, loss states and bottleneck behaviour, and an economic angle that translates these production-system effects into a per-part value indicator suitable for evaluating the cost of supply-chain interventions.

3.9. Cost layer for unit-cost evaluation

The unit-cost layer translates the simulation outputs produced by the SimPy model into a single economic indicator expressed in EUR per finished part. The cost layer is applied as a post-processing step over each replication and is used in Case C and in the sensitivity analysis of Section 4.5. It does not modify the simulation dynamics, the material-derived machining multipliers or the operation-event logic; instead, it interprets the finished output, the tool-change counts and the incident counts produced by the SimPy run through a transparent cost-accounting formulation. This separation between simulation dynamics and cost accounting allows the cost-parameter assumptions to be examined independently of the material model, and supports the structured sensitivity analysis presented in Section 4.5.

The unit-cost layer aggregates four components: machine-hour cost, tool-change cost, raw-material cost and probabilistic scrap cost. The total cost of one 8-hour shift is computed as:

$$C_{total} = C_{machine} + C_{tool} + C_{raw} + C_{scrap} \qquad \textit{Equation 17}$$

The unit cost per finished part is then:

$$c_{unit} = \frac{C_{total}}{N_{finished}} \qquad \textit{Equation 18}$$

where $N_{finished}$ is the number of finished parts produced during the shift.

Machine-hour cost. The machine-hour cost is calculated on a fully-allocated basis, in which each operation incurs its machine-hour rate for the full duration of the shift regardless of its busy, blocked, starved, tool-change or incident state. This is consistent with standard machine-hour costing practice, in which the rate covers depreciation, energy, maintenance, floor space and labour allocation and is incurred when the machine is operating. The total machine-hour cost is therefore:

$$C_{machine} = T_{shift} \cdot (r_{OP10} + r_{OP20} + r_{OP30}) \qquad \textit{Equation 19}$$

where T_{shift} = 8 h and r_i is the machine-hour rate of operation i. The rates assigned in this study are $r_{OP10} = r_{OP20}$ = 65 EUR/h and r_{OP30} = 100 EUR/h, reflecting the role of OP10 and OP20 as 3-axis milling and drilling operations and the role of OP30 as a heavy multi-axis finishing turning operation. The assigned values are taken from the upper-middle of the European commercial CNC machining-rate range [24],[25].

Tool-change cost. The tool-change cost is calculated by multiplying the operation-specific number of tool-change events recorded during the simulation by the operation-specific tool-change cost:

$$C_{tool} = n_{tc,OP10} \cdot p_{tool,OP10} + n_{tc,OP20} \cdot p_{tool,OP20} + n_{tc,OP30} \cdot p_{tool,OP30}$$

Equation 20

The tool-change costs $p_{tool,i}$ are derived from pricing for the cutting tools selected for each operation. For OP10, a face-mill assembly with eight inserts (SNMX2209ANTR-M18 MK1501 at 385 SEK per insert) is used; the effective tool-change cost is set to €134 per event, representing an averaged half-set replacement convention. For OP20, a solid carbide drill (SD207A-1000-062-10R1-P at 2,480 SEK) is used; the effective tool-change cost is set to €77 per event, representing the cost of one drill amortized over an average of 2.8 effective lifetimes through regrinding [26]. For OP30, a turning insert (CNMA190612 TK0501 at 335 SEK per insert) is used; the effective tool-change cost is set to €15 per event, representing an averaged half-insert convention that accounts for indexable use of multiple cutting edges per insert.

The tool prices used here are taken from the Seco commercial pricing platform and represent published list prices for the selected tool grades. Actual procurement prices in industrial settings vary based on volume agreements, framework contracts and supplier negotiations, and the list-price values used here are therefore representative reference figures rather than predicted procurement costs at any specific facility. The use of list prices is consistent with the role of the cost layer in this study as a methodological framework for translating material-derived production effects into a per-part economic indicator, rather than as a procurement-cost prediction tool for industrial implementation.

Raw-material cost. The raw-material cost is the cost of the grey cast iron casting that enters OP10 as input. Each finished part requires one casting, so the raw-material cost for the shift is:

$$C_{raw} = N_{finished} \cdot p_{raw}$$

Equation 21

The assigned value is p_{raw} = 30 EUR per casting, based on European grey cast iron market pricing for finished castings of approximately 5 kg. The market price of grey cast iron metal alone is reported at approximately 3 USD/kg [27]; the casting price used here includes the additional cost of sand-mould casting, energy and yield losses, which roughly doubles the per-kilogram cost relative to the metal price.

Scrap cost. Stochastic incidents recorded at OP30 during the simulation are assumed to produce scrap with a probability α = 0.30. The scrap loss per event is calculated as the raw-material cost of the lost casting plus the accumulated upstream processing cost up to the point of incident. The accumulated upstream processing cost is approximated as the sum of the OP10 cycle-time-weighted unit machine cost, the OP20 cycle-time-weighted unit machine cost, and one half of the OP30 cycle-time-weighted unit machine cost (representing a mid-process incident on average):

$$c_{scrap\text{-}per\text{-}event} = p_{raw} + \frac{t_{OP10} \cdot r_{OP10} + t_{OP20} \cdot r_{OP20} + 0.5 \cdot t_{OP30} \cdot r_{OP30}}{3600} \qquad \textit{Equation 22}$$

$$C_{scrap} = \alpha \cdot n_{incident,OP30} \cdot c_{scrap\text{-}per\text{-}event} \qquad \textit{Equation 23}$$

The assumed scrap conversion rate α = 0.30 is a modelling parameter rather than a measured industrial value. The sensitivity analysis in Section 4.5 examines the influence of this assumption by sweeping α from 0.10 to 0.50 and shows that the principal cost conclusions of Section 4.4 are robust to variation in this parameter within the swept range.

Table 7 summarizes the baseline cost-parameter assignments used in Case C and provides their sources. The sensitivity analysis in Section 4.5 examines variation around these baseline values.

Table 7.Baseline cost-parameter assignments for the unit-cost layer.

Parameter	Symbol	Value	Source
Machine-hour rate, OP10	r_{OP10}	65 EUR/h	Commercial CNC rate, 3-axis milling [25]
Machine-hour rate, OP20	r_{OP20}	65 EUR/h	Commercial CNC rate, 3-axis drilling [25]
Machine-hour rate, OP30	r_{OP30}	100 EUR/h	Commercial CNC rate, heavy turning [25]
Tool-change cost, OP10	$p_{\mathrm{tool,OP10}}$	134 EUR	Seco Suggest, face-mill half-set convention
Tool-change cost, OP20	$p_{\mathrm{tool,OP20}}$	77 EUR	Seco Suggest, drill amortized over 2.8 regrind lifetimes
Tool-change cost, OP30	$p_{\mathrm{tool,OP30}}$	15 EUR	Seco Suggest, insert half-index convention
Raw-material cost per part	p_{raw}	30 EUR	European cast iron market price, ~5 kg part [27]

Parameter	Symbol	Value	Source
Scrap conversion rate	α	0.30	Modelling assumption; sensitivity in Section 4.5

The cost layer is intentionally kept simple and transparent. It uses one rate per machine, one cost per tool-change event, one price per casting and one scalar conversion rate for scrap. More elaborate cost structures, such as time-of-day rates, indexable insert tracking, batch-dependent casting prices or detailed quality cost categories, are not modelled here. This simplification reflects the role of the cost layer in this study as a means of translating production-system effects into an economic indicator, rather than as an industrial cost-quotation tool. The reduced complexity of the cost layer also makes its outputs traceable to specific parameter assumptions, which is essential for the sensitivity analysis presented in Section 4.5.

4. Results

4.1. Parameterization and material-derived machining inputs

4.1.1. Representative compositional scenario inputs

Before the production-flow simulations are presented, the compositional scenarios defined in Chapter 3 are instantiated using representative numerical values. These values are not measured industrial batch chemistries. They are controlled scenario inputs selected within literature-supported grey cast iron composition ranges and arranged to represent distinct material-response tendencies. The purpose of this section is therefore to document how the material scenarios are converted into composition descriptors, material-response indicators and operation-specific machining multipliers for the SimPy simulation [2], [3], [7], [19], [20].

Table 8 lists the representative compositional values used to instantiate the scenario model. S0 is used as the reference material state. S1 is generated mainly by increasing the C-Si-P balance and therefore the carbon-equivalent-related drift [28]. S2 is generated by increasing sulphur and phosphorus and by shifting the Mn/S balance [7], [19]. S3 is generated by increasing residual elements associated with pearlite stabilization or carbide-forming tendency [2], [3], [5]. S4 combines the three tendencies and is used only as an upper-bound stress case.

Table 8.Representative compositional scenario inputs used in the reduced-order model, wt.%

element/scenario	S0 baseline	S1 CE/matrix drift	S2 local heterogeneity	S3 residual hardening	S4 combined adverse
C	3.35	3.45	3.34	3.32	3.42
Si	2.1	2.35	2.08	2.02	2.28
Mn	0.55	0.56	0.42	0.62	0.45
S	0.08	0.075	0.12	0.075	0.115
P	0.06	0.07	0.105	0.07	0.11
Cu	0.2	0.22	0.21	0.48	0.46
Sn	0.03	0.033	0.032	0.08	0.075

Cr	0.08	0.085	0.085	0.18	0.17
Ni	0.1	0.105	0.1	0.26	0.24
Mo	0.02	0.02	0.02	0.06	0.055
V	0.01	0.01	0.012	0.03	0.03
Ti	0.01	0.01	0.012	0.025	0.025

The values in Table 8 are representative scenario inputs rather than measured batch chemistries or material-grade specifications. They are used to instantiate the reduced-order material model and generate the descriptor and multiplier values reported in the following tables.

4.1.2. Composition descriptors and material-response indicators

The compositional values are first converted into reduced descriptors using the equations defined in Chapter 3. Table 9 summarizes the resulting CE, drift index, Mn/S ratio, Mn/S imbalance, residual hardening index and local heterogeneity index. These descriptors provide an intermediate layer between raw composition and machining response. In this way, the model avoids directly assigning a production loss to a single element.

Table 9.Composition descriptors and material-response indicators

scenario	CE	Idrift	Mn/S ratio	IMn/S	Ires	Ilocal	H*	M*
S0 baseline	4.07	0	6.875	0	0	0	1	1
S1 CE/matrix drift	4.2567	1.2444	7.4667	0.0826	0.1036	0.1086	1.039	1.0612
S2 local heterogeneity	4.0683	0.0111	3.5	0.6751	0.1146	0.9681	1.0217	1.0148
S3 residual hardening	4.0167	0.3556	8.2667	0.1843	2.592	0.1544	1.1667	1.1212
S4 combined adverse	4.2167	0.9778	3.913	0.5636	2.3702	0.9108	1.1803	1.1479

The descriptor results should reflect the intended role of each scenario. S1 is expected to show a higher drift index because it is constructed through the C-Si-P balance. S2 is expected to show a higher local heterogeneity index because it combines increased sulphur, increased phosphorus and a shifted Mn/S ratio. S3 is expected to show the highest residual hardening index because Cu, Sn, Cr, Ni, Mo, V and Ti are increased as a group. S4 is expected to produce the strongest combined response and is therefore treated as a stress case rather than a normal industrial condition.

4.1.3. Operation-specific machining multipliers

The material-response indicators are then translated into operation-specific machining multipliers. Table 10 shows the resulting k_{time}, k_{wear} and λ_{mult} values used as direct inputs to the SimPy model. These values are derived from the material-response indicators and the operation sensitivity logic defined in Chapter 3. They are therefore not manually assigned scenario outputs.

Table 10.Material-derived machining multipliers used in the SimPy model

scenario	operation	ktime	kwear	λ_{mult}
S0 baseline	OP10	1	1	1
S0 baseline	OP20	1	1	1
S0 baseline	OP30	1	1	1
S1 CE/matrix drift	OP10	1.0641	1.0804	1.2221
S1 CE/matrix drift	OP20	1.0333	1.0592	1.3077
S1 CE/matrix drift	OP30	1.0521	1.0859	1.2303
S2 local heterogeneity	OP10	1.0077	1.0155	1.2529
S2 local heterogeneity	OP20	1.0078	1.0278	1.9835
S2 local heterogeneity	OP30	1.0064	1.0352	1.6421
S3 residual hardening	OP10	1.0498	1.1395	1.2993
S3 residual hardening	OP20	1.0487	1.1987	1.5188
S3 residual hardening	OP30	1.0615	1.4111	1.4022
S4 combined adverse	OP10	1.0828	1.1744	1.564
S4 combined adverse	OP20	1.0655	1.2232	2.3418
S4 combined adverse	OP30	1.0854	1.4257	1.9463

The multiplier table shows how the same material scenario affects operations differently. S1 mainly increases processing-time-related multipliers, especially for operations sensitive to continuous machinability drift. S2 mainly increases incident intensity, particularly for constrained or event-sensitive operations. S3 mainly increases wear-related multipliers, especially at the heavy finishing operation. This confirms that the model does not bind one material mechanism to one operation. Instead, all scenarios can affect all operations, but their effects differ through operation-specific sensitivity coefficients.

These material-derived multipliers provide the input layer for the following production-flow simulations. The subsequent sections examine how these machining-level changes propagate through three production configurations: a configured bottleneck line, a capacity-adjusted line and an OP30 capacity-sweep experiment.

4.2. Case A: configured bottleneck line

The first production-flow experiment uses the configured bottleneck line. In this configuration, OP30 is kept as a single long-cycle finishing operation with high tool-life burden. The purpose of this case is not to identify an emerging bottleneck, because OP30 is intentionally configured as the dominant constraint. Instead, Case A is used to examine how material-derived machining multipliers affect throughput, tool-change burden, incident occurrence and upstream blocking when the line is structurally dominated by the final operation.

Table 11. Case A simulation summary for the configured bottleneck line

scenario	S0 baseline	S1 CE/matrix drift	S2 local heterogeneity	S3 residual hardening	S4 combined adverse
$finished_{mean}$	39.725	36.78	38.835	33.795	32.78
$finished_{std}$	0.5299	0.5029	0.457	0.5042	0.5225
$finishedq_{05}$	39	36	38	33	31.95
$finishedq_{95}$	40	37	39	34	33
$throughput_{mean}$	4.9656	4.5975	4.8544	4.2244	4.0975
$throughput_{std}$	0.0662	0.0629	0.0571	0.063	0.0653
Output loss vs S0 percent	0	7.4135	2.2404	14.9276	17.4827
OP10 $blocked_{mean}$	0.9553	0.9555	0.9561	0.9577	0.9577
OP20 $blocked_{mean}$	0.9729	0.974	0.9732	0.9752	0.9756
OP30 $active_{mean}$	0.9979	0.9979	0.9979	0.9979	0.9979
OP30 $toolchanges_{mean}$	13.96	14.96	13.995	17.96	17.95
$OP30incidents_{mean}$	0.6	0.685	0.965	0.695	0.99

Table 11 shows that the baseline scenario produces an average of 39.73 finished parts per 8-hour shift, corresponding to 4.96 parts per hour. Compared with the baseline, S1 CE/matrix drift reduces the mean output to 36.78 parts per shift, corresponding to an output loss of approximately 7.4%. S2 local heterogeneity produces a smaller mean output loss of approximately 2.2%, while S3 residual hardening and S4 combined adverse reduce the mean output by approximately 14.9% and 17.5%, respectively. This ranking is consistent with the material-derived machining multipliers reported in Section 4.1: S3 and S4 impose a stronger wear burden on OP30, while S2 mainly increases incident intensity without strongly increasing average processing time or wear.

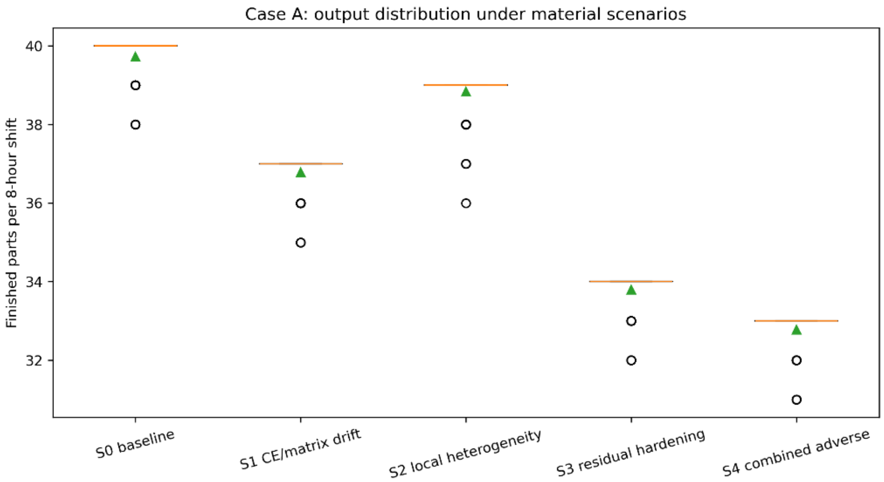

Figure 9.Output distribution under material scenarios in Case A. Boxplots are based on 200 Monte Carlo replications. Green triangles indicate mean values.

Figure 9 confirms the same pattern at the distribution level. The baseline case remains close to 40 finished parts per shift. S1 shifts the output distribution downward, while S3 and S4 produce the strongest output reduction. S2 remains closer to the baseline in terms of mean output. This indicates that, under a strongly configured OP30 bottleneck, residual hardening and combined adverse material conditions have a more direct effect on line output than local heterogeneity alone.

The state distribution further clarifies why the output response is dominated by OP30. In the representative baseline run shown in Figure 10, OP10 and OP20 spend most of the sampled time in the blocked state, while OP30 is almost continuously active, either processing parts or undergoing tool changes. This means that the upstream operations have more nominal capacity than the downstream finishing operation can absorb. As a result, upstream material-related effects are largely masked by blocking in Case A.

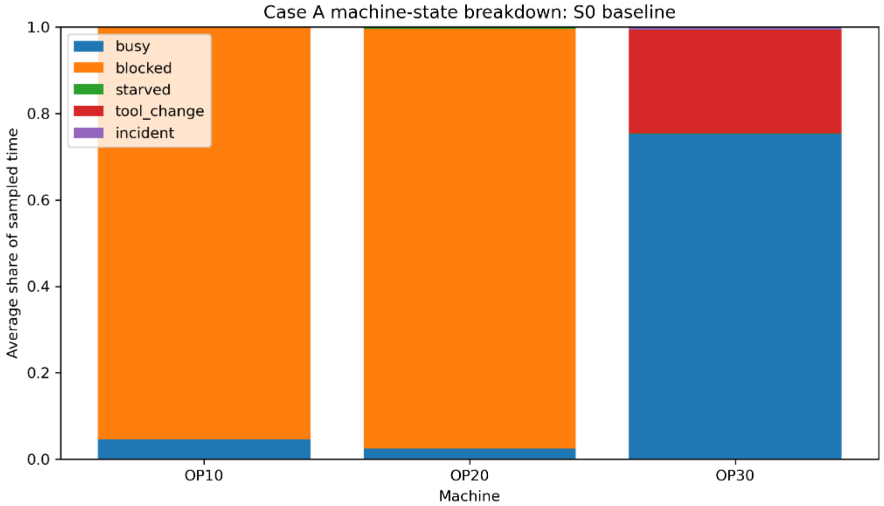

Figure 10.Machine-state breakdown for the baseline scenario in Case A. The figure shows a representative replication of the configured bottleneck line.

This state pattern shows that Case A is dominated by a structural bottleneck rather than by a dynamically shifting constraint. OP10 and OP20 are not heavily utilized as independent production constraints; their dominant loss state is blocking. OP30, by contrast, remains the controlling station. Therefore, changes in OP30 processing time, tool-life consumption and incidents are directly reflected in the total output of the line.

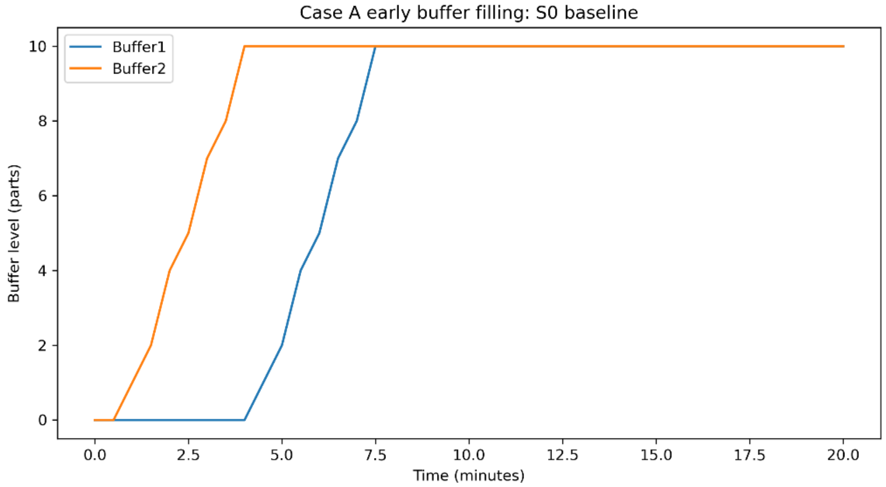

Figure 11.Early buffer filling behaviour for the baseline scenario in Case A. The figure shows the first 20 minutes of a representative replication.

The buffer trajectory in Figure 11 supports the same interpretation. Buffer2 fills first, followed by Buffer1. After both buffers reach their capacity limit, upstream stations become blocked for most of the shift. This shows how the configured OP30 bottleneck rapidly propagates upstream through finite buffers. The buffers absorb the initial flow imbalance only briefly; after that, they mainly reveal the structural capacity mismatch between the upstream operations and OP30.

Because OP30 dominates the configured bottleneck line, event burden at OP30 is the most important local mechanism behind the output losses observed in Case A. Figure 12 compares the mean number of OP30 tool changes and incidents per 8-hour shift across the material scenarios.

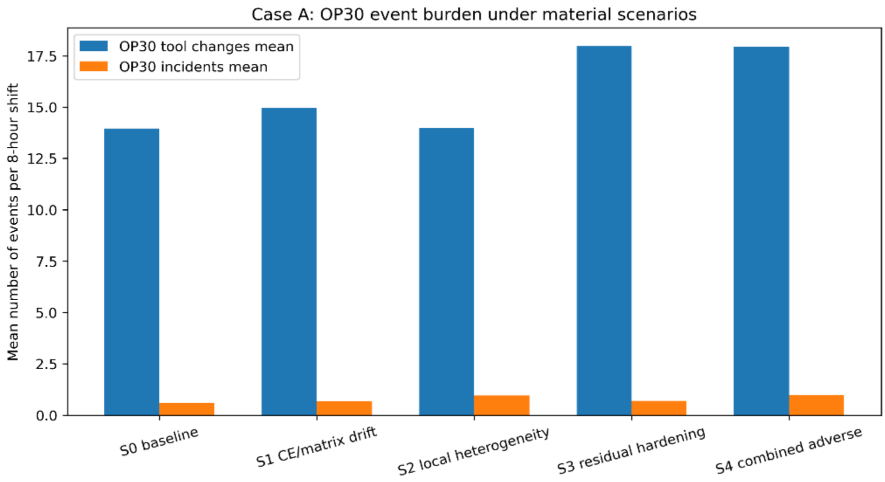

Figure 12.OP30 event burden under material scenarios in Case A. Values are averaged over 200 Monte Carlo replications.

The strongest difference between scenarios is observed in OP30 tool changes. S3 residual hardening and S4 combined adverse show the highest OP30 tool-change burden, which explains their larger output losses. S2 local heterogeneity increases the incident burden relative to the baseline, but the absolute number of incidents remains small compared with tool changes. Therefore, the effect of S2 on average output is limited in Case A. This does not mean that local heterogeneity is unimportant; rather, it indicates that the configured bottleneck structure masks part of the stochastic disturbance mechanism. A production configuration with less dominant OP30 capacity is therefore needed to examine whether local heterogeneity has a stronger effect on output variability and bottleneck behaviour.

Overall, Case A demonstrates both amplification and masking effects. Material conditions that increase OP30 processing time or tool-life consumption are amplified because OP30 directly controls line throughput. This is why S3 and S4 cause the largest output losses. At the same time, upstream or event-based disturbances are partly masked because OP10 and OP20 are blocked for most of the shift. Case A therefore provides a useful structural baseline, but it is not sufficient for studying dynamic bottleneck shifting. For this reason, the next simulation configuration relaxes the structural dominance of OP30 and examines whether the bottleneck becomes scenario-dependent.

4.3. Case B: cutting-speed sweep with production objective

Case B retains the single-machine line structure of Case A and uses the same material-derived machining multipliers. Its decision variable is the OP30 cutting speed, expressed as the normalized ratio $v_{ratio} = v_c / v_{baseline}$, where $v_{baseline}$ corresponds to the recommended cutting speed for the OP30 turning operation. The v_{ratio} is swept from 0.70 to 1.30 in steps of 0.05, generating 13 operating points per material scenario. The coupling between v_{ratio}, OP30 cycle time and OP30 tool life is governed by Taylor's tool-life equation (Section 3.7), with cycle time inversely proportional to v_{ratio} and tool life proportional to $v_{ratio}^{(-1/n)}$ at $n = 0.25$ (adopted as a literature-typical value for grey cast iron with carbide tooling [29]). For each (scenario, v_{ratio}) pair, 200 Monte Carlo replications of an 8-hour shift are executed. Case B therefore characterizes the response of finished output to OP30 cutting speed under each material scenario, and identifies the scenario-specific production-optimal v_{ratio}.

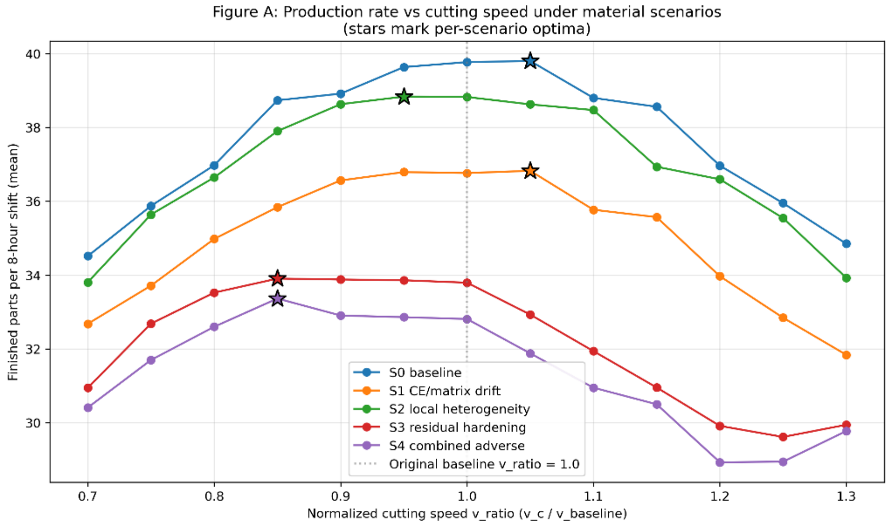

Figure 13. Production rate vs normalized OP30 cutting speed v_{ratio} under five material scenarios. Each marker is the mean finished output per 8-hour shift over 200 Monte Carlo replications. Stars mark the production-optimal v_{ratio} for each scenario. The dotted vertical line indicates the baseline (v_{ratio} = 1.0).

Figure 13 shows that the production-rate response to v_{ratio} is non-monotonic in all five material scenarios, producing a peak in the v_{ratio} range of 0.85–1.05. At low v_{ratio} the cutting time per part dominates and limits throughput; at high v_{ratio} the rapidly shortening tool life increases tool-change frequency and ultimately reduces throughput. This trade-off is the production-system manifestation of the speed–life relationship described by Taylor's equation. The non-monotonic behaviour is an emergent property of the Taylor coupling and is not imposed by the material multipliers, since the underlying k_{time}, k_{wear} and lambda multipliers are independent of v_{ratio}.

The location of the production-optimal v_{ratio} differs between scenarios. For S0 baseline and S1 CE/matrix drift, the production-optimal v_{ratio} is 1.05, indicating that the baseline speed is close to but slightly below the production-optimal speed under nominal material conditions. For S2 local heterogeneity, the production-optimal v_{ratio} shifts moderately to 0.95, reflecting the indirect coupling between incident intensity and effective cutting time. For S3 residual hardening and S4 combined adverse, the production-optimal v_{ratio} shifts to 0.85, a 20% reduction relative to the S0 optimum. The downward shift of the optimum under residual hardening reflects the elevated wear multiplier k_{wear} at OP30 (1.41 for S3, 1.43 for S4), which steepens the tool-life penalty at higher cutting speeds. Table 12 summarizes the production-optimal v_{ratio}, the corresponding mean output and the unavoidable production loss for each scenario.

Table 12.Production-optimal cutting speed and unavoidable production loss by material scenario.

Scenario	Optimal v_{ratio}	Mean output at optimum (parts/shift)	Output at v_{ratio} = 1.0 (parts/shift)	Unavoidable production loss vs S0 (%)
S0 baseline	1.05	39.80	39.77	0.00
S1 CE/matrix drift	1.05	36.83	36.77	7.48
S2 local heterogeneity	0.95	38.84	38.83	2.43
S3 residual hardening	0.85	33.91	33.80	14.81
S4 combined adverse	0.85	33.37	32.81	16.17

Three observations follow from Figure 13 and Table 12. First, the production-rate hierarchy of the material scenarios is preserved across the swept range: S0 ⩾ S2 > S1 > S3 ⩾ S4 holds at every v_{ratio}. Local heterogeneity remains close to the baseline in terms of mean output, while residual hardening and combined adverse conditions produce substantial output reductions of approximately 15% and 16% respectively. This confirms the central observation of Case A — that residual element accumulation has a stronger effect on mean production capacity than local heterogeneity — and shows that this hierarchy is robust to operating-point optimization.

Second, the production-optimal v_{ratio} shifts toward slower cutting speeds as material conditions worsen. The 0.20 leftward shift between S0 (v_{ratio} = 1.05) and S3 (v_{ratio} = 0.85) is a process-engineering decision of considerable magnitude, corresponding to a 19% reduction in nominal cutting speed. This shift is the emergent expression of the wear–speed coupling under Taylor's equation: when the wear multiplier increases, the marginal cost of additional cutting speed in terms of lost tool life outweighs the marginal benefit in terms of shortened cycle time at a lower threshold. The operating-point shift is therefore a quantitative consequence of the material model rather than a degree of freedom in the simulation.

Third, the difference between scenario output at v_{ratio} = 1.0 (the baseline) and at the scenario's own production-optimal v_{ratio} is small in absolute terms (typically less than 1% for S0–S3). This indicates that the production-rate response curves are relatively flat near their peaks, and that re-optimization of v_{ratio} yields limited improvement in production rate alone. The economic consequence of this flatness, examined in Case C, is that production-rate optimization and cost optimization produce different optimal cutting speeds, and that the unavoidable performance penalty of material variability is more clearly visible in cost units than in production units.

The Case B finding that the production-optimal v_{ratio} is scenario-dependent reproduces, within the present model, the classical machining-economics result that economic cutting speed lies below maximum cutting speed and depends on material properties. The present analysis extends this classical result to the production-system level by showing that the cutting-speed optimum is not only sensitive to nominal material grade but also to within-grade compositional variability that is normally classified as acceptable feedstock variation.

4.4. Case C: cutting-speed sweep with cost objective

Case C uses the same v_{ratio} sweep and the same simulation outputs as Case B, but evaluates each operating point through a unit-cost layer that converts the simulated finished output, tool-change counts and incident counts into a single economic indicator expressed in EUR per finished part. The purpose of Case C is to identify, for each material scenario, the cost-optimal cutting speed and the corresponding minimum achievable unit cost, and to quantify the unavoidable cost penalty that persists for each scenario even after per-scenario re-optimization of v_{ratio}. This unavoidable cost penalty is interpreted as an upper bound on the per-part value of any incoming-material quality improvement targeting the corresponding material mechanism. The unit-cost layer aggregates four components: machine-hour cost, tool-change cost, raw-material cost and a probabilistic scrap cost driven by stochastic incidents. Machine-hour rates are assigned at €65/h for OP10 and OP20 and €100/h for OP30, with the higher OP30 rate reflecting its role as a heavy finishing turning operation [24]. Tool-change costs per event are derived from Seco commercial pricing: €134 for an OP10 face-milling tool change (eight-insert face mill assembly, indexed at an effective half-set rate), €77 for an OP20 drill change (a 216 EUR solid carbide drill amortized over an average of 2.8 effective lifetimes through regrinding) and €15 for an OP30 turning insert change (29.20 EUR per insert, indexed at an effective half-insert rate per change). Raw material cost is assigned at €30 per part based on European grey cast iron market pricing for finished castings. Scrap cost per event is computed as the raw material cost plus the accumulated upstream processing cost at the time of incident, multiplied by an assumed scrap conversion rate of 30%. The robustness of the cost results to variation in these cost parameters is examined in Section 4.5.

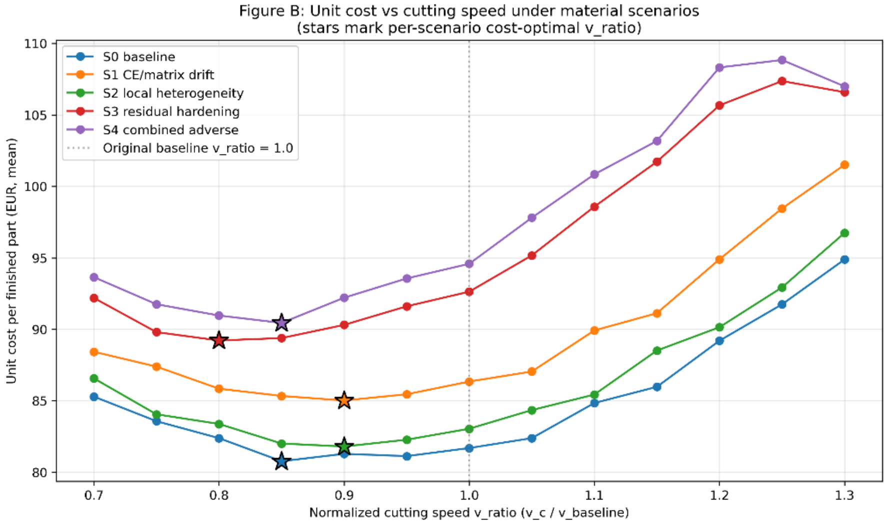

Figure 14.Unit cost per finished part as a function of normalized OP30 cutting speed v_{ratio}, under five material scenarios. Each marker is the mean unit cost over 200 Monte Carlo replications. Stars mark the cost-optimal v_{ratio} for each scenario. The dotted vertical line indicates the baseline (v_{ratio} = 1.0).

Figure 14 shows that the unit-cost response to v_{ratio} is U-shaped in all five material scenarios, with cost-optimal speeds in the range v_{ratio} = 0.80~ t0.90. The U-shape arises from the same speed–life trade-off that produced the production-rate peaks in Case B, but the cost objective penalizes both excessive cutting time (through machine-hour cost) and excessive tool consumption (through tool-change cost) more directly than the production-rate objective. As a result, the cost-optimal v_{ratio} lies systematically below the production-optimal v_{ratio} identified in Case B. For S0 baseline, the production optimum is at v_{ratio} = 1.05 while the cost optimum is at v_{ratio}= 0.85. For S3 residual hardening, the production optimum is at v_{ratio} = 0.85 while the cost optimum is at v_{ratio} = 0.80. This pattern reproduces, within the present model, the classical machining-economics result that the economic cutting speed lies below the maximum-production cutting speed [30].

Table 13 summarizes the cost-optimal v_{ratio}, the corresponding minimum achievable unit cost, and the unavoidable cost penalty relative to the S0 baseline minimum unit cost.

Table 13.Cost-optimal cutting speed and unavoidable cost penalty by material scenario.

Scenario	Cost-optimal v_{ratio}	Minimum unit cost (EUR/part)	Unavoidable cost penalty vs S0 (%)	Penalty in EUR per part
S0 baseline	0.85	80.76	0.00	0.00
S1 CE/matrix drift	0.90	85.02	5.27	4.26
S2 local heterogeneity	0.90	81.79	1.28	1.03
S3 residual hardening	0.80	89.21	10.46	8.45
S4 combined adverse	0.85	90.44	11.98	9.68

The four observations that follow from Figure 14 and Table 13 are central to the economic interpretation of the simulation framework.

First, the minimum achievable unit cost varies significantly across material scenarios. S0 baseline reaches €80.76 per part at its cost optimum, S2 local heterogeneity reaches €81.79 (a small €1.03 penalty), S1 CE/matrix drift reaches €85.02 (a €4.26 penalty), S3 residual hardening reaches €89.21 (a substantial €8.45 penalty) and S4 combined adverse reaches €90.44 (a €9.68 penalty). These differences represent the unavoidable economic consequence of the corresponding material conditions, in the sense that they persist even when the cutting speed is re-optimized for each scenario individually. An ideal process operator with full advance knowledge of the incoming material state and the freedom to set v_{ratio} independently for each batch would still face this cost increase.

Second, the unavoidable cost penalty hierarchy differs substantially in magnitude from the unavoidable production-rate loss hierarchy identified in Case B. S3 residual hardening produces a +10.46% cost penalty, an order of magnitude larger than S2 local heterogeneity at +1.28%. The eight-fold ratio between the S3 and S2 penalties indicates that residual element accumulation has a dominant economic consequence under the present cost structure, while local heterogeneity has a comparatively minor economic consequence after process re-optimization. This finding has direct implications for incoming-material quality control: a per-part investment ceiling of approximately €8.45 is justified for any supply-chain measure that reduces residual element accumulation toward the S0 baseline level, whereas the corresponding ceiling for measures targeting local heterogeneity is approximately €1.03 per part.

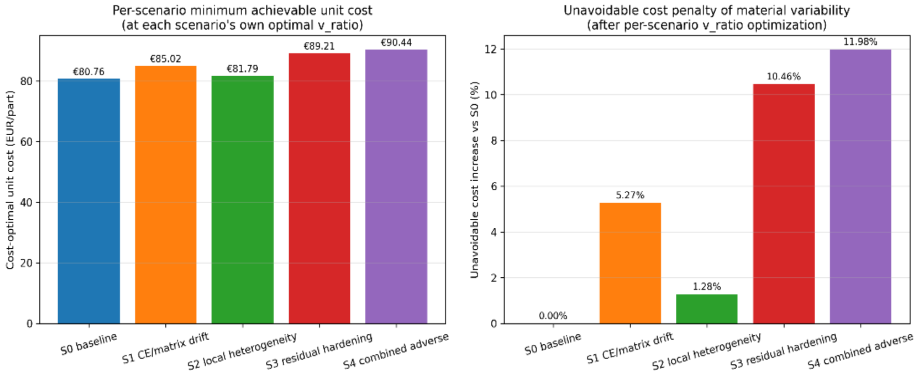

Figure 15.Per-scenario minimum achievable unit cost (left) and unavoidable cost penalty relative to S0 (right), under each scenario's own cost-optimal v_{ratio}. Bars are labelled with the absolute cost and the percentage penalty.

Third, the unit-cost decomposition at each scenario's cost-optimal v_{ratio} shows that machine-hour cost and raw-material cost together account for more than 90% of total unit cost in all scenarios, with tool-change cost contributing approximately 5–7% and scrap cost contributing less than 1%. Figure 16 shows this decomposition. The dominant contribution of machine-hour cost reflects the fact that OP30, as a heavy finishing operation with a baseline cycle time of 528 seconds, occupies the machine for an extended period per part. Material variability therefore propagates to unit cost primarily through machine-time amplification driven by throughput loss, rather than through direct tool or quality channels. This is an important interpretive result. Although the operational discourse around grey cast iron variability often emphasizes tool wear and surface-quality incidents, the cost decomposition indicates that the economic impact at the production-system level is dominated by the indirect throughput channel. A material condition that increases tool-life consumption affects unit cost not only through direct tool-change cost, but more substantially through the reduction in OP30 throughput that this consumption causes.

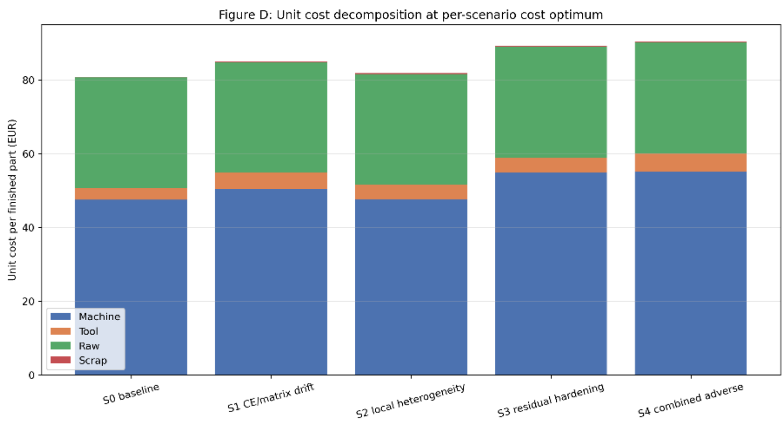

Figure 16. Decomposition of unit cost per finished part at each scenario's own cost-optimal v_{ratio}. Components shown are machine-hour cost, tool-change cost, raw-material cost and probabilistic scrap cost.

Fourth, the cost-optimal v_{ratio} shifts only moderately between material scenarios, with all values falling in the narrow range 0.80–0.90. This narrowness contrasts with the more pronounced shift seen in the production-optimal v_{ratio} (1.05 for S0, 0.85 for S3). The reason is that the cost objective is dominated by machine-hour cost, which is approximately monotonic in cycle time and therefore decreases with increasing v_{ratio} over much of the swept range. The tool-change cost penalty at high v_{ratio} pulls the cost-optimal speed downward, but the optimum is broad and relatively flat in cost terms. The narrowness of the cost-optimal v_{ratio} range is operationally significant: it indicates that a single conservative cutting-speed setting (approximately v_{ratio} = 0.85) is reasonably close to optimal for all material scenarios examined. Re-optimization of v_{ratio} on a per-batch basis is therefore not the principal lever for mitigating the cost consequences of material variability. The dominant lever is the underlying material condition itself.

The Case C results provide the principal quantitative answer to the research question of this thesis. Material variability in grey cast iron creates an unavoidable unit-cost penalty that ranges from +1.28% for local heterogeneity to +11.98% for combined adverse conditions, expressed in EUR per part as €1.03 to €9.68. The penalty for residual element accumulation alone is +10.46% or €8.45 per part. These values establish a quantitative upper bound on the per-part economic value of incoming-material quality improvements, and may provide a first-order estimate for evaluating the economic justification of supply-chain measures targeting specific material mechanisms.

4.5. Cost sensitivity Analysis

The cost results presented in Section 4.4 depend on four assigned cost parameters: machine-hour rates, tool-change costs, raw-material cost and scrap conversion rate. Each parameter carries some degree of uncertainty. Machine-hour rates vary across European manufacturing contexts, tool-change costs depend on insert indexing and procurement practices, raw-material prices fluctuate with market conditions, and the scrap conversion rate is an assumed rather than measured parameter. To examine whether the principal Case C conclusions are robust to variation in these parameters, this section reports the results of a structured sensitivity sweep over three of them: scrap conversion rate, raw-material cost per part, and OP30 tool-change cost. The OP10 and OP20 tool-change costs are scaled proportionally with the OP30 cost. Three values are examined for each parameter: scrap conversion rate at 0.10, 0.30 and 0.50; raw-material cost at €20, €30 and €40 per part; OP30 tool-change cost at €10, €15 and €30 per event. This produces a full factorial of 27 parameter combinations. For each combination, the cost-optimal v_{ratio} and minimum unit cost are re-computed for every material scenario using the same 200-replication Monte Carlo dataset as in Section 4.4. The unavoidable cost penalty of S3 residual hardening and S4 combined adverse relative to S0 is then evaluated for each combination. Table 14 summarizes the resulting penalty ranges. Across all 27 parameter combinations, the S3 residual hardening unavoidable cost penalty ranges from +9.05% to +12.42% with a median of +10.46%. The S4 combined adverse penalty ranges from +10.15% to +15.23% with a median of +12.12%. The penalty therefore remains within a band of approximately ±2 percentage points around the central estimate across a wide range of cost-parameter assumptions, including a ±66% range on the scrap conversion rate, a ±33% range on raw-material cost and a –33% to +100% range on tool-change cost.

Table 14.Range of unavoidable cost penalty across 27 cost-parameter combinations.

Scenario	Minimum penalty (%)	Median penalty (%)	Maximum penalty (%)	Penalty range (percentage points)
S3 residual hardening	9.05	10.46	12.42	3.37

Scenario	Minimum penalty (%)	Median penalty (%)	Maximum penalty (%)	Penalty range (percentage points)
S4 combined adverse	10.15	12.12	15.23	5.08

Two interpretive observations follow from Table 14. First, the qualitative conclusion of Case C — that residual hardening creates approximately a 10% unavoidable cost penalty and that combined adverse conditions create approximately a 12% unavoidable cost penalty — is robust to parameter uncertainty within the swept ranges. The order of magnitude of the penalty does not change, although the exact percentage shifts. The principal finding is therefore not an artefact of specific cost-parameter choices. Second, the scrap conversion rate has only a minor influence on the S3 and S4 penalties relative to the influence of raw-material cost and tool-change cost. Varying the scrap conversion rate across the full sweep range produces less than a 0.3 percentage-point change in the S3 penalty for any fixed combination of raw-material cost and tool-change cost. This is consistent with the unit-cost decomposition in Section 4.4, which showed that scrap cost contributes less than 1% of total unit cost at the cost-optimal v_{ratio}. The cost-driving parameters in the present model are therefore raw-material cost and tool-change cost, not the scrap conversion rate. The robustness of the headline penalty to scrap conversion rate variation indicates that the assumed value of 0.30 is not a critical assumption for the principal cost conclusion.

The sensitivity analysis also yields a robust per-part economic value range for incoming-material quality improvement. Across the swept parameter combinations, the per-part unavoidable cost penalty of S3 residual hardening ranges from €6.68 to €11.16, with a central estimate of €8.45. This range represents the per-part economic value ceiling for any supply-chain measure that reduces residual element accumulation from the S3 condition toward the S0 baseline condition, under the cost-parameter uncertainty examined.

The cost sensitivity analysis presented above varies the scrap conversion rate, raw-material cost and tool-change cost. The machine-hour rate is not explicitly varied, because in the fully-allocated machine-cost convention used here, the absolute increment $\Delta C_{machine}$ scales linearly with the machine-hour rate. Increasing OP30's hourly rate from €100 to €150 would, for example, increase S3's machine-channel increment from €7.39 to €11.09 per part, but would leave the tool and scrap channel increments unchanged. The qualitative conclusion that S3 residual hardening produces a substantially larger cost penalty than S2 local heterogeneity, and that the penalty propagates primarily through the machine-time channel, is therefore robust to machine-hour rate variation; only the absolute magnitude of the penalty shifts.

4.6. Cross-case synthesis

The three simulation configurations of this thesis provide complementary views of how material-derived machining changes propagate to production-system behaviour. Case A characterizes throughput response under a fixed line structure with cutting parameters held at nominal values; Case B characterizes the production-rate response across a range of OP30 cutting speeds; Case C extends the cutting-speed sweep with a unit-cost layer that produces a single economic indicator per operating point. Together, the three cases form a layered analysis in which each layer makes a different decision variable explicit: line structure in Case A, cutting speed under a production objective in Case B, and cutting speed under a cost objective in Case C. The cross-case synthesis presented here identifies four findings that emerge from the joint interpretation of these three layers and that are central to the principal conclusions of this thesis.

The first cross-case finding is that the relative impact ordering of material scenarios is preserved across all three configurations and across both production and cost objectives. In Case A under the nominal cutting speed, S3 residual hardening and S4 combined adverse produce the largest throughput reductions, S1 CE/matrix drift produces a moderate reduction, and S2 local heterogeneity produces only a small mean output loss. In Case B, this hierarchy is reproduced at every v_{ratio} in the swept range. In Case C, the same hierarchy reappears in the cost domain, with S3 and S4 producing the largest unavoidable cost penalties (+10.46% and +11.98%) and S2 producing the smallest (+1.28%). The consistency of the hierarchy across three configurations and two objectives indicates that the relative ordering reflects a property of the material model itself, rather than a property of any particular decision variable or production-system structure.

The second cross-case finding is that the production-optimal cutting speed and the cost-optimal cutting speed do not coincide for any material scenario. In Case B, the production-optimal v_{ratio} for S0 baseline is 1.05; in Case C, the corresponding cost-optimal v_{ratio} is 0.85. For S3 residual hardening, the production-optimal v_{ratio} is 0.85 and the cost-optimal v_{ratio} is 0.80. Across all five scenarios, the cost-optimal v_{ratio} lies systematically below the production-optimal v_{ratio}. Table 15 summarizes the per-scenario optima under the two objectives. The downward shift of the cost optimum reflects the fact that the cost objective penalizes both excessive cutting times, through machine-hour cost, and excessive tool consumption, through tool-change cost, while the production objective penalizes only the second of these channels. This finding indicates that a production-rate optimization strategy and a unit-cost optimization strategy would lead to different operating-point recommendations for the same material scenario, and that the choice between these two strategies is itself a decision variable in the management of material variability.

Table 15.Cross-case synthesis: per-scenario optimal cutting speed and minimum cost.

Scenario	Production-optimal v_{ratio} (Case B)	Production-optimal output (parts/shift)	Cost-optimal v_{ratio} (Case C)	Minimum unit cost (EUR/part)	Unavoidable cost penalty (%)
S0 baseline	1.05	39.80	0.85	80.76	0.00
S1 CE/matrix drift	1.05	36.83	0.90	85.02	5.27
S2 local heterogeneity	0.95	38.84	0.90	81.79	1.28
S3 residual hardening	0.85	33.91	0.80	89.21	10.46
S4 combined adverse	0.85	33.37	0.85	90.44	11.98

The third cross-case finding is that the dominant cost channel through which material variability reaches the economic objective is machine time, not tool consumption or quality loss. The unit-cost decomposition in Section 4.4showed that machine-hour cost and raw-material cost together account for more than 90% of unit cost at the cost-optimal v_{ratio} in every scenario, while tool-change cost contributes approximately 5–7% and scrap cost contribute less than 1%. This decomposition challenges a common framing in which the cost impact of grey cast iron variability is discussed primarily in terms of tool wear and surface-quality incidents. The Case C results indicate that, while material variability does increase the operational frequency of tool changes and incidents, the corresponding direct cost increments are small relative to the unit-cost penalty caused by reduced throughput at the bottleneck operation. Material variability therefore propagates to economics primarily through the indirect channel of machine-time amplification driven by throughput loss. A material condition that increases OP30 wear consumption affects unit cost not only through the direct tool-change cost of additional events, but more substantially through the reduction in OP30 throughput that those events cause. This interpretation has implications both for the engineering discourse around material variability — which often emphasizes tool wear as the primary cost concern — and for the design of supply-chain interventions, which should be evaluated against their effect on bottleneck throughput rather than only against their effect on tool consumption.

The fourth cross-case finding is that the per-part economic value of incoming-material quality improvement is highly scenario-specific and exhibits an approximately eight-fold ratio between the most economically severe and least economically severe mechanisms examined. Residual element accumulation, as represented by scenario S3, justifies a per-part supply-chain investment of up to €8.45, with the sensitivity analysis of Section 4.5 placing this value in the range €6.68 to €11.16 across cost-parameter uncertainty. Local heterogeneity, as represented by scenario S2, justifies a per-part investment of only €1.03. CE/matrix drift, as represented by S1, falls between these two values at €4.26 per part. This asymmetry has direct implications for the prioritization of material-control measures. Supply-chain interventions that target residual element accumulation, such as tighter feedstock screening for Cu, Sn, Cr, Ni, Mo, V and Ti, can be economically justified at substantially higher per-part costs than interventions that target Sulphur, phosphorus or Mn/S balance. This conclusion follows from the cost model presented here and is robust within the parameter ranges examined in the sensitivity analysis. The asymmetry between mechanisms therefore provides a quantitative basis for ranking material-control investments by expected economic return at the production-system level.

These four findings together provide the empirical basis for the discussion in Chapter 5, which interprets the results in relation to material control, machining planning, supply-chain decision-making and the role of cost-based modelling as a bridge between metallurgical variability and production-system economics. The simulation framework does not aim to predict an exact unit cost for a specific factory; rather, it provides a structured methodology to translate material-derived machining changes into a per-part economic indicator that can be evaluated against the cost of supply-chain alternatives.

4.7. Cost increment decomposition by channel

The unit-cost decomposition presented in Section 4.4 (Figure 16) characterises the total unit cost in each scenario at its cost-optimal v_{ratio}. It shows that machine-hour cost and raw-material cost together dominate the total cost in every scenario, with tool and scrap costs contributing smaller shares. While this composition is informative for understanding the overall cost structure, it does not directly identify *through which channel* material variability produces the unavoidable cost penalty relative to the baseline. The total-cost decomposition is dominated by components that are essentially scenario-independent (raw material is constant at €30 per part across all scenarios; machine-hour cost is fully allocated over the same shift duration), and these large but invariant components visually compress the smaller but variable channels through which the material effect actually propagates.

To complement the total-cost view, this section presents the decomposition of the cost *increment* — the unavoidable cost penalty itself — by channel. For each non-baseline scenario, the cost increment is computed at that scenario's own cost-optimal v_{ratio} relative to the S0 baseline at *its* cost-optimal v_{ratio}. The raw-material cost contribution to the increment is identically zero, because the same casting cost is incurred in every scenario regardless of material composition; the increment therefore reduces to three channels: machine-hour, tool-change and scrap. Figure 17 presents the result, with the absolute EUR increment shown on the left and the relative composition shown on the right.

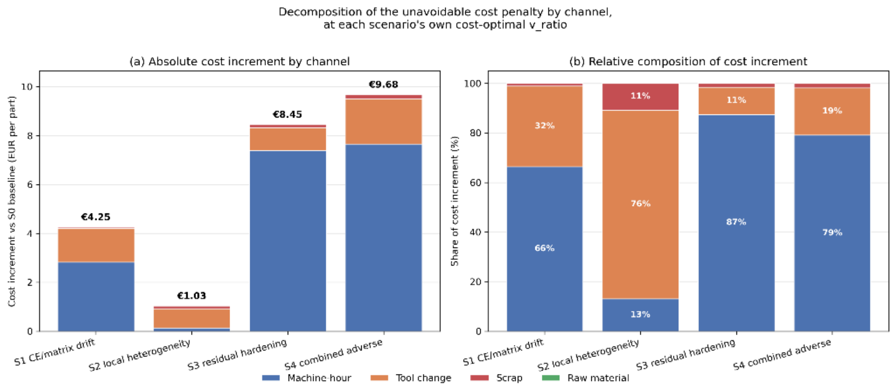

Figure 17. Decomposition of the unavoidable cost penalty by channel, at each scenario's own cost-optimal v_{ratio}. Panel (a) shows the absolute EUR increment relative to S0 baseline, with the total penalty labelled above each bar. Panel (b) shows the same data normalised to 100% to display the relative composition of the increment.

Figure 17 reveals that the three non-baseline mechanisms produce qualitatively different cost-increment signatures, despite all being expressed through the same modelling framework. Three distinct patterns emerge.

For S3 residual hardening, the cost increment is dominated by the machine-hour channel, which accounts for approximately 87% of the €8.45 per-part penalty. The tool-change channel contributes approximately 11%, and the scrap channel contributes less than 2%. This pattern confirms the interpretation developed in Sections 4.4 and 5.3: residual hardening reaches the cost objective primarily through machine-time amplification driven by throughput loss at OP30. The direct tool-change cost increase, although operationally visible as more frequent tool changes, contributes only a small fraction of the total economic penalty. The same pattern, slightly weakened, applies to S4 combined adverse, where the machine-hour channel accounts for approximately 79% and the tool-change channel for approximately 19% of the €9.68 per-part penalty.

For S2 local heterogeneity, the cost-increment signature is qualitatively different. The tool-change channel contributes approximately 76% of the €1.03 per-part penalty, the scrap channel contributes approximately 11%, and the machine-hour channel contributes only 13%. This composition is consistent with the underlying mechanism: local heterogeneity acts primarily through elevated incident intensity, which produces additional tool-change events and occasional scrap, but does not substantially reduce mean throughput once v_{ratio} has been re-optimized. The total penalty is small in absolute terms, but its composition is dominated by quality- and tool-related channels rather than by capacity loss. The S2 signature is therefore not a smaller version of the S3 signature; it is a structurally different cost pattern that reflects a different underlying material mechanism.

For S1 CE/matrix drift, the cost increment is approximately balanced between the machine-hour channel (66%) and the tool-change channel (32%), with negligible scrap contribution. This intermediate pattern is consistent with S1 representing a continuous batch-level shift that affects processing time and wear consumption moderately and uniformly across operations, without the sharp event-driven character of S2 or the dominant wear concentration of S3.

The decomposition of the cost increment therefore provides an additional layer of interpretive resolution beyond the total-cost decomposition. Two scenarios with similar total cost would be indistinguishable from a unit-cost perspective alone, but their cost-increment signatures may reveal that they reach the cost objective through entirely different production-system mechanisms. This distinction is operationally significant: a cost penalty dominated by the machine-hour channel calls for interventions that target throughput protection at the bottleneck operation, whereas a cost penalty dominated by the tool-change or scrap channel calls for interventions that target tool monitoring, quality inspection or supply-side reduction of local-defect drivers. Figure 17 makes this distinction visible at the per-scenario level and provides the empirical basis for the differentiated supply-chain recommendations summarised in Table 17 of Section 5.3.

5. Discussion

5.1. Interpretation of the material-to-production modelling framework

This study aimed to develop a structured modelling framework that connects material variability in grey cast iron with machining response and line-level production consequences. The simulation results should therefore be interpreted as scenario-based evidence of propagation mechanisms, rather than as direct production forecasts.

The modelling framework contains three main layers. The first layer translates representative compositional scenarios into reduced material-response descriptors. These descriptors do not attempt to describe the complete metallurgical state of the cast iron. Instead, they capture tendencies that are relevant for machining, such as carbon-equivalent drift, Mn/S imbalance, residual hardening tendency and local heterogeneity risk. This reduction is necessary because the available information does not support a full microstructure-resolved prediction model. However, it also provides a transparent way to connect compositional changes with machining-relevant behaviour.

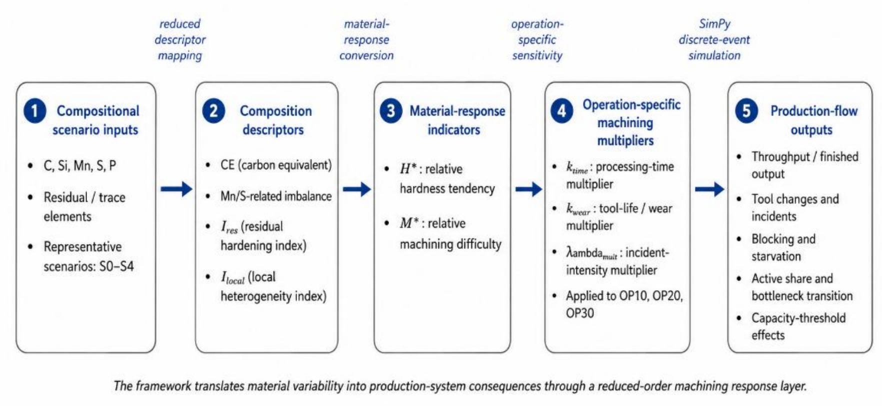

Figure 18.Material-to-production propagation framework

Figure 18 summarizes this propagation structure and clarifies how material-related uncertainty is translated into production-system indicators.

The second layer translates the material-response descriptors into operation-specific machining multipliers. These multipliers affect processing time, tool-life consumption and incident intensity. This step is important because the same material condition does not affect all machining operations in the same way. A batch-level matrix drift may influence general machinability across several operations, while residual hardening may be more strongly reflected in tool wear at a heavy finishing operation. Local heterogeneity may have a smaller effect on average processing time but a stronger effect on stochastic incidents. The model therefore avoids assigning a single material penalty to the whole line. Instead, material effects are propagated through operation-specific sensitivities.

The third layer propagates these machining-level changes through a discrete-event production-flow model. This makes it possible to observe not only finished output, but also machine-state behaviour, tool-change burden, incident occurrence, blocking and starvation. These system-level outputs are important because material variability does not necessarily appear as a local machining problem alone; it can amplify or attenuate as it propagates through buffers and downstream operations, depending on which operation is the constraining one and how its tool life and incident intensity respond to the material condition.

The fourth layer translates these production-flow effects into a per-part unit-cost indicator through a transparent cost-accounting layer covering machine-hour cost, tool-change cost, raw-material cost and probabilistic scrap cost. The unit-cost layer plays a specific methodological role: it does not change the simulation dynamics, but it aggregates the production-flow outputs into a single economic quantity that can be compared between material scenarios and against the cost of supply-chain alternatives. The cost layer therefore acts as the analytical bridge between production-system observation and supply-chain decision-making.

The three simulation cases demonstrate that the consequence of material variability is mediated by the operating-point of the line. Case A characterises throughput response under nominal cutting parameters; Case B introduces the OP30 cutting speed as a decision variable through a normalised ratio coupled to tool life via Taylor's equation, and identifies a production-optimal cutting speed for each material scenario; Case C extends the same sweep with the unit-cost layer and identifies a cost-optimal cutting speed and an unavoidable cost penalty for each scenario. The interpretation that emerges across the three cases is central to the thesis: material variability should not be understood only as a factor that increases cutting time or reduces tool life at one operation. Its practical economic consequence depends on the cutting-speed operating point at which the bottleneck operation is run, and on whether process-side re-optimization has been applied for the material condition.

The most important implication of the modelling framework is therefore that material variability can be treated as a production-system and economic input. In conventional planning, material properties are often considered mainly through quality specifications, machining recommendations or tool-life assumptions. The present framework suggests that material-related differences also influence the cost-optimal operating point of the line and the magnitude of the unavoidable per-part cost penalty that no amount of process-side adjustment can fully neutralise. This does not mean that the model replaces experimental machinability testing or facility-specific cost analysis. It provides a structuring framework in which material-control investments, process-side adjustments and production-flow consequences can be evaluated together on a common economic scale.

For this reason, the numerical results should be read comparatively. The absolute values of output, tool-change counts and unit cost depend on the assumptions used in the model. The more important findings are the relative differences between scenarios, the location of the cost-optimal cutting speed for each scenario and the magnitude of the unavoidable cost penalty that persists after per-scenario re-optimization. These results provide a basis for discussing how grey cast iron variability may influence machining planning, supply-chain decision-making and the cross-functional dialogue between metallurgy, production and procurement.

5.2. Material-specific interpretation of the simulation results

The four non-baseline material scenarios represent different types of material-derived machining difficulty. Their simulation results should therefore be interpreted not only as numerical output differences, but also as different production mechanisms. In this sense, the scenarios are not interchangeable disturbances. Each scenario represents a different way in which grey cast iron variability may enter the machining system.

Table 16.Material-specific interpretation of simulation scenarios

Scenario	Material interpretation	Dominant model signal	Production-system signature	Unavoidable cost penalty	Per-part intervention value (EUR)
S1 CE/matrix drift	Batch-level CE/matrix-related machinability drift	Higher I_{drift}; moderate k_{time}, k_{wear}	Moderate, systematic output loss	+5.27%	4.26
S2 local heterogeneity	Local S/P/Mn-S heterogeneity risk	Higher I_{local} higher λ_{mult}	Smaller mean output loss, higher incident burden	+1.28%	1.03
S3 residual hardening	Residual-element hardening / hard-phase tendency	Higher I_{res}, high OP30 k_{wear}	Large output loss, large cost penalty	+10.46%	8.45
S4 combined adverse	Combined stress case	High combined descriptor and multiplier values	Strongest system stress	+11.98%	9.68

As shown in Table 16, different simulation scenarios can be interpreted to different material variabilities. S1 CE/matrix drift represents a batch-level shift in the carbon-equivalent and matrix-related tendency of the material. In the model, this scenario produces a moderate but systematic reduction in output. This is consistent with the interpretation of S1 as a continuous machinability drift rather than a sudden disturbance mechanism. Its effect is distributed across the line through processing-time and wear multipliers, and it does not mainly appear as an incident-driven scenario.

The four non-baseline material scenarios represent different types of material-derived machining difficulty. Their simulation results should therefore be interpreted not only as numerical output differences, but also in terms of the underlying material mechanisms that they represent and the production-system response they evoke. The unit-cost results of Case C and the sensitivity analysis of Section 4.5 add a second layer of interpretation: each scenario can be characterised both by its production-system signature and by the per-part economic value of any supply-chain intervention that would reduce the corresponding material mechanism toward the S0 baseline state.

S1 CE/matrix drift. S1 represents a batch-level shift in the carbon-equivalent and matrix-related composition of the casting. The simulation framework treats this shift as a continuous drift that propagates through the I_{drift} descriptor and reaches the machining layer primarily through a moderate increase in effective cutting time. The corresponding cost effect is moderate: the unavoidable cost penalty is +5.27%, or €4.26 per part. The relatively modest economic consequence reflects the nature of S1 as a batch-level drift rather than a local-feature disturbance: process planners can to some extent compensate for batch-level drift through standard feed and depth-of-cut adjustments, and the residual penalty after v_{ratio} re-optimization represents the portion of the drift that cannot be fully neutralised by process control alone. The supply-chain implication is that batch-level CE-management measures, such as tighter foundry control of carbon and silicon levels, carry an economic value of approximately €4 per part at the production-system level.

S2 local heterogeneity. S2 represents local disturbances associated with sulphur effects, Mn/S imbalance, phosphorus-related brittleness or local hard spots. In the simulation, S2 acts primarily through the incident-intensity multiplier λ, producing short stochastic disturbances rather than continuous capacity reduction. The unavoidable cost penalty of S2 is the smallest of the four non-baseline scenarios at +1.28%, or €1.03 per part. This is a notable result because it modifies a common framing in which local heterogeneity is treated as a major source of machining difficulty. The simulation framework indicates that, under the present cost structure and within the parameter ranges examined, the average economic consequence of local heterogeneity at the production-system level is small once v_{ratio} has been re-optimized. Local heterogeneity remains an event-sensitive material state with significant operational consequences in terms of incident frequency, but the average per-part cost effect after process re-optimization is limited.

At the microstructural level, the mechanism corresponding to this signature is the local formation of hard-phase features such as manganese sulphide inclusions, phosphide-eutectic islands and local pearlite hardening, which act as point-wise disturbances along the cutting path. These features produce short, intense interactions with the cutting edge that are well represented in the present model as stochastic incidents through the λ multiplier, rather than as a uniform increase in cutting resistance through the k_{time} multiplier. The mapping of microstructural point-wise features onto a stochastic-incident channel in the production-system model is therefore consistent with the underlying metallurgical mechanism. The supply-chain implication is that broad-scope investments targeting sulphur, phosphorus or Mn/S balance can be economically justified up to approximately €1 per part at the production-system level. This figure is intentionally narrow: targeted investments addressing specific high-impact scrap pathways may carry a higher value and would warrant separate analysis with a more detailed quality-cost model.

S3 residual hardening. S3 represents the accumulation of residual and carbide-forming elements (Cu, Sn, Cr, Ni, Mo, V and Ti) and their contribution to pearlite stabilisation, hard-phase tendency and tool-life consumption at OP30. The simulation framework propagates S3 through both k_{wear} and a shift of the production-optimal cutting speed toward lower values. The resulting unavoidable cost penalty is the largest of the single-mechanism scenarios at +10.46%, or €8.45 per part. The sensitivity analysis of Section 4.5 places this value in the range €6.68 to €11.16 per part across cost-parameter uncertainty. This is the principal economic finding of the simulation framework. The mechanism through which S3 reaches the cost objective is consistent with experimental observations in the literature: Dias and Diniz [31] reported that grey cast irons within the same GG25 specification but with different pearlitic and ferritic concentrations exhibit substantially different milling tool lives, with the more pearlitic alloy producing faster tool wear. The simulation result that residual hardening produces a substantially larger production-system impact than local heterogeneity is therefore consistent with experimental machinability evidence. The underlying microstructural mechanism is the stabilisation of a pearlitic matrix and the suppression of ferrite formation by residual and carbide-forming elements such as copper, tin, chromium, nickel, molybdenum, vanadium and titanium. Pearlite is harder and more abrasive to cutting tools than ferrite, and the resulting increase in flank wear rate shortens the Taylor tool life of the cutting insert. In the present model this microstructural effect appears as an elevated k_{wear} multiplier at OP30, which propagates through the Taylor-equation coupling of Section 3.7 and produces both a leftward shift of the production-optimal cutting speed and an upward shift of the unit-cost curve. The connection between residual-element accumulation and pearlite stabilisation is well-established in the grey cast iron literature [28]; the present simulation framework translates this metallurgical mechanism into a production-system economic indicator.

The supply-chain implication is that tighter feedstock screening for residual elements, including incoming scrap charge control and selective alloy use, carries a per-part economic value of approximately €8 per part at the production-system level — substantially higher than the corresponding value for local-heterogeneity control.

S4 combined adverse. S4 is an upper-bound stress scenario rather than a typical industrial condition. It combines the adverse tendencies of S1, S2 and S3 and produces an unavoidable cost penalty of +11.98%, or €9.68 per part, with a sensitivity-range of €6.97 to €13.79 per part. The S4 penalty is only modestly larger than the S3 penalty alone, indicating that residual hardening accounts for the major share of the combined penalty and that the additional contributions of CE drift and local heterogeneity are partially absorbed once v_{ratio} has been re-optimized. From a supply-chain perspective, this suggests that combined adverse material conditions do not produce a strictly additive cost penalty across mechanisms; mitigation efforts focused on the dominant mechanism (residual hardening) capture most of the avoidable cost.

Taken together, the scenario-specific results indicate that different material mechanisms produce economically distinguishable production-system signatures, with an approximately eight-fold ratio between the most economically severe single mechanism (S3 residual hardening at €8.45 per part) and the least economically severe (S2 local heterogeneity at €1.03 per part). This ranking provides a quantitative basis for prioritising supply-chain interventions: investments in residual-element control offer substantially higher per-part economic returns than equivalent investments in sulphur, phosphorus or Mn/S balance under the cost structure examined. Treating all material deviations as a single generic "machinability reduction" would conflate these mechanisms and obscure the economic differences between them. The simulation framework presented in this thesis is therefore most useful not as a tool for predicting absolute unit costs, but as a structured framework for ranking material-control alternatives by their expected economic return at the production-system level.

5.3. Production-system implications and supply-chain decision framework

The production-system implication of the simulation results is that material variability affects the line not only through reduced throughput but through a coupled set of effects that reach the unit-cost objective through different channels and at different magnitudes. Three interpretive themes emerge from the joint reading of Cases A, B and C: the limited role of process-side re-optimization, the dominance of the indirect machine-time cost channel, and the asymmetric value of supply-chain interventions across material mechanisms.

The first implication is that v_{ratio} re-optimization, although available as a process-engineering response to material variability, is not the principal lever for mitigating its cost consequences. The cost-optimal v_{ratio} varies only between 0.80 and 0.90 across the four non-baseline scenarios, and the unit-cost curves are relatively flat near their minima. A single conservative cutting-speed setting therefore captures most of the achievable cost reduction from process-side optimization. The remaining unavoidable cost penalty, which ranges from +1.28% for S2 to +11.98% for S4, persists regardless of process-side adjustments. This result extends the classical machining-economics view, in which optimal cutting speed is treated primarily as a tool–material interaction parameter [32], by showing that batch-to-batch material variability creates a residual cost component that cannot be neutralised by process-side optimization alone.

The second implication is that the dominant channel through which material variability reaches the cost objective is machine-time amplification, not direct tool consumption or quality loss. The unit-cost decomposition of Section 4.4 shows that machine-hour cost and raw-material cost together account for more than 90% of unit cost at the cost-optimal v_{ratio}, while tool-change cost contributes 5–7% and scrap cost less than 1%. This decomposition has interpretive implications for the existing literature on grey cast iron machinability. The grey cast iron machining literature has traditionally emphasised tool wear and surface integrity as the primary engineering concerns associated with material variability [21], [31]. The present results do not contradict this emphasis at the cutting-zone level, but indicate that the corresponding cost consequences at the production-system level reach the economic objective primarily through the indirect channel of throughput loss at the bottleneck operation, rather than through the direct cost of additional tool changes. A material condition that increases OP30 wear consumption affects unit cost not only through the modest direct tool-change cost of additional events, but more substantially through the reduction in OP30 throughput that those events cause. This shifts the appropriate focus of engineering attention from tool-cost minimisation toward throughput protection at the constrained finishing operation.

The third implication concerns the asymmetric value of supply-chain interventions across material mechanisms. The unavoidable cost penalty ratio between the most economically severe single mechanism (S3 residual hardening at €8.45 per part) and the least economically severe (S2 local heterogeneity at €1.03 per part) is approximately eight to one. This asymmetry indicates that supply-chain investments are not equally valuable across material-control measures, and that ranking interventions by per-part economic return at the production-system level provides a quantitative basis for prioritisation. This asymmetry is supported by the cost-increment decomposition presented in Section 4.7, which shows that the two mechanisms reach the cost objective through structurally different channels: S3 through machine-time amplification, S2 through tool-change and scrap channels. Existing production-system literature has examined material-state variability primarily through statistical-process-control and process-capability frameworks [17]; the present framework extends this perspective by translating each material mechanism into a per-part economic value that can be directly compared with the cost of the corresponding supply-chain alternative.

These three implications together support a structured supply-chain decision framework, summarised in Table 17, that links each material mechanism to a recommended response and a per-part value ceiling.

Table 17.Supply-chain decision framework derived from the simulation results.

Dominant incoming material mechanism	**Production-system signature**	**Recommended process-side response**	**Recommended supply-chain response**	**Per-part value ceiling for supply-chain intervention**
S1 CE/matrix drift	Moderate, systematic throughput loss	Modest v_{ratio} adjustment (≈0.90); standard feed-rate compensation	Foundry-level control of carbon and silicon levels; batch-level CE specification	≈ €4 per part
S2 local heterogeneity	Smaller mean throughput loss, higher incident burden	Conservative v_{ratio} (≈0.90); reinforced real-time tool-condition monitoring; quality-gate inspection	Targeted control of S, P and Mn-S balance; selective inclusion screening	≈ €1 per part
S3 residual hardening	Large throughput loss; cost-	Conservative v_{ratio} (≈0.80); planned acceptance of	Tighter feedstock screening for residual elements (Cu, Sn, Cr,	≈ €8 per part (highest priority)

Dominant incoming material mechanism	Production-system signature	Recommended process-side response	Recommended supply-chain response	Per-part value ceiling for supply-chain intervention
	optimal speed shifted to ≈0.80	higher tool-change frequency	Ni, Mo, V, Ti); incoming-scrap charge control; selective alloy use	
S4 combined adverse	Strongest system stress; cost-optimal speed at ≈0.85	Conservative v_{ratio} (≈0.85); combined process-side measures	Combined supply-chain measures targeting residual elements first	≈ €10 per part

The decision framework in Table 17 is the principal industrial-relevance output of the simulation study. It is not intended as a prescriptive recommendation for a specific facility, but as a structured framework that links material mechanisms to economically calibrated supply-chain responses. The per-part value ceilings should be interpreted as upper bounds for the economic justification of the corresponding intervention; the actual investment decision in any specific industrial setting will depend on the cost and feasibility of the intervention itself, on the volume produced and on factors outside the scope of the present model. Within these limits, the framework provides a quantitative starting point for the dialogue between machining-process planners, metallurgy specialists and supply-chain managers that the original consultation interviews of Section 3.2 identified as currently uncommon in everyday production.

5.4. Cross-functional implications and the role of cost-based modelling

The decision framework presented in Section 5.3 has a deliberate cross-functional orientation. Each row of Table 17 links a material mechanism, identified through composition and microstructure, to a process-side response that is the domain of the machining engineer, and to a supply-chain response that is the domain of metallurgy specialists and procurement managers. The per-part value ceiling in the rightmost column places these three perspectives on a common economic scale. This structure addresses an observation that emerged from the industrial practitioner consultation described in Section 3.2: in many industrial settings, decisions about material specification, cutting parameters, tool selection, line layout and incoming-quality control are made within separate functional teams that rarely share a common quantitative basis for trade-off analysis.

The simulation framework is intended to operate as a structuring layer between these teams rather than as a prediction tool for any one of them in isolation. The framework does not replace foundry-level metallurgical control, machining process planning or tool-life testing. It provides a translation between them. A metallurgist analysing residual-element control measures and a process planner adjusting OP30 cutting parameters can both be reasoning about the same per-part economic indicator, computed under transparent and shared assumptions. A supply-chain manager evaluating the cost of tighter incoming-material specifications has a quantitative reference point against which the cost of the specification can be compared.

The framework also has implications for how material certificates and incoming-batch information are used in everyday production. In many settings, material certificates are treated primarily as compliance documents verifying that the batch falls within accepted specification windows. The simulation results suggest that the same information, when read through the propagation chain of Sections 3.5 to 3.7, carries production-system implications that are economically meaningful at the per-part level. A batch with elevated residual-element content within specification may still produce an unavoidable cost penalty several times larger than a batch with elevated S and P content within specification. Treating both deviations as equally acceptable, because both are within specification, conflates two situations with substantially different economic consequences.

A second implication concerns the role of tool-life planning. Tool life is commonly treated as a fixed planning parameter for a given operation. The simulation framework supports a complementary view in which tool life is conditional on the incoming material state, and the production-system response is itself conditional on tool life. Heavy finishing operations such as OP30, which combine long cycle times, tool-life sensitivity and tight quality requirements, are particularly susceptible to this material-state conditionality. A modest reduction in effective tool life at OP30, when amplified by the bottleneck role of the operation, can produce a disproportionate effect on per-part unit cost.

A third implication concerns the monitoring of local heterogeneity. The cost analysis indicates that S2 produces a small mean cost penalty after process re-optimization, but the underlying production-system signature is dominated by stochastic incidents rather than continuous capacity reduction. Mean-value indicators such as average throughput or average cycle time are therefore not adequate for detecting or managing local heterogeneity. Event-oriented indicators — abnormal tool failures, vibration alarms, quality-gate interventions, short stoppages and rework frequency — are more directly aligned with the production-system signature of this material mechanism. The industrial consultation of Section 3.2 confirmed that such event-oriented indicators are already in use in practice, but typically without a quantitative link to incoming-material characteristics.

The role of cost-based modelling in this study is therefore to provide a common quantitative language for these cross-functional decisions. The unit-cost layer of Section 3.9 is intentionally simple and transparent; it is not a procurement-cost prediction tool for industrial implementation. Its value lies in making the trade-off between material-control investment and process-side response explicit and economically calibrated. Within the boundaries of the simulation framework, this translation is the principal industrial-relevance contribution of the thesis.

5.5. Model assumptions and interpretation boundaries

The simulation results should be interpreted within the assumptions and boundaries of the modelling framework. The purpose of the model is to examine how material-derived machining changes may propagate through a constrained production line and reach an economic indicator. The model is not intended to replicate a specific industrial facility, to provide absolute predictions of throughput or unit cost, or to substitute for experimental cutting trials. This section identifies the principal assumptions and limitations of the framework and discusses their implications for the interpretation of the numerical results.

The first limitation concerns the material scenarios. The compositional scenarios used in this thesis are representative and literature-supported, but they are not measured batch chemistries from a specific foundry. The scenarios are intended to capture distinguishable material conditions rather than to predict the actual frequency or severity of any specific industrial batch state. The scenario S4 in particular is an upper-bound stress case rather than a typical industrial condition.

The second limitation concerns the reduced material descriptors. The descriptors used in the model — carbon equivalent, Mn/S-related imbalance, residual hardening index and local heterogeneity index — are reduced-order representations of complex metallurgical phenomena. They aggregate effects that, in a fully physical model, would require explicit description of graphite morphology, matrix microstructure, inclusion populations and segregation patterns. The descriptors used here are sufficient to distinguish between the four scenarios examined, but they would require refinement and calibration against measured microstructural data before any quantitative application to specific industrial batches.

The third limitation concerns the conversion from material descriptors to machining multipliers. The processing-time multiplier k_{time}, the wear multiplier k_{wear} and the incident-intensity multiplier λ_{mult} are modelling parameters derived from logical reasoning about the relative sensitivity of each operation to each material mechanism. They are not fitted to experimental cutting trials. The numerical values used in the simulation should therefore be read as logically consistent rather than as empirically calibrated. A related point concerns Taylor's tool-life equation, which is used in Section 3.7 both as a conceptual link between material response and tool life and as the operational coupling between OP30 cutting speed and OP30 tool life in Case B and Case C. The Taylor exponent n is assigned a single value (0.25) consistent with the cemented-carbide-on-grey-cast-iron literature range, but is treated as fixed across all material scenarios. In practice, the wear sensitivity to cutting speed may vary with material condition; treating n as fixed therefore tends to underestimate the speed sensitivity of tool life under the more adverse scenarios and is conservative for the principal cost conclusions.

A related point concerns the robustness of the quantitative conclusions to the cost-parameter assumptions used in Section 3.9. The cost layer relies on assigned values for machine-hour rates, tool-change costs, raw-material cost and the scrap conversion rate, none of which is calibrated from a specific industrial facility. The sensitivity analysis of Section 4.5 examined this concern by sweeping the scrap conversion rate over the range 0.10 to 0.50, the raw-material cost over €20 to €40 per part and the OP30 tool-change cost over €10 to €30 per event. Across the 27 resulting parameter combinations, the unavoidable cost penalty of S3 residual hardening relative to S0 baseline remained within the range +9.05% to +12.42%, and the corresponding penalty for S4 combined adverse remained within +10.15% to +15.23%. The qualitative conclusion that residual hardening produces an unavoidable cost penalty on the order of 10% and that combined adverse conditions produce an unavoidable cost penalty on the order of 12% is therefore robust to wide variation in the cost-parameter assumptions. The exact percentage shifts with the assigned parameters, but the order of magnitude does not. The per-part economic value of an incoming-material quality improvement targeting residual elements correspondingly remains in the range €6.68 to €11.16 across the swept parameter combinations, with a central estimate of €8.45.

The fourth limitation concerns the discrete-event production model. The SimPy model represents the production line as a simplified serial system with OP10, OP20, OP30 and finite buffers. This structure is sufficient to examine the propagation of material-derived disturbances through a constrained line, but it does not represent the full complexity of an industrial machining cell. Real cells may include parallel resources, alternative routings, shared tooling, operator-shared stations, transport delays, planned preventive maintenance, shift-change effects and inter-cell coupling. These features were judged outside the scope of the present study, which focuses on the propagation mechanism rather than on facility-specific dynamics.

The fifth limitation concerns the modelling of quality consequences. In the present simulation, local heterogeneity is represented primarily through incident intensity. This captures the idea of sudden disturbances but does not explicitly model scrap, rework, dimensional deviation or surface integrity loss. The cost layer compensates for this in part by treating a fraction of OP30 incidents as scrap events, but it does not represent quality consequences in detail. The S2 local heterogeneity scenario in particular may carry industrial costs through quality channels that are not fully captured by the present model. A more detailed quality-cost model would refine the S2 penalty estimate and could plausibly raise it above the value of €1.03 per part reported in Section 4.4.

The sixth limitation concerns the cost layer itself. The cost layer is intentionally simple and transparent, using one rate per machine, one cost per tool-change event, one price per casting and one scalar conversion rate for scrap. More elaborate cost structures — time-of-day rates, indexable insert tracking, batch-dependent casting prices, detailed quality-cost categories, opportunity cost of capacity loss, energy cost and labour allocation — are not modelled. The unit-cost values reported in Section 4.4 are therefore reference figures suitable for comparative analysis between material scenarios, not procurement-cost predictions for any specific facility. The sensitivity analysis of Section 4.5 demonstrates that the comparative cost conclusions are robust to variation in the cost-parameter values, but it does not address the structural simplification of the cost layer itself.

The seventh limitation concerns the use of Monte Carlo replications. The simulation uses 200 replications per (scenario, v_{ratio}) combination, with shorter sensitivity sweeps using equivalent sampling. This is sufficient to obtain stable mean values and clear comparative orderings between scenarios, but it does not exhaustively characterise the tail behaviour of the production output distribution. Rare-event consequences, such as catastrophic tool failure or compound disturbance sequences, are not specifically addressed.

The industrial practitioner consultation described in Section 3.2 should also be interpreted within its limited scope. It involved a single practitioner and was used to support practical grounding and plausibility interpretation rather than to provide statistical evidence about industrial practice. The consultation contributes to the framing of the study but does not constitute a structured empirical investigation of industrial machining behaviour.

These assumptions and limitations do not invalidate the modelling framework. They define its intended use. The model is best understood as an exploratory and explanatory tool, suitable for testing how a defined material change may propagate through a constrained production line and what magnitude of economic consequence may result. The principal interpretation boundary is the distinction between mechanistic insight and industrial prediction: the thesis provides mechanistic insight into how grey cast iron variability can influence production-system economics, but it does not provide industrial predictions for any specific facility. Within this boundary, the modelling results identify which relationships should be investigated further, which production indicators may be sensitive to material variability and which supply-chain interventions carry the highest expected economic return at the production-system level.

5.6. Future work

The assumptions discussed above also define the principal directions for future work. The present thesis develops a reduced-order material-to-production modelling framework with a cost layer, and demonstrates its behaviour through representative scenarios with structured sensitivity analysis. A natural next step is to replace representative assumptions with measured material, machining, production and cost data.

The first direction is to use real batch-level grey cast iron composition data. The compositional scenarios used in this thesis were constructed to represent different material-response mechanisms. Future work could collect chemical composition data from multiple industrial batches and calculate the same descriptors used in this study, including carbon equivalent, Mn/S-related imbalance, residual hardening index and local heterogeneity index. This would make it possible to examine whether the scenario space used in the thesis corresponds to actual industrial batch variation and to characterise the empirical distribution of each material mechanism in a specific foundry-to-machining supply chain.

The second direction is to connect the compositional descriptors with measured material properties and microstructural observations. Hardness, tensile strength, matrix fraction, graphite morphology and carbide tendency could be measured for the same batches. This would allow the reduced descriptors to be evaluated against physical material responses rather than treated only as modelling indicators. Such validation would be especially important for the residual hardening and local heterogeneity indices, because these mechanisms involve microstructural features that composition alone cannot fully characterise.

The third direction is to calibrate the machining multipliers through cutting trials. In the current model, the processing-time multiplier, the wear multiplier and the incident-intensity multiplier are scenario-sensitivity parameters derived from logical reasoning rather than from experimental cutting data. Future experiments could measure cutting forces, tool wear, surface quality, vibration and tool-life variation for different grey cast iron batches and use the results to estimate operation-specific relationships between material descriptors and machining response. In particular, the tool-life response of OP30-like finishing operations should be investigated experimentally, because the simulation results indicate that finishing tool-life consumption is the principal mechanism through which residual hardening reaches the unit-cost objective. Within the same direction, the Taylor exponent n could be estimated for each material condition rather than treated as a fixed value, removing one of the conservative simplifications of the present cost analysis.

The fourth direction is to validate the production-flow model using shop-floor data. Machine-state logs, tool-change records, quality alarms, production counts, cycle-time records and buffer or work-in-process observations could be compared with simulation outputs for the same production line. This would test the model not only at the machining level, but also at the production-system level. If a material batch shows elevated residual-hardening tendency in its descriptors, future work could examine whether the production line actually experiences increased tool-change frequency at OP30, reduced effective throughput at OP30, and the magnitude of unit-cost increase that the cost layer would predict for the observed material condition.

The fifth direction is to expand the production model to include quality-related losses in more detail. The current cost layer represents quality consequences through a single scalar scrap conversion rate applied to OP30 incidents. Real local heterogeneity may produce a richer set of quality consequences, including scrap, rework, inspection delays, surface defects and dimensional deviations, distributed across multiple operations. Adding quality loops, inspection delays and explicit rework paths to the simulation, together with calibrated quality-cost categories, would refine the S2 cost estimate and could plausibly raise it above the value of €1.03 per part reported in this thesis. The S2 result of the present model should therefore be interpreted as a lower bound on the economic consequence of local heterogeneity, pending such an extension.

The sixth direction is to refine the cost layer with industrial cost data. The present cost layer uses commercial list prices for tooling, market-level reference values for raw material and assigned values for machine-hour rates and scrap conversion. Future work could replace these with facility-specific cost data including time-of-day rates, indexable insert tracking, batch-dependent casting prices, detailed quality-cost categories, opportunity cost of capacity loss, energy cost and labour allocation. This would convert the cost layer from a comparative analytical tool into a facility-specific decision-support tool. The sensitivity analysis of Section 4.5 indicates that the qualitative conclusions of the present framework are robust to wide variation in the assigned cost parameters, so the refinement would primarily improve absolute accuracy rather than overturn the comparative ranking of material mechanisms.

The seventh direction is to connect the framework with production planning and scheduling. If material batches can be classified according to their expected mechanism profile through their composition descriptors, this classification could be used in scheduling decisions. Batches with elevated residual-hardening tendency could be scheduled during periods of higher machining-cost tolerance or with reinforced tool-change planning. Batches with elevated local heterogeneity could be scheduled during periods of enhanced incident-monitoring coverage. This would convert material characterisation from a compliance-verification activity into an active production-planning input, complementing the supply-chain decision framework of Section 5.3 with a scheduling-layer extension.

An eighth direction is to extend the exploratory practitioner consultation of Section 3.2 into a more systematic interview study involving material engineers, machining engineers, production planners, quality engineers and supply-chain managers. The decision framework of Table 17 in Section 5.3 is in principle testable against industrial practice: whether per-part value ceilings of the magnitudes derived here correspond to investment thresholds that industrial decision-makers consider justified would be a useful empirical follow-up.

A final direction is to develop the modelling chain into an industrial decision-support workflow. Such a workflow could start from measured batch composition, calculate material-response descriptors, estimate machining multipliers, simulate production-flow consequences, apply a facility-specific cost layer and report expected per-part cost penalties together with the corresponding supply-chain value ceilings. This would require the calibration steps listed in directions one through six, but the structure developed in the present thesis provides a starting point for such a tool.

Overall, future work should aim to move from a scenario-based explanatory model toward a calibrated industrial decision-support model. The most important requirement is not additional simulation complexity, but stronger empirical connection between material data, cutting response, production records and facility cost structures. If these connections can be established, grey cast iron material variability could be incorporated more systematically into machining planning, capacity analysis, production scheduling and supply-chain decision-making, with explicit per-part economic accounting at each step.

5.7. Chapter summary

This chapter discussed the meaning, industrial relevance and limitations of the simulation results presented in Chapter 4. The main purpose was to move from numerical results to interpretation. The material-to-production framework developed in this thesis should be understood as an explanatory and structuring modelling chain rather than as a direct industrial prediction model. Its value lies in showing how representative grey cast iron material scenarios can be translated, first into material-response descriptors, then into operation-specific machining multipliers, then into discrete-event production-flow consequences, and finally into a per-part economic indicator suitable for evaluating supply-chain alternatives.

The discussion first interpreted the simulation results scenario by scenario. CE/matrix drift behaves as a batch-level shift in machining difficulty that partially admits process-side compensation through cutting-speed re-optimization, leaving an unavoidable cost penalty of approximately +5% relative to baseline. Local heterogeneity behaves as a stochastic disturbance mechanism, dominated by incident intensity rather than by continuous capacity reduction; its average cost penalty is the smallest of the non-baseline scenarios at approximately +1%, but its cost-increment composition is dominated by tool-change and scrap channels rather than by machine-time amplification, indicating a structurally distinct cost pattern from the other mechanisms. Residual hardening produces the largest single-mechanism cost penalty at approximately +10%, driven by elevated wear consumption at the heavy finishing operation and the corresponding leftward shift of both the production-optimal and cost-optimal cutting speeds. The combined adverse scenario reaches approximately +12%, only modestly above the residual-hardening alone, indicating that residual hardening accounts for the major share of the combined penalty.

The chapter then drew three production-system implications from the joint reading of Cases A, B and C. First, process-side re-optimization of OP30 cutting speed is not the principal lever for mitigating the cost consequences of material variability; the cost-optimal cutting speeds for all scenarios lie in the narrow range 0.80–0.90, and the unit-cost curves are relatively flat near their minima. Second, the dominant channel through which material variability reaches the cost objective is machine-time amplification driven by throughput loss at the bottleneck operation, not direct tool consumption or quality loss. The unit-cost decomposition shows that machine-hour cost and raw-material cost together account for more than 90% of unit cost in all scenarios. Third, the per-part economic value of supply-chain interventions is highly asymmetric across material mechanisms, with an approximately eight-fold ratio between the value of residual-element control (€8.45 per part) and the value of local-heterogeneity control (€1.03 per part).

The chapter then presented a supply-chain decision framework that links each material mechanism to a recommended process-side response, a recommended supply-chain response, and a per-part value ceiling for the corresponding intervention. The framework is intended to operate as a structuring layer between metallurgy specialists, machining engineers, supply-chain managers and production planners, providing a common economic reference point for cross-functional trade-off discussions. The unit-cost values produced by the framework are reference figures suitable for comparative analysis between material scenarios rather than procurement-cost predictions for any specific facility.

The chapter also clarified the boundaries of the model. The compositional scenarios are representative rather than measured industrial batches. The material descriptors and machining multipliers are reduced-order modelling tools rather than experimentally calibrated metallurgical or cutting models.

The cost layer is intentionally simple and transparent, and the simulation excludes several factors that may be important in specific facilities, including detailed quality consequences, parallel resources, alternative routings, preventive maintenance and shift-change effects. The numerical results should therefore be interpreted comparatively and mechanistically, not as direct predictions for a specific factory. Within these boundaries, the sensitivity analysis of Section 4.5 confirms that the principal cost conclusions are robust to wide variation in the assigned cost-parameter values.

Finally, the chapter outlined future work needed to develop the framework further. The most important next steps are to use measured batch composition data, validate material descriptors against microstructural observations, calibrate the machining multipliers through experimental cutting trials, refine the cost layer with facility-specific cost data, and compare production-flow simulation outputs with shop-floor records. With such calibration, the framework could evolve from a scenario-based explanatory model into a calibrated industrial decision-support tool.

Overall, the discussion supports the central argument of the thesis: grey cast iron material variability is not adequately characterised as a single generic machinability loss. It propagates through distinct mechanisms with substantially different production-system signatures and substantially different per-part economic consequences. Through its effects on tool life, incident frequency and bottleneck-operation throughput, material variability creates an unavoidable cost penalty whose magnitude and distribution across mechanisms can be quantified and used as a basis for prioritising supply-chain interventions.

6. Conclusion

6.1. Main conclusions

This thesis investigated how material variability in grey cast iron propagates from compositional differences through machining response to production-system behaviour and unit-cost impact. The study developed a reduced-order modelling framework that connects representative compositional scenarios with material-response descriptors, operation-specific machining multipliers, discrete-event production-flow simulation and a transparent cost layer. The purpose was not to predict the exact output or unit cost of a specific industrial line, but to examine how different material-related mechanisms influence throughput, tool consumption, incident frequency and the per-part economic indicator of production, and to translate these effects into a structured basis for supply-chain decision-making.

The first main conclusion is that material variability should not be treated as a single generic machinability loss. The four non-baseline scenarios examined in this thesis produce distinguishable production-system signatures and substantially different per-part economic consequences. CE/matrix drift behaves as a stable batch-level shift in machining difficulty and produces an unavoidable cost penalty of approximately +5.27% (€4.26 per part). Local heterogeneity behaves as a stochastic disturbance mechanism with high event burden but limited mean cost impact, producing an unavoidable cost penalty of approximately +1.28% (€1.03 per part). Residual hardening produces a systematic capacity-channel effect at the heavy finishing operation, generating the largest single-mechanism unavoidable cost penalty at +10.46% (€8.45 per part). Combined adverse material conditions reach +11.98% (€9.68 per part). These differences are robust to wide variation in the cost-parameter assumptions, as demonstrated by the sensitivity analysis of Section 4.5.

The second main conclusion is that process-side optimization of cutting speed is not the principal lever for mitigating the cost consequences of material variability. The cost-optimal v_{ratio} for the four non-baseline scenarios lies in the narrow range 0.80–0.90, with relatively flat unit-cost curves near each minimum. A single conservative cutting-speed setting captures most of the achievable cost reduction from process-side response, and the unavoidable cost penalties listed above persist regardless of further cutting-speed adjustment. This result extends the classical machining-economics view of cutting-speed optimization by showing that batch-to-batch material variability creates a residual cost component that cannot be neutralised by process-side optimization alone.

The third main conclusion is that the dominant channel through which material variability reaches the unit-cost objective is machine-time amplification driven by throughput loss at the bottleneck operation, not direct tool consumption or quality loss. The unit-cost decomposition shows that machine-hour cost and raw-material cost together account for more than 90% of unit cost in all scenarios, while tool-change cost contributes 5–7% and scrap cost less than 1%. This identifies throughput protection at the bottleneck operation as the principal engineering focus when material variability is the concern, rather than the more commonly discussed focus on tool-cost minimisation.

The fourth main conclusion is that the per-part economic value of supply-chain interventions is highly asymmetric across material mechanisms. The approximately eight-fold ratio between the value of residual-element control (€8.45 per part) and the value of local-heterogeneity control (€1.03 per part) provides a quantitative basis for prioritising supply-chain investments. Within the cost structure examined and across the sensitivity ranges of Section 4.5, supply-chain measures targeting residual-element accumulation in incoming castings carry substantially higher per-part economic return than supply-chain measures targeting sulphur, phosphorus or Mn/S balance.

Taken together, these conclusions support the central argument of the thesis. Grey cast iron material variability propagates through distinct mechanisms with quantifiably different production-system signatures and quantifiably different per-part economic consequences. The framework developed here translates this distinction into a structured supply-chain decision basis: for each identified material mechanism, the per-part value ceiling for any corresponding incoming-quality improvement is established, and the relative priority of competing supply-chain investments can be ranked on a common economic scale.

6.2. Answer to the research question

The three research questions stated in Chapter 1.3 can be answered through the modelling framework and simulation results developed in this thesis.

Research question 1 asked how structured compositional variability in grey cast iron can be represented through reduced material-response indicators. This thesis addressed this question by constructing a descriptor layer between raw composition and machining response. Instead of treating each element as an isolated input, the model groups compositional tendencies according to their expected material-response mechanisms. Carbon-equivalent-related drift is represented through CE and a drift index. Mn/S imbalance, sulphur and phosphorus contributions are represented through a local heterogeneity index. Residual and carbide-forming elements such as Cu, Sn, Cr, Ni, Mo, V and Ti are represented through a residual hardening index. These descriptors are not intended to predict the complete metallurgical state of the cast iron, but they provide a transparent way to express batch-level drift, local heterogeneity risk and residual hardening tendency in a form that can be propagated into machining and production analysis. The descriptor layer therefore provides a structured intermediate representation in which compositional variability can be characterised by mechanism rather than by raw chemistry alone.

Research question 2 asked how these material-response indicators can be translated into operation-specific machining-response parameters. This thesis addressed this question by introducing three machining-response multipliers: the processing-time multiplier k_{time}, the wear multiplier k_{wear} and the incident-intensity multiplier λ_{mult}. These parameters allow the same material scenario to affect different operations differently. OP10 is treated as more sensitive to general machinability drift, OP20 as more sensitive to localised heterogeneity and incident risk, and OP30 as more sensitive to residual hardening and tool-life burden. This operation-specific translation is essential because material variability does not enter the production line as a single uniform penalty. It affects processing time, tool-life consumption and stochastic incidents through different mechanisms and with different operation-level sensitivities, and the resulting differentiation is the foundation for the production-flow and cost analyses that follow.

Research question 3 asked how the resulting machining-response parameters affect production-flow behaviour in a serial machining line. This thesis addressed this question through three complementary simulation configurations, each making a different decision variable explicit. Case A examined production-flow behaviour under a fixed line structure with nominal cutting parameters, establishing the throughput and event-burden baseline. Case B introduced the OP30 cutting speed as a decision variable through a normalised ratio v_{ratio} and identified the production-optimal v_{ratio} for each material scenario, finding that the production-optimal cutting speed shifts toward lower values under more adverse material conditions (from v_{ratio} = 1.05 for the baseline scenario to v_{ratio} = 0.85 for residual hardening). Case C extended the cutting-speed sweep with a transparent unit-cost layer, identifying the cost-optimal v_{ratio} for each scenario and the unavoidable cost penalty that persists after per-scenario re-optimization. The unavoidable cost penalty ranges from +1.28% for local heterogeneity to +11.98% for combined adverse conditions, with residual hardening alone reaching +10.46%. The sensitivity analysis of Section 4.5 confirms that these magnitudes are robust to wide variation in cost-parameter assumptions. Together, the three cases show that the production-system consequence of material variability extends beyond throughput loss and reaches the unit-cost objective primarily through machine-time amplification at the bottleneck operation, with substantially asymmetric magnitudes across mechanisms.

Taken together, the answers to the three research questions demonstrate that grey cast iron material variability can be linked to production-system consequences and to per-part economic outcomes through an intermediate machining-response layer. The relationship is neither direct nor linear: compositional tendencies first influence material-response descriptors, these descriptors then influence operation-specific machining parameters, the machining changes propagate through the discrete-event production system, and the production-system outputs are aggregated into a per-part economic indicator. The economic consequence depends not only on the material scenario, but on the cost structure of the line and on whether process-side cutting-speed adjustment has been re-optimised for the material condition.

6.3. Contributions of the thesis

The first contribution of this thesis is the development of a transparent four-layer material-to-production modelling chain. The framework connects representative grey cast iron composition scenarios to reduced material descriptors, then to operation-specific machining multipliers, then to discrete-event production-flow simulation outputs, and finally to a per-part economic indicator through a transparent cost layer. The propagation is fully explicit at each layer, and the simplifications used at each layer are documented. This provides a structured way to analyse material variability as a production-system and supply-chain issue rather than only as a metallurgical or machining issue.

The second contribution is the differentiation between distinct material-risk mechanisms. The thesis shows that batch-level drift, local heterogeneity and residual hardening should not be grouped into a single generic machinability factor. They reach the production-system level through different channels and reach the unit-cost objective at substantially different magnitudes. The approximately eight-fold ratio between the per-part economic value of residual-element control and the per-part economic value of local-heterogeneity control, identified within the present cost structure and sensitivity ranges, is a direct expression of this differentiation.

The third contribution is the introduction of a cost-based interpretation of material variability at the production-system level. By coupling the discrete-event production-flow model with a transparent cost layer and a Taylor-equation cutting-speed sweep, the framework translates production-system effects into a per-part economic indicator that can be directly compared with the cost of supply-chain alternatives. The unavoidable cost penalty, defined as the cost increase that persists after per-scenario re-optimization of the cutting speed, identifies the portion of material-variability cost that process-side response alone cannot address. The corresponding per-part value ceiling for incoming-material quality improvement is established for each material mechanism.

The fourth contribution is a supply-chain decision framework that links each material mechanism to a recommended process-side response, a recommended supply-chain response and a per-part value ceiling for the corresponding intervention. The framework provides a common quantitative basis for cross-functional decisions between metallurgy specialists, machining engineers, supply-chain managers and production planners. It is intended as a structuring layer between these functions rather than as a prescriptive recommendation for any specific facility. The principal industrial-relevance output of the thesis is this framework rather than any single numerical value.

A fifth contribution, methodological in nature, is the demonstration that the propagation chain from material composition to per-part unit cost is computationally tractable using freely available open-source tools (Python and SimPy) and a transparent cost-accounting layer. This supports the use of the framework as a basis for future extensions, including facility-specific calibration, refined quality-cost modelling and integration with production-planning systems.

6.4. Limitations and future outlook

The results should be interpreted within the boundaries of the model. The material scenarios are representative rather than measured industrial batches. The material descriptors and machining multipliers are reduced-order modelling tools rather than experimentally calibrated metallurgical or cutting models. The Taylor exponent n used to couple cutting speed with tool life in

Cases B and C is treated as a fixed value across all scenarios, rather than as a material-dependent parameter. The production line is simplified into a serial SimPy model with three operations and finite buffers. The cost layer is intentionally simple and transparent, using commercial list prices for tooling, market-level reference values for raw material, and assigned values for machine-hour rates and scrap conversion. Several factors that may be important in specific facilities — including detailed quality consequences, parallel resources, alternative routings, preventive maintenance, shift-change effects, time-of-day rates, indexable insert tracking, batch-dependent casting prices, and opportunity cost of capacity loss — are not modelled.

The unit-cost values produced by the framework should therefore be interpreted as reference figures suitable for comparative analysis between material scenarios, not as procurement-cost predictions for any specific facility. The sensitivity analysis of Section 4.5 demonstrates that the comparative cost conclusions are robust to wide variation in the assigned cost-parameter values across a 27-combination sweep, with the S3 unavoidable cost penalty consistently in the range +9.05% to +12.42% and the corresponding per-part value of residual-element control in the range €6.68 to €11.16. The qualitative conclusions of the framework are therefore not artefacts of specific cost-parameter choices, although the exact numerical values shift with the assigned parameters.

Future work should focus on calibration and validation. Real batch-level composition data could be used to test whether the proposed material descriptors reflect actual industrial variation. Cutting trials could be used to calibrate the relationships between material state, processing time, tool wear and incident risk, and to estimate the material-dependence of the Taylor exponent that is treated as fixed in the present model. Shop-floor data such as machine-state logs, tool-change records, quality alarms, production counts and unit-cost records could then be used to validate the production-flow and cost layers jointly.

A further development would be to include quality-related losses in more detail, beyond the single-rate scrap conversion of the present cost layer. This would be especially important for evaluating local heterogeneity, whose per-part cost penalty in the present framework is approximately €1.03 but could plausibly increase substantially under a more detailed quality-cost model. Extension of the framework with facility-specific cost data and integration with production-planning and scheduling systems would convert the framework from a comparative analytical tool into a facility-specific decision-support tool, while retaining the cross-functional structuring role identified in this thesis as its principal industrial-relevance contribution.

Data and Code Availability Statement

The complete simulation code, parameter files, raw Monte Carlo replication outputs and aggregated result tables underlying the figures and conclusions of this thesis are publicly available in an open-source GitHub repository at https://github.com/ChenghanQian/Cast-Iron-Machining-Simulation (DOI:10.5281/zenodo.20370260). The repository contains a Jupyter notebook that reproduces all numerical results reported in Chapter 4 and Appendix C, together with a README file documenting the reproducibility procedure, the random-seed convention and the expected wall-clock time. The commit referenced in this thesis is 153ebae08cf9901d15254c1bb52c119f315281fb, accessed 2026-05-25. The code is released under the MIT License. No restricted, proprietary or personally identifiable data are used in this study; all material composition and cost-parameter values are documented within the repository.

References

[1] C. Soffritti, L. Calzolari, S. Pepi, A. Fortini, M. Merlin, and G. L. Garagnani, ‘Metallurgical and Statistical Approaches to the Study of Cast Iron Street Furniture’, Metall. Mater. Trans. A, vol. 52, no. 3, pp. 1127–1141, Mar. 2021, doi: 10.1007/s11661-021-06135-6.

[2] P. Futáš, A. Pribulová, G. Fedorko, and V. Molnár, ‘Influence of Steel Scrap in the Charge on the Properties of Gray Cast Iron’, ISIJ Int., vol. 57, no. 2, pp. 374–379, 2017, doi: 10.2355/isijinternational.ISIJINT-2016-270.

[3] K. Ankamma, ‘Effect of Trace Elements (Boron and Lead) on the Properties of Gray Cast Iron’, *J. Inst. Eng. India Ser. D*, vol. 95, no. 1, pp. 19–26, Apr. 2014, doi: 10.1007/s40033-013-0031-3.

[4] J. Sertucha and J. Lacaze, ‘Casting Defects in Sand-Mold Cast Irons—An Illustrated Review with Emphasis on Spheroidal Graphite Cast Irons’, *Metals*, vol. 12, no. 3, p. 504, Mar. 2022, doi: 10.3390/met12030504.

[5] X. Zhang, G. Ma, M. Liu, and Z. Li, ‘Removal of Residual Element Tin in the Ferrous Metallurgy Process: A Review’, *Metals*, vol. 9, no. 8, p. 834, Jul. 2019, doi: 10.3390/met9080834.

[6] I. Riposan, M. Chisamera, S. Stan, C. Ecob, and D. Wilkinson, ‘Role of Al, Ti, and Zr in Gray Iron Preconditioning/Inoculation’, *J. Mater. Eng. Perform.*, vol. 18, no. 1, pp. 83–87, Feb. 2009, doi: 10.1007/s11665-008-9260-2.

[7] R. Gundlach, M. Meyer, and L. Winardi, ‘Influence of Mn and S on the Properties of Cast Iron Part III—Testing and Analysis’, *Int. J. Met.*, vol. 9, no. 2, pp. 69–82, Apr. 2015, doi: 10.1007/BF03355617.

[8] J. Lacaze, ‘Trace Elements and Graphite Shape Degeneracy in Nodular Graphite Cast Irons’, *Int. J. Met.*, vol. 11, no. 1, pp. 44–51, Jan. 2017, doi: 10.1007/s40962-016-0115-6.

[9] X. Ding, X. Li, Q. Feng, W. Matthias, and S. Huang, ‘Microstructure evolution in grey cast iron during directional solidification’, *Int. J. Miner. Metall. Mater.*, vol. 24, no. 8, pp. 884–890, Aug. 2017, doi: 10.1007/s12613-017-1474-6.

[10] L. F. De Santiago-Méndez, M. D. J. Castro-Román, M. Herrera-Trejo, H. Mancha-Molinar, and B. Bravo, ‘Influence of S and Mn Initial Concentrations on the Graphite Branching in Gray Cast Iron as Quantified by 2D Image Analysis’, *Materials*, vol. 18, no. 21, p. 4837, Oct. 2025, doi: 10.3390/ma18214837.

[11]K. Odum, N. Raymond, D. Pell, and M. Soshi, ‘Surface feature formation mechanism during finish milling of gray cast iron’, *Int. J. Adv. Manuf. Technol.*, vol. 92, no. 1–4, pp. 459–469, Sep. 2017, doi: 10.1007/s00170-017-0162-z.
[12]K. Odum and M. Soshi, ‘Surface Formation Study Using a 3-D Explicit Finite Element Model of Machining of Gray Cast Iron’, *Procedia CIRP*, vol. 45, pp. 111–114, 2016, doi: 10.1016/j.procir.2016.02.168.
[13]F. Schultheiss, V. Bushlya, F. Lenrick, D. Johansson, S. Kristiansson, and J.-E. Ståhl, ‘Tool Wear Mechanisms of pcBN tooling during High-Speed Machining of Gray Cast Iron’, *Procedia CIRP*, vol. 77, pp. 606–609, 2018, doi: 10.1016/j.procir.2018.08.201.
[14]V. Bushlya, R. Lindvall, F. Lenrick, L. Magnusson Åberg, R. M’Saoubi, and J.-E. Ståhl, ‘Effect of ageing on machining performance of grey cast iron and its compensation by cutting speed management’, *CIRP Ann.*, vol. 73, no. 1, pp. 45–48, 2024, doi: 10.1016/j.cirp.2024.04.061.
[15]M. El Rayes, A. Abbas, A. Al-Abduljabbar, A. Ragab, F. Benyahia, and A. Elkaseer, ‘Investigation and Statistical Analysis for Optimizing Surface Roughness, Cutting Forces, Temperature, and Productivity in Turning Grey Cast Iron’, *Metals*, vol. 13, no. 6, p. 1098, Jun. 2023, doi: 10.3390/met13061098.
[16]L. D. C. Ng Corrales, M. P. Lambán, M. E. Hernandez Korner, and J. Royo, ‘Overall Equipment Effectiveness: Systematic Literature Review and Overview of Different Approaches’, *Appl. Sci.*, vol. 10, no. 18, p. 6469, Sep. 2020, doi: 10.3390/app10186469.
[17]M. Subramaniyan, A. Skoogh, H. Salomonsson, P. Bangalore, and J. Bokrantz, ‘A data-driven algorithm to predict throughput bottlenecks in a production system based on active periods of the machines’, *Comput. Ind. Eng.*, vol. 125, pp. 533–544, Nov. 2018, doi: 10.1016/j.cie.2018.04.024.
[18]X. Su, J. Lu, C. Chen, J. Yu, and W. Ji, ‘Dynamic Bottleneck Identification of Manufacturing Resources in Complex Manufacturing System’, *Appl. Sci.*, vol. 12, no. 9, p. 4195, Apr. 2022, doi: 10.3390/app12094195.
[19]R. Panai Ersbo, ‘Compositional Analysis of Feedstock Materials for Cast Iron Production’, Bachelor’s thesis, Lund University, Lund, Sweden, 2026. Accessed: May 02, 2026. [Online]. Available: http://lup.lub.lu.se/student-papers/record/9223412
[20]M. Tabaszewski, P. Twardowski, M. Wiciak-Pikuła, N. Znojkiewicz, A. Felusiak-Czyryca, and J. Czyżycki, ‘Machine Learning Approaches for Monitoring of Tool Wear during Grey Cast-Iron Turning’, *Materials*, vol. 15, no. 12, p. 4359, Jun. 2022, doi: 10.3390/ma15124359.

[21] A. Laouissi, M. A. Yallese, A. Belbah, S. Belhadi, and A. Haddad, 'Investigation, modeling, and optimization of cutting parameters in turning of gray cast iron using coated and uncoated silicon nitride ceramic tools. Based on ANN, RSM, and GA optimization', *Int. J. Adv. Manuf. Technol.*, vol. 101, no. 1–4, pp. 523–548, Mar. 2019, doi: 10.1007/s00170-018-2931-8.

[22] A. Dos Reis *et al.*, 'Tool life assessment of high strength cast iron alloys in dry face milling operations', *J. Manuf. Process.*, vol. 111, pp. 180–193, Feb. 2024, doi: 10.1016/j.jmapro.2024.01.014.

[23] 'ASM handbook. 16: Machining', 9. ed., 4. print., Materials Park, Ohio: ASM International, 1999.

[24] G. Boothroyd, P. Dewhurst, and W. A. Knight, *Product design for manufacture and assembly*, 3. ed. in Manufacturing engineering and materials processing, no. 74. Boca Raton, Fla.: CRC Press, 2011.

[25] M. Alsafy, M. Atia, and J. Maher, 'A Cost Estimation Model for Machining Operations; an ANN Parametric Approach', 2017, *Unpublished*. doi: 10.13140/RG.2.2.15220.81287.

[26] V. P. Astakhov, *Drills: science and technology of advanced operations*, 1st edition. Boca Raton, [Florida]: CRC Press/Taylor & Francis, 2014.

[27] M. F. Ashby, *Materials selection in mechanical design*, 4th ed. Burlington, MA: Butterworth-Heinemann, 2011.

[28] D. M. Stefanescu, Ed., *ASM Handbook. Volume 1A, Cast iron science and technology*. Materials Park, OH: ASM International, 2017.

[29] S. Kalpakjian and S. R. Schmid, *Manufacturing engineering and technology*, Seventh edition. Upper Saddle River, NJ: Pearson, 2014.

[30] M. P. Groover, *Fundamentals of modern manufacturing: materials, processes, and systems*, Seventh edition. Hoboken, NJ: John Wiley & Sons, Inc., 2020.

[31] L. R. M. Dias and A. E. Diniz, 'Effect of the gray cast iron microstructure on milling tool life and cutting force', *J. Braz. Soc. Mech. Sci. Eng.*, vol. 35, no. 1, pp. 17–29, Mar. 2013, doi: 10.1007/s40430-013-0004-3.

[32] C. J. McCauley, *Machinery's handbook: a reference book for the mechanical engineer, designer, manufacturing engineer, draftsman, toolmaker, and machinist*, 26th ed. New York: Industrial press, 2000.

Appendix A: Simulation parameters & reference data

This appendix documents the numerical parameters used throughout the simulation framework. All values are reported at their baseline assignments. Variations explored in the sensitivity analysis of Section 4.5 are described separately in Appendix C.

A.1 Baseline operation parameters

The baseline machining parameters for the three operations are listed in Table A.1. Cutting parameters comes from typical values reviewed in literature and data from the Seco Tools website.; For OP30, the nominal reference tool life is approximately 38.9 minutes. However, the effective tool-life threshold used in the simulation is set to 24 minutes. This value is interpreted as a conservative planned replacement threshold. In this way, the simulation represents preventive tool-change practice, where the tool is replaced before the nominal end-of-life condition is reached. All times are expressed in seconds; the incident rate λ_0 is expressed in incidents per cutting hour.

Table A.1. Baseline operation parameters used in the SimPy simulation.

Parameter	OP10 face milling	OP20 drilling	OP30 heavy turning
Cut time (s)	6.38	2.29	528.00
Auxiliary time (s)	5.00	4.00	0.00
Setup allocation (s)	9.00	6.30	9.00
Nominal tool life (s)	660	600	1439
Tool change time (s)	60	30	500
Nominal incident rate λ_0 (per cutting hour)	0.100	0.100	0.083
Mean incident duration (s)	120	180	300
Baseline cutting speed v_{baseline} (m/min)	343	155	385

A.2 Cost-layer parameters

The cost-layer parameters used in Section 3.9 and applied in Cases C and 4.5 are listed in Table A.2. The currency is EUR; the SEK/EUR conversion uses an assumed rate of 1 SEK ≈ 0.087 EUR.

Table A.2. Baseline cost-layer parameters.

Parameter	Value	Source
Machine-hour rate OP10	65 EUR/h	Commercial European CNC rate for 3-axis milling
Machine-hour rate OP20	65 EUR/h	Commercial European CNC rate for 3-axis drilling
Machine-hour rate OP30	100 EUR/h	Commercial European CNC rate for heavy multi-axis turning
Tool-change cost OP10	134 EUR	Seco face-mill assembly, half-set indexing convention
Tool-change cost OP20	77 EUR	Seco solid carbide drill (216 EUR list price), amortised over 2.8 regrind cycles at 60% effective life
Tool-change cost OP30	15 EUR	Seco CNMA190612 insert (29.20 EUR per insert), half-edge indexing convention
Raw casting cost	30 EUR/part	European cast iron market, ~5 kg part
Scrap conversion rate α	0.30	Modelling assumption; sensitivity in Section 4.5
Shift duration	8 h	Standard production-day reference

A.3 Tooling specifications

The cutting tools selected for the three operations are listed in Table A.3, with their product codes, item numbers and list prices at the time of consultation (January–February 2026). Prices are reported in SEK as obtained from the Seco website, with EUR equivalents at the assumed exchange rate.

Table A.3. Tooling specifications.

Operation	Tool type	Product code	Item number	List price (SEK)	List price (EUR)
OP10	Face-mill holder	R220.54-0125-22-6A	03156812	12,300	1,070
OP10	Face-mill insert	SNMX2209ANTR-M18 MK1501	10312406	385	33.5
OP20	Solid carbide drill	SD207A-1000-062-10R1-P	03046378	2,480	216
OP30	Turning holder	PCLNR4040S19	02468695	2,080	181
OP30	Turning insert	CNMA190612 TK0501	03064386	335	29.2

A.4 Taylor module parameters

The Taylor tool-life equation $V \cdot T^n = C$ is used in Cases B and C to couple the OP30 cutting speed with the OP30 tool life. The Taylor exponent is assigned the value $n = 0.25$, consistent with literature reports for cemented-carbide machining of grey cast iron. The normalised cutting-speed ratio $v_{ratio} = v_c / v_{baseline}$ is swept from 0.70 to 1.30 in steps of 0.05, generating 13 operating points per material scenario. At $v_{ratio} = 1.0$ the cut time and tool life reduce to the baseline values reported in Table A.1. The cycle-time relation is $t_{cut}(v_{ratio}) = t_{cut,base} / v_{ratio}$ and the tool-life relation is $T(v_{ratio}) = T_{base} \cdot v_{ratio}^{(-1/n)}$. Material effects are applied through k_{wear} as $T_{eff} = T(v_{ratio}) / k_{wear}$.

A.5 Simulation execution parameters

The Monte Carlo replication settings and simulation runtime parameters are listed in Table A.4.

Table A.4. Simulation execution parameters.

Parameter	Value
Simulation horizon per replication	8 hours (28,800 s)
Buffer 1 capacity	10 parts
Buffer 2 capacity	10 parts
State-monitoring sample interval	30 s
Process noise distribution	Triangular(0.98, 1.00, 1.02) on cycle time
Monte Carlo replications, Case A	200
Monte Carlo replications, Case B (per v_{ratio})	200
Monte Carlo replications, Case C (per v_{ratio})	200
Monte Carlo replications, sensitivity (per combination)	200 (reusing Case C samples)
Random-seed convention	Per-replication offset, reproducible

Appendix B: Material scenario derivation

This appendix documents the complete numerical chain from compositional inputs to operation-specific machining multipliers. The reduced material-response model of Sections 3.5 to 3.6 is applied to each of the five compositional scenarios. Intermediate values are reported at full precision for traceability.

B.1 Compositional scenarios (wt.%)

Table B.1. Element composition for the five scenarios (weight percent).

Element	S0 baseline	S1 CE drift	S2 local heterogeneity	S3 residual hardening	S4 combined adverse
C	3.35	3.45	3.34	3.32	3.42
Si	2.10	2.35	2.08	2.02	2.28
Mn	0.55	0.56	0.42	0.62	0.45
S	0.080	0.075	0.120	0.075	0.115
P	0.060	0.070	0.105	0.070	0.110
Cu	0.20	0.22	0.21	0.48	0.46
Sn	0.030	0.033	0.032	0.080	0.075
Cr	0.080	0.085	0.085	0.180	0.170
Ni	0.100	0.105	0.100	0.260	0.240
Mo	0.020	0.020	0.020	0.060	0.055
V	0.010	0.010	0.012	0.030	0.030
Ti	0.010	0.010	0.012	0.025	0.025

B.2 Scaling factors and weights

Table B.2. Scaling factors and weights of the reduced material-response model.

Parameter group	Element / index	Value
Scaling factor s	CE	0.150
	S	0.035
	P	0.035
	Cu	0.250
	Sn	0.040
	Cr	0.080
	Ni	0.160
	Mo	0.040
	V	0.020
	Ti	0.020
Residual hardening weights	Cu	0.35
	Sn	0.55
	Cr	0.45
	Ni	0.20
	Mo	0.35
	V	0.25
	Ti	0.20
Local heterogeneity weights	Mn/S imbalance	0.45
	S+ excess	0.30
	P+ excess	0.25

B.3 Operation sensitivity coefficients

Table B.3. Operation sensitivity coefficients for the three machining multipliers.

Operation	Channel	I_{drift} coefficient	I_{local} coefficient	I_{res} coefficient
OP10	k_{time}	0.050	0.006	0.012
	k_{wear}	0.060	0.010	0.045
	λ_{mult}	0.15	0.25	0.08
OP20	k_{time}	0.025	0.006	0.015
	k_{wear}	0.040	0.020	0.070
	λ_{mult}	0.15	1.00	0.12
OP30	k_{time}	0.040	0.004	0.018
	k_{wear}	0.055	0.018	0.150
	λ_{mult}	0.12	0.65	0.10

Multiplier guardrails: $k_{time} \in [1.0, 1.25]$, $k_{wear} \in [1.0, 1.60]$, $\lambda_{mult} \in [1.0, 4.0]$.

B.4 Derived descriptors and material-response indicators

Table B.4. Material descriptors and indicators by scenario.

Scenario	CE	I_{drift}	Mn/S ratio	$I_{Mn/S}$	I_{res}	I_{local}	H*	M*
S0 baseline	4.0700	0.0000	6.8750	0.0000	0.0000	0.0000	1.0000	1.0000
S1 CE drift	4.2233	1.0222	7.4667	0.0826	0.1407	0.0372	1.0349	1.0517
S2 local heterogeneity	4.0613	0.0578	3.5000	0.6745	0.0540	0.5012	1.0123	1.0103
S3 residual hardening	3.9933	0.5111	8.2667	0.1843	2.2964	0.0830	1.1632	1.1183
S4 combined adverse	4.1900	0.8000	3.9130	0.5634	2.1843	0.4174	1.1631	1.1340

B.5 Operation-specific machining multipliers

Table B.5. Per-scenario, per-operation machining multipliers as used in the SimPy simulation.

Scenario	Operation	k_{time}	k_{wear}	λ_{mult}
S0 baseline	OP10	1.0000	1.0000	1.0000
	OP20	1.0000	1.0000	1.0000
	OP30	1.0000	1.0000	1.0000
S1 CE/matrix drift	OP10	1.0531	1.0747	1.1631
	OP20	1.0277	1.0510	1.1937
	OP30	1.0438	1.0750	1.1714
S2 local heterogeneity	OP10	1.0036	1.0098	1.1336
	OP20	1.0042	1.0138	1.5077
	OP30	1.0030	1.0153	1.3327
S3 residual hardening	OP10	1.0498	1.1395	1.2993
	OP20	1.0487	1.1987	1.5188
	OP30	1.0615	1.4111	1.4022
S4 combined adverse	OP10	1.0828	1.1744	1.5640
	OP20	1.0655	1.2232	2.3418
	OP30	1.0854	1.4257	1.9463

Appendix C: Full simulation results

This appendix provides the complete numerical results behind the figures and summary tables of Chapter 4. Means and standard deviations are computed over 200 Monte Carlo replications per condition.

C.1 Case A: configured bottleneck line

Table C.1. Case A complete results by scenario (200 replications).

Scenario	Finished mean	Finished std	Q05	Q95	OP30 tool changes (mean)	OP30 incidents (mean)	Output loss vs S0 (%)
S0 baseline	39.77	0.50	39.00	41.00	13.96	0.66	0.00
S1 CE/matrix drift	36.77	0.51	36.00	38.00	14.85	0.78	7.54
S2 local heterogeneity	38.83	0.48	38.00	40.00	14.15	1.10	2.36
S3 residual hardening	33.80	0.55	33.00	35.00	18.73	0.92	15.01
S4 combined adverse	32.81	0.62	31.95	34.00	19.21	1.28	17.50

C.2 Case B: cutting-speed sweep, production-rate response

Table C.2. Mean finished output per shift for each (scenario, v_{ratio}) combination, Case B.

v_{ratio}	S0	S1	S2	S3	S4
0.70	34.51	32.69	33.81	30.97	30.40
0.75	35.87	33.71	35.61	32.69	31.72
0.80	36.97	34.97	36.66	33.51	32.59
0.85	38.74	35.85	37.91	33.92	33.37
0.90	38.91	36.58	38.65	33.89	32.90
0.95	39.62	36.78	38.84	33.86	32.86
1.00	39.77	36.77	38.83	33.80	32.81
1.05	39.80	36.83	38.63	32.91	31.84
1.10	38.79	35.77	38.50	31.95	30.97
1.15	38.55	35.54	36.92	30.95	30.45
1.20	36.95	34.00	36.58	29.93	28.97
1.25	35.96	32.83	35.55	29.63	28.95
1.30	34.85	31.82	33.94	29.95	29.81

C.3 Case C: cutting-speed sweep, unit-cost response

Table C.3. Mean unit cost per finished part (EUR) for each (scenario, v_{ratio}) combination, Case C.

v_{ratio}	S0	S1	S2	S3	S4
0.70	85.24	88.44	86.54	92.16	93.67
0.75	83.55	87.36	84.07	89.78	91.74
0.80	82.39	85.89	83.37	89.21	90.97
0.85	80.76	85.34	82.00	89.36	90.44
0.90	81.29	85.02	81.79	90.27	92.19
0.95	81.06	85.43	82.19	91.55	93.54
1.00	81.65	86.34	83.07	92.61	94.55
1.05	82.36	87.01	85.39	95.16	97.81
1.10	84.83	89.94	85.45	98.62	100.84
1.15	86.02	91.11	88.50	101.66	103.18
1.20	89.20	94.86	90.10	105.65	108.32
1.25	91.78	98.45	92.90	107.36	108.86
1.30	94.86	101.45	96.78	106.53	106.99

C.4 Cost sensitivity sweep (27 combinations)

Table C.4. Cost-sensitivity sweep results: S3 and S4 unavoidable cost penalty across 27 parameter combinations.

α (scrap rate)	Raw cost (EUR)	OP30 tool cost (EUR)	S0 min unit cost	S3 min unit cost	S4 min unit cost	S3 penalty (%)	S4 penalty (%)
0.10	20	10	69.52	76.31	77.20	9.77	11.05
0.10	20	15	71.63	78.85	80.05	10.07	11.76
0.10	20	30	73.75	82.80	84.78	12.28	14.96
0.10	30	10	79.63	86.85	88.12	9.07	10.66
0.10	30	15	80.69	89.21	90.49	10.55	12.14
0.10	30	30	82.86	92.94	94.85	12.16	14.47
0.10	40	10	89.52	97.62	98.60	9.05	10.15

α (scrap rate)	Raw cost (EUR)	OP30 tool cost (EUR)	S0 min unit cost	S3 min unit cost	S4 min unit cost	S3 penalty (%)	S4 penalty (%)
0.10	40	15	90.66	100.06	101.32	10.37	11.76
0.10	40	30	92.30	103.91	105.36	12.58	14.15
0.30	20	10	69.62	76.41	77.50	9.75	11.31
0.30	20	15	71.78	78.97	80.31	10.01	11.88
0.30	20	30	73.82	82.94	85.02	12.35	15.16
0.30	30	10	79.66	86.92	88.28	9.11	10.82
0.30	30	15	80.76	89.21	90.44	10.46	11.98
0.30	30	30	82.93	93.04	95.07	12.19	14.64
0.30	40	10	89.66	97.82	98.88	9.10	10.29
0.30	40	15	90.79	100.16	101.51	10.32	11.81
0.30	40	30	92.42	104.06	105.58	12.59	14.24
0.50	20	10	69.78	76.55	77.80	9.70	11.49
0.50	20	15	71.92	79.16	80.55	10.07	11.99
0.50	20	30	73.90	83.07	85.15	12.42	15.23
0.50	30	10	79.78	87.10	88.45	9.18	10.86
0.50	30	15	80.92	89.39	90.66	10.47	12.04
0.50	30	30	83.01	93.21	95.29	12.29	14.79
0.50	40	10	89.80	98.03	99.16	9.17	10.42
0.50	40	15	90.91	100.32	101.69	10.36	11.86
0.50	40	30	92.51	104.21	105.76	12.65	14.32

Summary statistics across the 27 combinations:

S3 unavoidable cost penalty: min 9.05%, median 10.46%, max 12.42%

S4 unavoidable cost penalty: min 10.15%, median 12.12%, max 15.23%

Appendix D: Source code

This appendix presents the core algorithmic implementations of the framework. The complete codebase is published as an open-source repository at https://github.com/ChenghanQian/Cast-Iron-Machining-Simulation (commit hash 153ebae08cf9901d15254c1bb52c119f315281fb, accessed 2026-05-25) with a Zenodo DOI for archival reference(10.5281/zenodo.20370260). The repository contains the full Jupyter notebook, requirements file, raw simulation outputs (CSV) and a README with reproducibility instructions. The five code listings below provide the methodologically significant components of the framework; auxiliary code for file I/O, plotting and Excel export is omitted from the appendix but is included in the repository.

D.1 Taylor module (Section 3.7)

The Taylor module computes the OP30 cycle time and nominal tool life at any normalised cutting speed v_{ratio}. At v_{ratio} = 1.0 the function returns the baseline values reported in Appendix A.1, confirming numerical identity with the calibration anchor.

```
pythonTAYLOR_N = 0.25


OP30_BASELINE = {
    "cut_time_base": 528.00,      # seconds
    "tool_life_base": 1439.00,    # seconds
    "tool_change_time": 500,
    "setup_alloc": 9.00,
    "aux_time": 0.00,
    "lambda0": 0.083,
    "incident_mean": 300,
}


def      op30_params_at_vratio(v_ratio,      baseline=OP30_BASELINE,
n=TAYLOR_N):
    """
    Compute OP30 cut time and tool life at the given v_ratio.
```

```
    cut_time(v) = cut_time_base / v_ratio
    Taylor:  T(v) = T_base * v_ratio^(-1/n)
    """
    if v_ratio <= 0:
        raise ValueError("v_ratio must be positive")
    return {
        "cut_time": baseline["cut_time_base"] / v_ratio,
        "tool_life": baseline["tool_life_base"] * (v_ratio ** (-1.0
/ n)),
        "tool_change_time": baseline["tool_change_time"],
        "setup_alloc": baseline["setup_alloc"],
        "aux_time": baseline["aux_time"],
        "lambda0": baseline["lambda0"],
        "incident_mean": baseline["incident_mean"],
    }
```

D.2 Composition-to-multiplier chain (Sections 3.5–3.6)

The composition-to-multiplier chain maps a compositional scenario through descriptors and material-response indicators into operation-specific machining multipliers. The full chain is implemented as three composable functions.

```
pythonimport numpy as np


def carbon_equivalent(comp):
    return comp["C"] + (comp["Si"] + comp["P"]) / 3.0


def positive_excess(value, base, scale, cap=3.0):
    return min(max(0.0, (value - base) / scale), cap)


def composition_to_descriptors(comp, base_comp, scaling, residual_w,
local_w):
    """Composition -> reduced descriptors I_drift, I_res, I_local,
etc."""
    ce = carbon_equivalent(comp)
    base_ce = carbon_equivalent(base_comp)
    i_drift = abs((ce - base_ce) / scaling["CE"])

    mn_s_ratio = comp["Mn"] / max(comp["S"], 1e-9)
    base_mn_s = base_comp["Mn"] / max(base_comp["S"], 1e-9)
    i_mn_s = abs(np.log(mn_s_ratio / base_mn_s))

    s_plus = positive_excess(comp["S"], base_comp["S"], scaling["S"])
    p_plus = positive_excess(comp["P"], base_comp["P"], scaling["P"])

    i_res = sum(
        residual_w[el] * positive_excess(comp[el], base_comp[el],
scaling[el])
        for el in residual_w
```

```
    )
    i_local = local_w["MnS"]*i_mn_s + local_w["S"]*s_plus +
local_w["P"]*p_plus

    return {"I_drift": i_drift, "I_res": i_res, "I_local": i_local}



def descriptors_to_indicators(desc):
    """Descriptors -> relative material-response indicators H*,
M*."""
    H_star = (1.0 + 0.025*desc["I_drift"]
              + 0.060*desc["I_res"] + 0.015*desc["I_local"])
    M_star = (1.0 + 0.045*desc["I_drift"]
              + 0.040*desc["I_res"] + 0.010*desc["I_local"])
    return {**desc, "H_star": H_star, "M_star": M_star}



def indicators_to_multipliers(ind, op_sens):
    """Indicators -> operation-specific multipliers k_time, k_wear,
lambda_mult.

    op_sens contains, for one operation, the nine sensitivity
coefficients
    {time_drift, time_local, time_res, wear_..., lambda_...}.
    """
    ktime = (1.0 + op_sens["time_drift"]*ind["I_drift"]
             + op_sens["time_local"]*ind["I_local"]
             + op_sens["time_res"]*ind["I_res"])
    kwear = (1.0 + op_sens["wear_drift"]*ind["I_drift"]
             + op_sens["wear_local"]*ind["I_local"]
             + op_sens["wear_res"]*ind["I_res"])
    lam = (1.0 + op_sens["lambda_drift"]*ind["I_drift"]
           + op_sens["lambda_local"]*ind["I_local"]
```

```
        + op_sens["lambda_res"]*ind["I_res"])
    # Reduced-order guardrails
    ktime = min(max(ktime, 1.0), 1.25)
    kwear = min(max(kwear, 1.0), 1.60)
    lam = min(max(lam, 1.0), 4.0)
    return {"ktime": ktime, "kwear": kwear, "lambda_mult": lam}
```

D.3 SimPy production-flow model (Section 3.8)

The SimPy model represents one worker per operation. Each worker runs an event loop that consumes parts from its input store, processes them under the scenario-specific multipliers, consumes tool life, possibly triggers a tool change and a stochastic incident, and finally transfers the part downstream.

```
pythonimport simpy, math, random


class StationWorker:
    def __init__(self, env, station, input_store, output_store,
                 op_params, scenario_params, rng, sink=None):
        self.env, self.station = env, station
        self.input_store,    self.output_store    =    input_store,
output_store
        self.op_params,    self.scenario_params    =    op_params,
scenario_params
        self.rng, self.sink = rng, sink
        self.tool_consumed = 0.0
        self.parts_processed = 0
        self.tool_changes = 0
        self.incidents = 0
        env.process(self.run())


    def run(self):
        while True:
            part = yield self.input_store.get()
```

```
            ktime = self.scenario_params["ktime"]
            kwear = self.scenario_params["kwear"]
            lam_mult = self.scenario_params["lambda_mult"]

            cut_eff = self.op_params["cut_time"] * ktime
            aux_eff = self.op_params["aux_time"] * ktime
            noise = self.rng.triangular(0.98, 1.00, 1.02)
            proc_time    =    (cut_eff    +    aux_eff)    *    noise    +
self.op_params["setup_alloc"]

            yield self.env.timeout(proc_time)
            self.parts_processed += 1

            # Tool-life consumption
            self.tool_consumed += cut_eff * kwear
            if self.tool_consumed >= self.op_params["tool_life"]:
                yield
self.env.timeout(self.op_params["tool_change_time"])
                self.tool_changes += 1
                self.tool_consumed -= self.op_params["tool_life"]

            # Stochastic incident
            lam_eff = self.op_params["lambda0"] * lam_mult
            p_inc = 1.0 - math.exp(-lam_eff * (cut_eff / 3600.0))
            if self.rng.random() < p_inc:
                dt        =        self.rng.expovariate(1.0        /
self.op_params["incident_mean"])
                yield self.env.timeout(dt)
                self.incidents += 1

            # Transfer downstream
            if self.output_store is not None:
```

```
                yield self.output_store.put(part)
            else:
                self.sink["finished"] += 1



def         run_one_replication(scenario_params,        op30_params,
op_params_other,
                        sim_time=8*3600, buffer_cap=10, seed=0):
    rng = random.Random(seed)
    env = simpy.Environment()
    raw = simpy.Store(env, capacity=100000)
    buf1 = simpy.Store(env, capacity=buffer_cap)
    buf2 = simpy.Store(env, capacity=buffer_cap)
    raw.items.extend(range(100000))
    sink = {"finished": 0}

    StationWorker(env, "OP10", raw, buf1,
                  op_params_other["OP10"],  scenario_params["OP10"],
rng)
    StationWorker(env, "OP20", buf1, buf2,
                  op_params_other["OP20"],  scenario_params["OP20"],
rng)
    StationWorker(env, "OP30", buf2, None,
                  op30_params,     scenario_params["OP30"],     rng,
sink=sink)

    env.run(until=sim_time)
    return sink["finished"]
```

D.4 Cost-layer module (Section 3.9)

The cost layer converts each replication's finished output, tool-change counts and incident counts into a per-part unit cost through fully-allocated machine-hour accounting, list-price tool costing, market-rate raw-material costing and probabilistic scrap costing.

```
pythonCOST_PARAMS = {
    "rate_OP10": 65.0, "rate_OP20": 65.0, "rate_OP30": 100.0,    # EUR/h
    "tool_OP10": 134.0, "tool_OP20": 77.0, "tool_OP30": 15.0,    # EUR/event
    "raw_per_part": 30.0,                                        # EUR/part
    "alpha_scrap": 0.30,                                         # incident->scrap
    "shift_hours": 8.0,
}


def compute_unit_cost(replication_result, op_params_full, cost=COST_PARAMS):
    """
    Convert one SimPy replication output into a unit cost.
    replication_result keys: finished, OP30_tool_changes, OP30_incidents, ...
    """
    finished = replication_result["finished"]
    if finished <= 0:
        return None

    # Machine cost: fully allocated over shift
    c_machine = cost["shift_hours"] * (
        cost["rate_OP10"] + cost["rate_OP20"] + cost["rate_OP30"]
    )

    # Tool cost: per recorded change event
```

```
    c_tool = (
        replication_result.get("OP10_tool_changes", 0) * cost["tool_OP10"]
        + replication_result.get("OP20_tool_changes", 0) * cost["tool_OP20"]
        + replication_result["OP30_tool_changes"] * cost["tool_OP30"]
    )

    # Raw material cost
    c_raw = finished * cost["raw_per_part"]

    # Scrap cost: alpha * incidents * (raw + half OP30 processing)
    op30_partial = (op_params_full["OP30"]["cut_time"] / 2.0) / 3600.0
    scrap_per_event = cost["raw_per_part"] + op30_partial * cost["rate_OP30"]
    c_scrap = (cost["alpha_scrap"]
               * replication_result["OP30_incidents"]
               * scrap_per_event)

    return (c_machine + c_tool + c_raw + c_scrap) / finished
```

D.5 v~ratio~ sweep main loop (Sections 4.3–4.4)

The main sweep loop iterates over the five compositional scenarios and the 13 cutting-speed ratios, executing 200 Monte Carlo replications per combination and recording finished output, tool-change counts, incident counts and the resulting unit cost.

```
pythonimport pandas as pd

VRATIO_LIST = [round(0.70 + 0.05*i, 3) for i in range(13)]
N_REPS = 200

rows = []
for scenario_name, scenario_params in SCENARIOS.items():
```

```
    for v_ratio in VRATIO_LIST:
        op30_params = op30_params_at_vratio(v_ratio)
        for rep in range(N_REPS):
            seed = 30000 + rep + int(v_ratio * 10000)
            result = run_one_replication(
                scenario_params=scenario_params,
                op30_params=op30_params,
                op_params_other=OP_PARAMS_BASE,
                seed=seed,
            )
            unit_cost = compute_unit_cost(result, OP_PARAMS_FULL)
            rows.append({
                "scenario": scenario_name,
                "v_ratio": v_ratio,
                "replication": rep,
                "finished": result["finished"],
                "OP30_tool_changes": result["OP30_tool_changes"],
                "OP30_incidents": result["OP30_incidents"],
                "unit_cost": unit_cost,
            })

sweep_df = pd.DataFrame(rows)


# Aggregate per (scenario, v_ratio)
summary = (sweep_df.groupby(["scenario", "v_ratio"])
                   .agg(finished_mean=("finished", "mean"),
                        unit_cost_mean=("unit_cost", "mean"))
                   .reset_index())


# Find per-scenario cost optimum
cost_optima                                                          =
(summary.loc[summary.groupby("scenario")["unit_cost_mean"].idxmin()]
```

```
.reset_index(drop=True))
```

The output of this sweep, when seeded as documented in the repository, reproduces the production-rate values of Table C.2, the unit-cost values of Table C.3 and the cost-optima reported in Section 4.4 within numerical precision.