Computers & Industrial Engineering 216 (2026) 112010

Contents lists available at ScienceDirect

Computers & Industrial Engineering

journal homepage: www.elsevier.com/locate/caie

Multi-objective optimal acquisition and production planning of 3D printing technologies under uncertainty

Federica Tomelleri *, Matteo Brunelli

Department of Industrial Engineering, University of Trento, via Sommarive, 9, Trento, 38123, Italy

ARTICLE INFO

Keywords:
Multi-objective optimization
Additive manufacturing
Uncertainty
Technology acquisition
Production planning

ABSTRACT

The rapid evolution of Additive Manufacturing (AM) technologies offers significant opportunities for process improvement, yet selecting the most suitable 3D printers remains challenging due to the high specialization and cost of equipment tailored to specific applications such as prototyping. This complexity is further compounded by uncertainty in future demand and production capacity. To support informed decision-making, this study introduces a stochastic multi-objective optimization framework for the acquisition and utilization of AM technologies over a multi-period planning horizon. The proposed model simultaneously: (i) minimizes the total discounted acquisition cost, (ii) maximizes the probability of satisfying fluctuating demand, and (iii) maximizes the probability of maintaining feasibility with respect to both machine capacity and human supervision availability. Uncertainty is addressed through a Monte Carlo simulation approach and, given the high-dimensional and nonlinear nature of the problem, the solution strategy relies on an evolutionary algorithm to efficiently explore trade-offs among competing goals. The framework is examined using industry-inspired performance data, demonstrating its ability to identify Pareto-optimal solutions and highlighting its potential as a basis for future decision-support applications in AM system planning.

1. Introduction

Additive Manufacturing (AM) technologies are transforming the industrial landscape by enabling the production of complex geometries, customized components, and small batch sizes with unprecedented design freedom (Franco et al., 2020; Jimo et al., 2025). Unlike traditional subtractive methods, which rely on material removal, AM constructs parts layer by layer directly from digital models, offering significant advantages in material efficiency and time-to-market (Gibson et al., 2021; Kanishka & Acherjee, 2023). Consequently, AM is increasingly adopted across a wide range of industrial sectors, including but not limited to aerospace, biomedical, automotive, and tooling (Kanishka & Acherjee, 2023; Pourhejazy et al., 2025). AM is employed not only for prototyping but also for the production of functional components, patient-specific implants, lightweight structures, and end-use parts in low-volume settings (Jimo et al., 2025). Importantly, AM is also recognized as a key enabler for Industry 4.0, facilitating digital manufacturing integration and flexible production paradigms (Rüßmann et al., 2015; Zheng et al., 2021).

Despite these opportunities, the large-scale implementation of AM poses non-trivial decision-making challenges. Production managers must navigate a heterogeneous portfolio of technologies, ranging from low-cost, high-throughput polymer systems to capital-intensive metal platforms operating in controlled environments. Investment decisions are typically guided by machine datasheets, yet translating these technical attributes into long-term production capacity remains difficult. Selecting an inappropriate technology — such as prioritizing build volume over feature resolution, or minimizing acquisition cost at the expense of high operating and post-processing requirements — may lock firms into suboptimal capabilities over multiple planning periods. Moreover, part-to-machine assignment is inherently intertwined with acquisition choices, since machines procured today constrain future production configurations. For example, investments in high-resolution photopolymer systems may favor aesthetic or biomedical parts while creating bottlenecks for structurally demanding components that require powder-based or metal technologies. These strategic and tactical decisions are further complicated by demand volatility and variability in processing times due to setup, part geometry, post-processing, and maintenance, making it difficult to assess whether a given fleet of 3D printers can reliably meet service-level targets over time. Addressing these challenges requires integrated optimization frameworks that go beyond static technology selection.

This study introduces a multi-objective optimization framework to analyze the long-term planning of 3D printing acquisition and part-to-machine assignment, incorporating variability in both demand and

* Corresponding author.
E-mail addresses: federica.tomelleri@unitn.it (F. Tomelleri), matteo.brunelli@unitn.it (M. Brunelli).

https://doi.org/10.1016/j.cie.2026.112010
Received 29 September 2025; Received in revised form 31 March 2026; Accepted 31 March 2026
Available online 1 April 2026
0360-8352/© 2026 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

production capacity. The uncertainty is handled through Monte Carlo simulation, enabling the estimation of probabilistic performance indicators across a range of future scenarios. The framework simultaneously optimizes three primary objectives: (1) Minimizing time-discounted acquisition costs; (2) Maximizing the probability of meeting variable customer demand; and (3) Maximizing the probability that the resulting production plan remains feasible with respect to machine capacity and human supervision constraints. Given the presence of conflicting objectives, the solution approach relies on the Non-dominated Sorting Genetic Algorithm II (NSGA-II) (Deb et al., 2002). This evolutionary algorithm is particularly well suited to discrete, nonlinear decision spaces and enables effective exploration of the Pareto front, offering a diverse set of trade-off solutions. The proposed model is validated using a realistic experimental setup with synthetic data reflecting plausible AM configurations, acquisition costs, and stochastic demand and capacity profiles. The numerical results provide insight into the solution space and show how different acquisition and production strategies perform under uncertainty, suggesting directions for future decision-support approaches in AM.

The paper is structured as follows. Section 2 reviews the relevant literature, highlighting the research gaps addressed in this work. Section 3 introduces the deterministic linear programming formulation that serves as the foundation of the model. Section 4 extends this formulation to incorporate uncertainty through a stochastic programming approach. Section 5 describes the solution strategy based on evolutionary multi-objective optimization. Section 6 presents an experimental setup to validate the model using representative data. Finally, Section 7 discusses model extensions and implementation aspects, and Section 8 concludes the paper.

2. Literature review

AM is enhancing manufacturing supply chain flexibility through its ability to support on-demand, customized, and distributed production (Akbari, 2023; Franco et al., 2020). However, large-scale implementation presents a complex set of decision-making challenges. A key issue lies in the acquisition of AM technologies, which involve high capital investment and significant risk (Kanishka & Acherjee, 2023; Zhai et al., 2025). The current AM market is characterized by a proliferation of heterogeneous technologies, such as Powder Bed Fusion (PBF), Material Extrusion (MEX), and Directed Energy Deposition (DED), each optimized for specific applications and material requirements (Algunaid & Liu, 2022; Gibson et al., 2021). Moreover, the rapid evolution of AM — with frequent advancements in materials, processes, and equipment — makes 3D printer selection increasingly complex, requiring a nuanced understanding of technical specifications and long-term strategic objectives (Algunaid & Liu, 2022; Jimo et al., 2025). Several studies have proposed Multi-Criteria Decision Analysis (MCDA) methods to support various types of selection decisions in AM, including technology, machine, and service provider evaluation (Framinan et al., 2023; Pourhejazy et al., 2025; Tomelleri et al., 2025). However, although dynamic and stochastic extensions of MCDA exist, these methods are typically most effective for trade-off assessment in settings characterized by a limited number of alternatives/solutions known a priori and a relatively small decision space, and most applications in the AM literature focus on single-period formulations.

Another critical challenge concerns the assignment of parts to machines (Pinto et al., 2024; Zheng et al., 2021). Modern AM facilities increasingly rely on distributed networks, where different 3D printers with heterogeneous capabilities must be coordinated effectively (Framinan et al., 2023). In this context, the inherently multi-objective nature of AM system planning, where economic and operational goals often conflict, requires the application of Pareto-based optimization techniques capable of identifying trade-offs rather than single-point optima. For example, Ransikarbum et al. (2017) developed a multi-objective optimization model for part-to-printer assignment and scheduling in AM networks, minimizing operating cost, load imbalance, tardiness, and unprinted parts. More recently, Altekin and Bukchin (2022) proposed an optimization model for production planning in multi-machine systems, jointly optimizing part-to-job and job-to-machine assignments under cost and makespan objectives. Within this stream of research, several studies have increasingly adopted metaheuristic approaches. For instance, Lu et al. (2023) proposed a fast pixel-based packing algorithm to efficiently solve AM production planning problems under irregular 2D constraints, achieving sub-second feasibility checks and significant makespan reductions. In a similar direction, Aloui and Hadj-Hamou (2021) tackled a scheduling and nesting problem under technological constraints in AM systems, proposing a heuristic approach that integrates production time estimation with a mixed-integer model to minimize delays and maximize machine utilization. However, as highlighted by Framinan et al. (2023), most studies in this field address isolated decision layers, providing only partial support to practitioners, whereas the integration of strategic and tactical decisions has mainly been investigated in other domains. For instance, Amorim et al. (2019) combined optimization with metamodel-based simulation to jointly determine emergency medical services station locations (strategic) and vehicle allocations (tactical), while Luo et al. (2021) developed a forest supply chain model linking long-term forest management with sawmill operations. Similarly, Aazami and Saidi-Mehrabad (2021) proposed a bi-level, multi-period production–distribution planning model for perishable products under seller–buyer competition, accounting for demand sensitivity to price and freshness, and solved it using a hybrid genetic algorithm and Benders decomposition. Along the same line, Makui et al. (2016) developed a robust, multi-site aggregate production planning model for products with very limited lifecycles, integrating production, workforce, and inventory decisions under demand and cost uncertainty. These studies showed the benefits of integration but rely on domain-specific assumptions (e.g., facility location, forest rotations, or perishable product) that cannot be easily adapted to AM. This highlights the need for AM-specific frameworks that capture its unique features, such as heterogeneous machine capabilities, high capital costs, and uncertain production performance.

These challenges are further exacerbated by several sources of uncertainty inherent in the operational integration of AM technologies. Chief among these is demand fluctuation, which is particularly critical in low-volume or customized production contexts, where future orders are difficult to forecast accurately (Homem-de Mello & Bayraksan, 2014; Pinto et al., 2024; Sgarbossa et al., 2021). This uncertainty impacts capacity planning and return-on-investment calculations, complicating decisions regarding the scale and timing of machine acquisitions. In parallel, machine reliability introduces additional sources of uncertainty. Rather than benefiting from the maturity and standardization typical of conventional manufacturing systems, AM technologies still exhibit significant variability in production times, operating costs, and output quality (Gibson et al., 2021; Pinto et al., 2024; Zhai et al., 2025). These factors affect production efficiency, but also the ability to meet lead time and quality requirements (Kanishka & Acherjee, 2023; Pourhejazy et al., 2025). Traditional deterministic models fail to capture this variability, offering limited guidance in real-world contexts (Marmarelis & Ghanem, 2020). Recent developments in operations research emphasized the need for stochastic and robust optimization approaches to address planning challenges in AM (Sgarbossa et al., 2021; Stittgen & Schleifenbaum, 2021; Yılmaz, 2020). These methods incorporate uncertainty through a variety of modeling strategies, including scenario analysis, stochastic programming, and Monte Carlo simulation, enabling adaptive and risk-aware decision-making. However, most existing studies consider either demand uncertainty or operational variability in isolation, overlooking their interdependence and missing the opportunity to address both dimensions within an integrated planning framework.

Other recent contributions have explored the role of AM hubs as strategic components in distributed networks, facilitating efficient production and delivery of customized or spare parts. For instance, Ahmed

Table 1
Comparison of related studies with the proposed model.

Study	AM-focused	Strategic decisions	Tactical decisions	Multi-objective	Uncertainty
Tomelleri et al. (2025)	✓	✓(machine selection)	✗	✗	✓(printing parameters & costs)
Pourhejazy et al. (2025)	✓	✓(service provider selection)	✗	✗	✗
Ransikarbum et al. (2017)	✓	✗	✓(single-type allocation)	✓	✗
Altekin and Bukchin (2022)	✓	✗	✓(single-type allocation)	✓	✗
Lu et al. (2023)	✓	✗	✓(part allocation & packing)	✗	✗
Aloui and Hadj-Hamou (2021)	✓	✗	✓(scheduling & nesting)	✗	✗
Amorim et al. (2019)	✗	✓(station location)	✓(vehicle allocation)	✗	✓(demand & travel times)
Luo et al. (2021)	✗	✓(forest management)	✓(sawmill operations)	✗	✓(discount rate, harvest cost & volume)
Aazami and Saidi-Mehrabad (2021)	✗	✓(production–distribution planning)	✓(pricing & allocation)	✗	✓(demand & freshness)
Makui et al. (2016)	✗	✓(production planning)	✓(inventory & workforce)	✗	✓(demand & costs)
Ahmed et al. (2023)	✓	✓(supply chain network)	✗	✗	✓(demand & supply)
Strong et al. (2019)	✓	✓(facility location & capacity)	✗	✗	✓(demand & capacity)
Gao et al. (2024)	✓	✓(hub routing & delivery)	✓(production scheduling)	✗	✗
Proposed model	✓	✓(technology acquisition)	✓(multi-type allocation)	✓	✓(demand & processing times)

et al. (2023) developed a three-stage stochastic optimization model to evaluate the strategic role of AM hubs in traditional manufacturing supply chains, showing how AM can act as a responsive and cost-effective recourse option under demand uncertainty, and proposing an analytical allocation rule for optimal sourcing decisions. Strong et al. (2019) proposed a two-stage p-median location model to optimize the placement and capacity of hybrid metal AM hubs in the U.S., integrating AM with post-processing services like heat treatment and machining. Due to the complexity of these problems, metaheuristic approaches are commonly adopted also in this field. For example, Gao et al. (2024) investigated a mobile AM hub operating under a "printing-while-delivering" paradigm, and proposed an approach combining GAs and branch-and-bound to optimize production and delivery coordination. These studies aim to integrate AM into supply chains, with a specific focus on the strategic localization of hubs. Our model, by jointly addressing technology selection and optimal part allocation decisions under uncertainty, aims to complement these contributions. A comparative overview of related works and our framework is provided in Table 1.

3. An integer linear programming formulation

To support decision-making in AM system planning, we develop an optimization model that guides the selection and utilization of 3D printers over a finite, multi-period planning horizon. As an initial step, we present the deterministic formulation of the model, in which uncertain parameters are assumed to be fixed. This formulation establishes the structure of the problem, which is later extended to incorporate stochasticity. Throughout the paper, we adopt the notation $\mathbb{Z}_{\geq}$ and $\mathbb{Z}_{>}$ for the sets of non-negative and positive integers, respectively. The complete list of sets, parameters, and variables used in the model is provided in Appendix A.

The model is defined over three primary sets. The set of machine types, $\mathcal{M} = \{1, \dots, M\}$, includes all AM technologies under consideration, with each element representing a specific 3D printer model characterized by its own technical and operational features. The set of parts, $\mathcal{P} = \{1, \dots, P\}$, represents homogeneous families of products to be manufactured, such as prototypes, functional components, or biomedical devices. The planning horizon is discretized into a finite set of time periods, $\mathcal{T} = \{1, \dots, T\}$, where each period represents, for instance, a quarter or a year. At each time step, the model determines

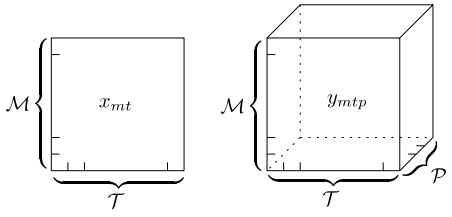

Fig. 1. Visual representation of the integer-valued decision variables.

two categories of decisions. The first pertains to *technology acquisition*: the variable $x_{mt} \in \mathbb{Z}_{\geq}$ denotes the number of machines of type $m \in \mathcal{M}$ acquired at the beginning of period $t \in \mathcal{T}$, which directly affects both production capacity and capital investment. The second concerns *workload allocation*: the variable $y_{mpt} \in \mathbb{Z}_{\geq}$ indicates the number of units of part $p \in \mathcal{P}$ produced using machines of type $m \in \mathcal{M}$ during period $t \in \mathcal{T}$. These quantities form a two-dimensional matrix and a three-dimensional tensor, as illustrated in Fig. 1. The total number of integer decision variables is therefore $|\mathcal{M}| \cdot |\mathcal{T}| \cdot (|\mathcal{P}| + 1)$ and corresponds to the degree of freedom of the optimization problem. For example, consider a setting with two machine types (`MEX-100`, `PBF-450`), three part families (`Gear`, `Bracket`, `Implant`), and a four-period planning horizon. The resulting decision space encompasses 6 technology-product combinations over 4 successive time periods, leading to 8 acquisition variables and 24 production variables.

The primary objective of this deterministic model is to minimize the total acquisition cost of AM equipment over the entire planning horizon, while accounting for the time value of money through a discounting mechanism. Let $c_m > 0$ denote the acquisition cost of a single machine of type m, and let $r > 0$ be the discount rate. The cost incurred by acquiring machines in period t is discounted accordingly to reflect its present value. The resulting objective function can be expressed as:

$$\min f_1 = \sum_{m \in \mathcal{M}} \sum_{t \in \mathcal{T}} \frac{c_m x_{mt}}{(1+r)^t} \tag{O1}$$

To avoid unrealistic investment profiles and ensure economic feasibility, the model also considers a per-period budget constraint. Specifically, we assume that a maximum amount of money is available in each time period, denoted by $B_t > 0$. This reflects capital rationing and cash

flow management. The total cost of machine acquisitions in period t is constrained not to exceed the available budget B_t:

$$\sum_{m\in\mathcal{M}} c_m x_{mt} \le B_t \quad \forall\, t \in \mathcal{T} \tag{C1}$$

To determine how many machines are effectively available for production in each time period, we introduce an auxiliary variable a_{mt}, which denotes the number of machines of type m that are operational in period t. This variable accounts for both newly acquired technologies and pre-existing machines. Specifically, machine availability at time t can arise from two contributing components: (i) machines that were already available at the beginning of the planning horizon and have *not* yet reached the end of their operational life, and (ii) machines acquired at earlier stages of the planning horizon whose operational lifetime still includes period t. Let $s_{mt} \in \mathbb{Z}_{\ge}$ represent the number of machines of type m from the initial inventory that are still active in period t. Additionally, any machine of type m acquired at time $t' \le t$ contributes to availability at time t if it has not yet exceeded its lifetime L_m. The total number of available machines is therefore given by:

$$a_{mt} = s_{mt} + \sum_{t'=\max\{1,t-L_m+1\}}^{t} x_{mt'} \quad \forall\, m \in \mathcal{M},\ t \in \mathcal{T} \tag{C2}$$

To illustrate, consider a scenario involving two MEX-100 machines available at the start of the planning horizon: one with two remaining operational periods and the other with only one. In $t = 1$, both machines are active, so $a_{\text{MEX-100},1} = 2$. In $t = 2$, only the first machine remains active, yielding $a_{\text{MEX-100},2} = 1$, assuming no new acquisitions. If a new MEX-100 machine is acquired in $t = 2$, and its lifetime is $L_{\text{MEX-100}} = 3$, it will contribute to machine availability in periods $t = 2, 3, 4$.

The model also incorporates organizational preferences concerning technological diversity. To reduce managerial complexity, avoid supplier fragmentation, and streamline maintenance, training, and technical support, the number of distinct machine brands that can be operated concurrently is constrained. Let $\mathcal{M} = S_1 \cup \cdots \cup S_q$ denote a partition of the set of machine types, where each subset S_h corresponds to the group of machines belonging to the hth brand. Let u be the maximum number of brands allowed to be active in any given period. To enforce this restriction, we introduce binary variables $b_{ht} \in \{0, 1\}$, where $b_{ht} = 1$ indicates that brand h is active in period t, i.e., at least one machine belonging to brand h is operational. This relationship is enforced through the following constraints:

$$\sum_{m\in S_h} a_{mt} \le M \cdot b_{ht} \quad \forall\, h = 1, \ldots, q,\ \ \forall\, t \in \mathcal{T} \tag{C3}$$

$$\sum_{h=1}^{q} b_{ht} \le u \quad \forall\, t \in \mathcal{T} \tag{C4}$$

(C3) ensures that b_{ht} is set to 1 whenever at least one machine from brand h is in use during period t, using a sufficiently large constant M as a big-M parameter. (C4) then limits the total number of active brands in each period to at most u.

A key component of the deterministic formulation is demand satisfaction over the entire planning horizon. For each p and each t, a known demand level $d_{pt} \in \mathbb{Z}_{\ge}$ must be met to ensure service level compliance. However, due to technological differences among AM systems, not all machines can produce every part. For instance, a polymer extrusion system may be suitable for prototyping plastic components but inadequate for fabricating metal parts. To capture these constraints, we define a compatibility set $\mathcal{I} \subseteq \mathcal{M} \times \mathcal{P}$, which includes all admissible machine–part pairs (m, p). Consequently, production is restricted to feasible assignments, and the total number of units produced for each part must meet or exceed the demand:

$$\sum_{\substack{m\in\mathcal{M}\\(m,p)\in\mathcal{I}}} y_{mpt} \ge d_{pt} \quad \forall\, p \in \mathcal{P},\ t \in \mathcal{T} \tag{C5}$$

Moreover, to ensure that the production plan is operationally feasible, the model includes a capacity constraint that limits the workload assigned to each machine type. For every m and t, the total amount of time required to produce all compatible parts must not exceed the maximum processing capacity provided by the machines of type m that are active in that period. This requirement is formally expressed as:

$$\sum_{\substack{p\in\mathcal{P}\\(m,p)\in\mathcal{I}}} \tau_{mp} \cdot y_{mpt} \le H_m \cdot a_{mt} \quad \forall\, m \in \mathcal{M},\ t \in \mathcal{T} \tag{C6}$$

Here, $\tau_{mp} \in \mathbb{R}_{>}$ denotes the time required to process one unit of p using m, while H_m indicates the number of hours available per technology of that type. For example, suppose two machines of type PBF-450 are active in period $t = 2$, and each provides $H_m = 150$ h of usable time. Then, the total time required to produce all compatible parts using this machine type must not exceed 300 h.

Machine operations require also human supervision. Let $\rho_m \in [0, 1]$ represent the fraction of processing time on machine type m that must be supervised by an operator. Let W_t denote the total number of available human hours in t, determined by the number of operators and their individual workloads. The model enforces that the total supervised processing time remains within the available workforce capacity:

$$\sum_{(m,p)\in\mathcal{I}} \rho_m \cdot \tau_{mp} \cdot y_{mpt} \le W_t \quad \forall\, t \in \mathcal{T} \tag{C7}$$

If machines of type MEX-100 require $\rho_m = 0.6$ and the planned production on these machines amounts to 100 h of processing in t, then 60 human hours are needed. If $W_1 = 200$, this leaves 140 h of supervision capacity for other machines.

Overall, the objective function (O1), along with constraints (C1)–(C7), defines an integer linear programming problem solvable with standard optimization tools. This formulation provides a clear foundation for the model but assumes full knowledge of all parameters, a condition seldom met in practice; the next section addresses this limitation by introducing uncertainty to support more risk-aware decisions.

4. A stochastic programming extension

To address the inherent variability of real-world manufacturing environments, we extend the linear deterministic model to a stochastic programming formulation. Here, several key parameters are modeled as continuous random variables rather than fixed quantities. We use a tilde notation to explicitly denote the uncertain nature of these parameters. For example, customer demand is represented as $\tilde{d}_{pt}$, acknowledging that order volumes may fluctuate due to changes in market conditions or customer behavior. Similarly, the processing time $\tilde{\tau}_{mp}$ accounts for variations in how long it takes to manufacture a part, which may depend on factors such as operator experience, material inconsistencies, or minor machine faults. Lastly, the total capacity $\tilde{H}_m$ available from machines of type m in each period incorporates the possibility of unplanned downtime, preventive maintenance, or efficiency losses, all of which can affect how much work a machine can realistically perform.

The corresponding constraints from the deterministic model are reformulated into their probabilistic counterparts, also known as *chance constraints* (Charnes & Cooper, 1959), which require that each condition be satisfied with a minimum specified probability. For instance, the demand satisfaction constraint (C5) becomes:

$$\mathbb{P}\left(\sum_{\substack{m\in\mathcal{M}\\(m,p)\in\mathcal{I}}} y_{mpt} \ge \tilde{d}_{pt}\right) \ge \alpha_{\text{demand}} \quad \forall\, p \in \mathcal{P},\ t \in \mathcal{T} \tag{C5-s}$$

Here, $\alpha_{\text{demand}} \in [0, 1]$ represents the minimum acceptable probability that the demand for part p in period t will be fulfilled.

Similarly, the machine capacity constraint (C6) is reformulated to account for the cumulative stochastic processing times required to manufacture the assigned parts. Rather than approximating the total processing time as the product of the expected processing time and the number of units produced, we explicitly model each unit q as an

independent realization of the random variable $\tilde{\tau}_{mp}$.[1] Accordingly, the total processing time becomes the sum of y_{mpt} independent random variables. The resulting chance constraint is expressed as:

$$\mathbb{P}\left(\sum_{\substack{p\in\mathcal{P}\\(m,p)\in\mathcal{I}}}\sum_{q=1}^{y_{mpt}}\tilde{\tau}_{mp}^{(q)} \le \tilde{H}_m \cdot a_{mt}\right) \ge \alpha_{\text{capacity}} \quad \forall\, m\in\mathcal{M},\ t\in\mathcal{T} \tag{C6-s}$$

Likewise, the human supervision constraint (C7) is also extended to incorporate the stochastic nature of processing times. Since each unit produced contributes independently to the total supervised workload, the constraint aggregates over all such units, weighted by the supervision ratio ρ_m. The probabilistic formulation is:

$$\mathbb{P}\left(\sum_{(m,p)\in\mathcal{I}}\sum_{q=1}^{y_{mpt}}\rho_m\cdot\tilde{\tau}_{mp}^{(q)} \le W_t\right) \ge \alpha_{\text{capacity}} \quad \forall\, t\in\mathcal{T} \tag{C7-s}$$

Although constraints (C6-s) and (C7-s) are expressed separately for clarity, they are inherently related: both ensure that the total processing workload—one in terms of machine time, the other in terms of operator time—remains within feasible limits in each period. For this reason, we combine their probabilistic requirements into a single joint capacity measure, defined as $\alpha_{\text{capacity}} \in [0,1]$.

At this stage, a key modeling decision concerns how to treat the threshold probabilities α_{demand} and α_{capacity}. In many applications, these values are set exogenously based on service-level agreements or industry standards (e.g., 95% or 99% service reliability) (Charnes & Cooper, 1959). However, such an approach may lead to over-conservative solutions, requiring costly investments in excess capacity to hedge against unlikely events. Instead, in our model we treat both α_{demand} and α_{capacity} as variables. This allows the optimization process to explore the trade-offs between minimizing costs and ensuring system robustness, helping decision-makers balance economic efficiency with risk tolerance. To this end, we reformulate the problem as a tri-objective optimization model. The first objective remains the minimization of total discounted acquisition cost, as expressed in (O1). The second and third objectives aim to maximize the likelihood that the demand is met and that the plant capacity is not exceeded. Specifically, we seek to:

$$\max\ f_2 = \alpha_{\text{demand}} \tag{O2}$$

$$\max\ f_3 = \alpha_{\text{capacity}} \tag{O3}$$

At the same time, to guarantee that the resulting solutions remain meaningful from an operational perspective, we introduce lower bounds $\underline{\alpha}_{\text{demand}}$ and $\underline{\alpha}_{\text{capacity}}$, representing minimum acceptable reliability levels. These thresholds, whose values can be specified depending on the application context, allow degenerate or operationally unacceptable solutions (e.g., zero-production policies) to be excluded from the trade-off analysis. They can also be set to zero when a full exploration of the feasible space is desired. These requirements are enforced in the model through the followings:

$$\alpha_{\text{demand}} \ge \underline{\alpha}_{\text{demand}} \tag{C8-s}$$

$$\alpha_{\text{capacity}} \ge \underline{\alpha}_{\text{capacity}} \tag{C9-s}$$

[1] Consider an optimization problem in which the uncertain quantity is the price $\tilde{p}$ per ton of a given material. Since each ton represents an identical good purchased at the same price, the total cost can be expressed as $\tilde{p}\cdot x$, where x is the number of tons. In contrast, in our setting, each p denotes a family of homogeneous but *non-identical* items that may exhibit different processing times. Accordingly, the processing time required to produce y_{mpt} units of p on m is modeled as:

$$\sum_{q=1}^{y_{mpt}}\tilde{\tau}_{mp}^{(q)},$$

which captures variability across individual realizations of the processing time.

Finally, the full stochastic optimization problem becomes:

$$\begin{aligned}
\text{minimize}\quad & \langle f_1, -f_2, -f_3\rangle \\
\text{subject to}\quad & \sum_{m\in\mathcal{M}} c_m\cdot x_{mt} \le B_t \quad \forall\, t\in\mathcal{T} \\
& a_{mt} = s_{mt} + \sum_{t'=\max\{1,t-L_m+1\}}^{t} x_{mt'} \\
& \sum_{m\in S_h} a_{mt} \le M\cdot b_{ht} \quad \forall\, h=1,\dots,q,\ t\in\mathcal{T} \\
& \sum_{h=1}^{q} b_{ht} \le u \quad \forall\, t\in\mathcal{T} \\
& \mathbb{P}\left(\sum_{\substack{m\in\mathcal{M}\\(m,p)\in\mathcal{I}}} y_{mpt} \ge \tilde{d}_{pt}\right) \ge \alpha_{\text{demand}} \quad \forall\, p\in\mathcal{P},\ t\in\mathcal{T} \\
& \mathbb{P}\left(\sum_{\substack{p\in\mathcal{P}\\(m,p)\in\mathcal{I}}}\sum_{q=1}^{y_{mpt}}\tilde{\tau}_{mp}^{(q)} \le \tilde{H}_m\cdot a_{mt}\right) \ge \alpha_{\text{capacity}} \quad \forall\, m\in\mathcal{M},\ t\in\mathcal{T} \\
& \mathbb{P}\left(\sum_{(m,p)\in\mathcal{I}}\sum_{q=1}^{y_{mpt}}\rho_m\cdot\tilde{\tau}_{mp}^{(q)} \le W_t\right) \ge \alpha_{\text{capacity}} \quad \forall\, t\in\mathcal{T} \\
& x_{mt}, a_{mt} \in \mathbb{Z}_{\ge} \quad \forall\, m\in\mathcal{M},\ t\in\mathcal{T} \\
& y_{mpt} \in \mathbb{Z}_{\ge} \quad \forall (m,p)\in\mathcal{I},\ t\in\mathcal{T} \\
& b_{ht} \in \{0,1\} \quad \forall\, h=1,\dots,q,\ t\in\mathcal{T} \\
& \underline{\alpha}_{\text{capacity}} \le \alpha_{\text{capacity}} \le 1 \\
& \underline{\alpha}_{\text{demand}} \le \alpha_{\text{demand}} \le 1
\end{aligned} \tag{1}$$

Fig. 2 shows a conceptual representation of the optimization problem (1). The framework produces a Pareto front of solutions trading off cost, service level, and feasibility under uncertainty. For practical purposes, three-dimensional Pareto fronts are sufficiently transparent for an expert to select a final solution. However, if this is not the case, a multi-attribute value theory model (Howard & Abbas, 2015) could be used to obtain a full ranking of the non-dominated solutions.

5. Algorithmic solution

Given the multi-objective and stochastic nature of the proposed model, the solution strategy is based on the NSGA-II algorithm (Deb et al., 2002), implemented via MATLAB's `gamultiobj` function.[2] NSGA-II is a population-based metaheuristic that evolves candidate solutions through selection, crossover, and mutation, while preserving the multi-objective structure of the problem via fast non-dominated sorting and elitism. A crowding-distance metric is used to maintain diversity and mitigate premature convergence. The algorithm is well suited to problems with discrete variables and stochastic components, as it does not rely on gradient information or convexity assumptions (Ma et al., 2023). It has been widely applied in resource allocation and supply chain optimization under uncertainty, further supporting its suitability for the present problem setting (Kabiri et al., 2022; Shi & Xiong, 2024).

Each individual in the population represents a complete candidate production and acquisition plan and is encoded as a single concatenated decision vector i, which contains both the acquisition variables (x_{mt}) and the production variables (y_{mpt}). During initialization, all decision variables are randomly generated within predefined discrete bounds. To avoid obviously infeasible starting points, production quantities are assigned only to machine–part pairs that are both technically compatible and operationally available in the corresponding time period. This procedure ensures that the initial population satisfies basic structural constraints and provides a diverse set of feasible solutions from which the evolutionary search can proceed. Each candidate solution is then evaluated through a custom fitness function. The first component of

[2] https://it.mathworks.com/help/gads/gamultiobj.html (retrieved on: 25/03/2026).

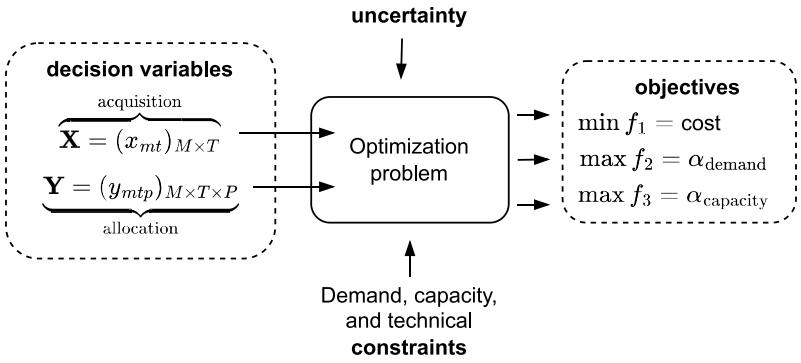

Fig. 2. Graphical representation of the proposed tri-objective stochastic optimization problem.

the function computes the total discounted acquisition cost according to (O1), using a constant discount rate over the planning horizon. The remaining components estimate the reliability of the solution with respect to uncertain demand and capacity. Rather than treating the target probabilities α_{demand} and α_{capacity} as fixed inputs, they are computed *ex post* as performance indicators that measure how often a given solution satisfies the stochastic constraints. To estimate these probabilities, we adopt a simulation-based procedure inspired by the Sample Approximation approach (Luedtke & Ahmed, 2008; Pagnoncelli et al., 2009). For each individual and at each generation, N_{sim} independent Monte Carlo scenarios are generated. In every scenario, customer demand ($\tilde{d}_{pt}$) and available machine capacity ($\tilde{H}_m$) are sampled once from their respective probability distributions, which may be specified analytically (e.g., Gamma or Normal) or derived empirically from historical data. Processing times are treated differently, as they capture variability at the level of individual produced units. For this reason, the processing time ($\tilde{\tau}_{mp}$) is sampled independently for each unit of part p assigned to machine type m, allowing different units of the same part to require different amounts of time.

Once all uncertain parameters are sampled for a given scenario, the candidate solution is checked against the stochastic constraints. Demand satisfaction is verified by comparing the total planned production with the sampled demand levels for each part and period. Capacity feasibility is assessed by summing the sampled processing times across all produced units and verifying that the total workload does not exceed the available machine capacity and human supervision resources. For each scenario, binary indicators record whether the demand and capacity constraints are satisfied. After all N_{sim} have been evaluated, the empirical probabilities α_{demand} (O2) and α_{capacity} (O3) are computed as the fraction of scenarios in which the corresponding constraints are met. To ensure that hard constraints are always respected, a feasibility check is embedded into the fitness evaluation. Candidate solutions that violate these structural constraints are penalized in their objective values, making them less likely to be selected for reproduction in subsequent generations.

This simulation-based evaluation framework provides a flexible and transparent way to incorporate real-world uncertainty into the optimization process without relying on closed-form expressions or restrictive distributional assumptions (Homem-de Mello & Bayraksan, 2014). It naturally supports mixed-integer decision variables, non-convex search spaces, and a wide range of probabilistic behaviors, including skewed and heavy-tailed distributions, which are commonly observed in industrial contexts where demand is irregular and processing times are strictly positive and asymmetric.[3] The full evaluation procedure is summarized in Algorithm 1.

[3] The Monte Carlo approach adopted here is agnostic with respect to the choice of distributional families, and different distributions can be used to model different sources of uncertainty.

Algorithm 1 Fitness Evaluation of a Candidate Solution

Input: Decision vector i, model parameters, N_{sim}
Decode i into acquisition matrix x_{mt} and production tensor y_{mpt}
Compute active machines a_{mt} based on initial inventory s_{mt} and lifetime L_m
Compute discounted acquisition cost according to (O1)
Initialize counters $n_{\text{demand}} \leftarrow 0$, $n_{\text{capacity}} \leftarrow 0$
for $k = 1$ to N_{sim} **do** ▷ Monte Carlo scenarios
  Sample demand $\tilde{d}_{pt}^{(k)}$ and capacity $\tilde{H}_m^{(k)}$
  Sample processing times $\tilde{\tau}_{mp}^{(q)}$ for all produced units
  if demand constraint (C5-s) is satisfied **then**
    $n_{\text{demand}} \leftarrow n_{\text{demand}} + 1$
  end if
  if capacity constraints (C6-s) and (C7-s) are satisfied **then**
    $n_{\text{capacity}} \leftarrow n_{\text{capacity}} + 1$
  end if
end for
Compute $\alpha_{\text{demand}} \leftarrow \frac{n_{\text{demand}}}{N_{\text{sim}}}$
Compute $\alpha_{\text{capacity}} \leftarrow \frac{n_{\text{capacity}}}{N_{\text{sim}}}$
Check hard constraints (budget, brand cardinality, compatibility, minimum service levels)
if constraints are violated **then**
  Assign penalty
end if
Return: Objective vector $f = [\text{Cost}+\text{Penalty}, -\alpha_{\text{demand}}+\text{Penalty}, -\alpha_{\text{capacity}}+\text{Penalty}]$

Following fitness evaluation, the offspring population is generated via crossover and mutation applied to selected parents. Crossover is performed using a uniform strategy in which each gene in the decision vector has a fixed probability of being inherited from either parent. Then, since this recombination may produce infeasible solutions (e.g., nonzero production on inactive or incompatible machines) a custom repair operator is applied to restore structural feasibility by adjusting production values based on machine availability. Genetic diversity is further enhanced through mutation, where each gene is randomly perturbed with a small probability. Acquisition variables are adjusted within predefined bounds, while production variables are mutated only if the associated machine is active and compatible.

At the end of the evolutionary process, the final population yields a diverse set of non-dominated solutions approximating the Pareto front. Each solution reflects a distinct trade-off among the three objectives (O1), (O2), and (O3).

6. Validation

The validation is conducted on a dataset designed to be both realistic and interpretable. This enables a controlled analysis of the model's behavior and facilitates a clear interpretation of the resulting decision structures. Specifically, we simulate a representative industrial

Table 2
Representation of unit processing times for each machine–part combination. Each entry reports the parameters of a Gamma distribution $\tilde{\tau}_{mp} \sim \Gamma(k, \theta)$ [min/unit], where k is the shape parameter and θ the scale parameter. Low mean values correspond to efficient combinations (✓), higher means to inefficient ones (—), and ✗ denotes incompatible combinations.

Machine	Support (mechanical) p_1	Housing (electronics) p_2	Aesthetic (design) p_3	Gear (test) p_4	Drone Case (lightweight) p_5	Prosthesis (custom-fit) p_6	Connector (pipe) p_7	Bracket (load-bearing) p_8
m_1	✓(5, 2.5)	✓(5, 2.5)	—(6, 4)	—(6, 4.5)	✓(5, 3)	—(6, 5)	✓(5, 2.5)	✗[a]
m_2	—(6, 4)	✓(5, 3.5)	✓(5, 3.5)	—(6, 5)	—(6, 5.5)	✓(5, 4)	—(6, 4.5)	✗[a]
m_3	✓(5, 2.2)	—(6, 3.8)	—(6, 3.8)	✓(5, 3)	✓(5, 3)	—(6, 4.5)	✓(5, 2.2)	—(6, 5)
m_4	✓(5, 1.5)	✓(5, 1.5)	—(6, 3)	✓(5, 2)	✓(5, 2)	—(6, 3.5)	—(6, 3)	—(6, 4)
m_5	—(6, 8)	—(6, 8)	—(6, 7)	—(6, 7)	✗[b]	—(6, 9)	—(6, 7)	✓(5, 6)
m_6	—(6, 2.5)	—(6, 2.5)	—(6, 3)	—(6, 3)	—(6, 2.5)	—(6, 3)	✓(5, 1.2)	✓(5, 1.5)
m_7	—(6, 3.5)	—(6, 3.5)	✓(5, 2.5)	✓(5, 2.5)	—(6, 3.5)	✓(5, 3)	—(6, 3.5)	—(6, 4)

[a] Plastic- and resin-based technologies are not suitable for producing structural metal brackets.
[b] Producing drone casings using this metal-based 3D printer is economically impractical.

scenario in the context of an AM facility operating as a prototyping hub, involving the acquisition and deployment of 3D printers over a planning horizon of three years. The decision process is structured across $T = 6$ consecutive periods, each corresponding to a six-month interval. The simulated environment includes $M = 7$ different AM technologies, ranging from well-established plastic-based systems to advanced metal printing solutions. Specifically, the considered portfolio includes Fused Deposition Modeling (FDM), Stereolithography (SLA), and Multi Jet Fusion (MJF) for polymer-based applications, as well as Direct Metal Laser Sintering (DMLS) and Binder Jetting for metal parts.

To reflect a diverse product mix, $P = 8$ different part types are considered, representative of common AM applications in both prototyping and functional production. Examples include a component for mechanical tests, a housing for electronic modules, and an aesthetic piece for industrial design. Some parts are intended for lightweight applications, such as the casing of a small drone, while others, like the medium-load structural bracket, are meant to be manufactured using stronger materials. Therefore, in line with practical production constraints, the machine–part compatibility matrix includes 3 hard exclusions based on material mismatch. To preserve modeling flexibility and avoid overly constraining the decision space, all other machines are assumed to be technically capable of producing every part, but their efficiency varies considerably: for instance, the bracket can be produced by Selective Laser Sintering (SLS), but much more efficiently by DMLS. This approach allows the model to implicitly discourage suboptimal technology–to-part assignments through longer production times, rather than through strict infeasibility. A summary of the assumed production times for each combination is provided in Table 2.

Although the proposed framework can accommodate different probability distributions, in this study the key stochastic parameters — namely $\tilde{d}_{pt}$, $\tilde{\tau}_{mp}$, and $\tilde{H}_m$ — are modeled using Gamma distributions. These are defined by only two parameters, making them both simple and flexible. Unlike the normal distribution, they are asymmetric and restricted to non-negative values, which makes them well suited for modeling variables such as demand and processing times, as discussed in Section 5. Compared to the exponential distribution, they allow for a non-zero mode. Moreover, they capture the possibility of rare but extreme adverse events, especially when the expected value is small. Under these assumptions, demand exhibits moderate variability across periods and part types, with average totals ranging from approximately 1.250 to 2.000 units per period. The peak occurs in t_5, which also shows the highest variability, highlighting the increased complexity of planning toward the end of the horizon. In contrast, t_1 has the lowest average demand and variability. Overall, the trend suggests a gradual ramp-up in production.

Since the planning horizon is three years, whereas industrial 3D printers typically operate over longer time spans, both acquisition costs and machine lifetimes are rescaled relative to commercial values. The reference data are reported in Appendix B, together with other typical characteristics of the considered technologies. This adjustment ensures consistency with availability and renewal dynamics within the planning horizon; otherwise, all machines would remain active after acquisition, eliminating the distinction between purchased and available units. Moreover, each machine belongs to a unique brand, and the production system is initialized with a small fleet of pre-existing 3D printers that are still within their usable life. Summary information on the main characteristics of the available technologies is provided in Table 3. Finally, human operators are assumed to be available on a full-time basis, and the budget in each period is scaled to represent the investment capacity of a mid-sized facility seeking to progressively expand its AM capabilities.

Table 3
Summary of machine-related parameters. Each entry of column s_{mt} reports the number of available machines and their remaining operational lifetime.

Machine	c_m [€]	L_m [periods]	s_{mt}	ρ_m
m_1	2.000	2	–	0.2
m_2	3.000	2	–	0.3
m_3	5.000	4	1×[2]	0.2
m_4	8.000	4	–	0.2
m_5	12.000	6	–	0.5
m_6	10.000	6	2×[1,3]	0.5
m_7	7.000	4	–	0.3

The configuration adopted in this case study is inspired by empirical observations from industrial AM facilities, capturing realistic parameter trends. However, since the primary purpose of this analysis is to demonstrate the functionality and flexibility of the proposed model, the lower bounds on the service levels, $\underline{\alpha}_{\text{demand}}$ and $\underline{\alpha}_{\text{capacity}}$, are set to zero, allowing for a full exploration of the feasible trade-off space.

To solve the problem, we adopt a reconfigured version of the NSGA-II algorithm, as detailed in Section 5. The evolutionary process is set with a population size of 200 and a maximum of 200 generations. The algorithm was run on an *MSI Vector 16 AI A2XWHG-200IT* laptop equipped with an *Intel Core Ultra 7 255HX* processor and 32 GB of RAM. The simulation-based evaluation was performed using $N_{\text{sim}} = 10.000$. The entire optimization process took approximately 47 min to complete.

Pareto fronts can be monitored in real time, and to assess convergence, we employ the hypervolume indicator, a well-established metric in multi-objective optimization that measures the volume of the objective space dominated by the current Pareto front (Guerreiro et al., 2021). The objective space is defined as $[0, \sum_{t\in\mathcal{T}} B_t] \times [0, 1] \times [0, 1]$, corresponding to the range of discounted acquisition cost and the feasible domains of the two probability-based objectives. This setup allows us to track, at each generation, the proportion of the normalized space covered by non-dominated solutions. A hypervolume value of zero indicates that no feasible solution has yet been found, while higher values reflect improved exploration of the objective space.

To assess the stability of the convergence behavior, we conducted a series of experiments in which both the $\tilde{d}_{pt}$ and $\tilde{\tau}_{mp}$ parameters were

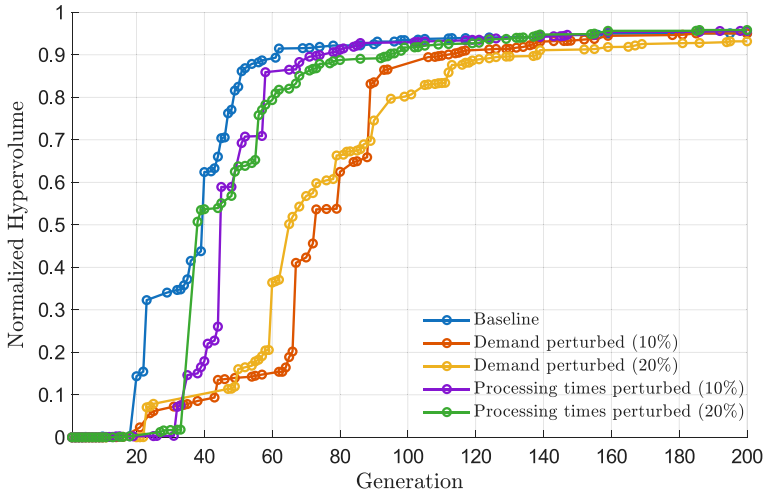

Fig. 3. Evolution of the hypervolume indicator across generations for the baseline instance and for perturbed scenarios.

Table 4
Qualitative summary of the objectives optimized in each region of the Pareto front. A ✓ indicates an emphasized objective, whereas a ✗ indicates one that is traded off.

Region	Cost (O1)	Demand (O2)	Capacity (O3)
(A)	✓	✗	✓
(B)	✗	✓	✓
(C)	✓	✓	✗

multiplicatively perturbed using random noise. Specifically, the perturbation factors ε were modeled as independent lognormally distributed random variables with unit expected value and standard deviation σ_ε. In a first series of tests, the scale parameters of the Gamma demand models were perturbed for all part–period pairs. In a second series, the nominal unit processing times were perturbed for all machine–part combinations. The noise intensity was varied in the range $\sigma_\varepsilon \in 0.10, 0.20$, and the corresponding results, together with the unperturbed baseline, are reported in Fig. 3. Across all instances, the hypervolume curves exhibit consistent growth patterns and stabilize within the first 140 generations, indicating that the observed convergence behavior is robust to moderate uncertainty in both demand and processing-time parameters and that the algorithm converges efficiently toward a well-formed Pareto front.

As illustrated in Fig. 4, the optimization yields a diverse set of non-dominated solutions spanning a wide range of trade-offs. Within this set, three meaningful regions can be identified in the objective space. Region (A) contains solutions with low demand reliability ($\alpha_{\text{demand}} < 0.5$) but high capacity feasibility ($\alpha_{\text{capacity}} > 0.5$). A representative case is the degenerate zero-decision solution, where no acquisitions or production are planned: cost is minimized and capacity constraints are trivially satisfied, but demand is never met. Region (B) represents the desirable operating zone in which both probabilities exceed 0.5. These solutions typically require higher investment but achieve reliable performance by jointly satisfying demand and capacity constraints; the top-right point corresponds to the most risk-averse configuration with $\alpha_{\text{demand}} = \alpha_{\text{capacity}} = 1$. Region (C), in contrast, includes solutions with good demand coverage but insufficient capacity reliability. Although economically attractive, these configurations rely on aggressive production plans that are unlikely to be feasible under stochastic capacity fluctuations. Overall, as summarized in Table 4, solutions in each region tend to emphasize two out of the three objectives.

To gain deeper insight into each region, we examine three representative solutions located at the extremes of the Pareto front—one from each region (A), (B), and (C) in Fig. 4. While the point at

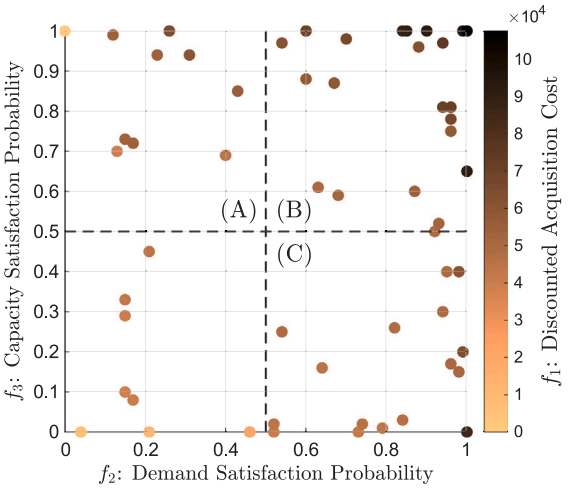

Fig. 4. Bidimensional representation of the Pareto front. Each point represents a non-dominated solution, plotted in terms of (O2) and (O3), with color indicating the corresponding acquisition cost (O1). The front is divided into three meaningful regions: (A), (B), and (C).

the top-left corner corresponds to the trivial zero-decision solution, it does not reflect the typical behavior of region (A). We instead select the adjacent point, which maintains high capacity feasibility while introducing minimal production. This is contrasted with the solution from the lower-right, representative of region (C), and the one from the upper-right corner, representative of region (B). From a methodological perspective, analyzing all three regions helps to clarify how the optimization framework explores extreme trade-offs between cost, demand satisfaction, and capacity feasibility. In practice, however, a production planner would typically focus on solutions within region (B), which achieve high levels of both α_{demand} and α_{capacity}. Finally, Fig. 5 shows the acquisition plans and active machine profiles for these solutions. The timeline includes an initial period t_0, representing the pre-existing inventory: 1 m_3 (SLS) with 2 remaining periods, and 2 m_6 (Binder Jetting) with 1 and 3 periods left. This comparison highlights the strategic differences in technology choices and capacity planning.

An important aspect common to all three solutions is the constraint limiting the number of active brands per period to 6, which prevents the simultaneous deployment of all technologies. The presence of pre-existing machines further steers acquisition decisions toward the same machine types, reflecting a preference for fleet standardization. As a

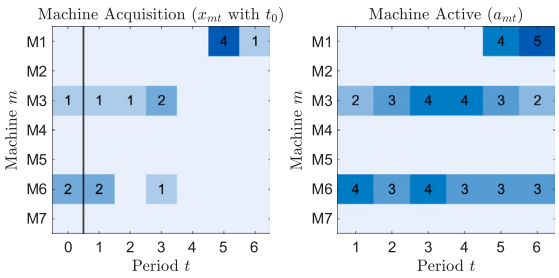

(a) Representative solution from region (A). Objectives: $\alpha_{\text{demand}} = 0{,}12 \quad \alpha_{\text{capacity}} = 0.99 \quad \text{Cost} = 53.382.$

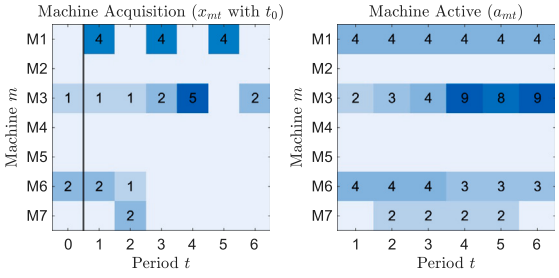

(b) Representative solution from region (B). Objectives: $\alpha_{\text{demand}} \approx 1 \quad \alpha_{\text{capacity}} \approx 1 \quad \text{Cost} = 107.579.$

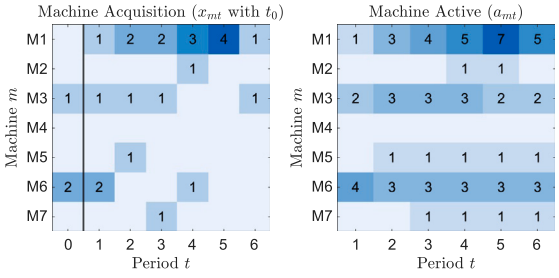

(c) Representative solution from region (C). Objectives: $\alpha_{\text{demand}} \approx 1 \quad \alpha_{\text{capacity}} = 0 \quad \text{Cost} = 85.706.$

Fig. 5. Machine acquisition plan (left) and resulting machine availability (right) for selected solutions from regions (A), (B), and (C). The fictitious time t_0 represents the initial inventory.

result, m_4 (MJF) is excluded in all cases. Although m_4 is not the most expensive option, it belongs to a higher-cost tier and is less attractive than more economical alternatives capable of producing the same parts with comparable speed (see Table 2) and similar human supervision requirements (refer to Table 3).

- The solution selected from region (A) prioritizes cost minimization and feasibility with respect to capacity constraints (Fig. 5(a)). This strategy relies primarily on technologies capable of producing all part types — specifically, m_3 and m_6 — with production assigned to the fastest available machines. For example, the structural bracket (p_8), which cannot be produced by m_1 (FDM) and is inefficient on non-metal printers, is assigned to m_6, which can also efficiently manufacture the pipe connector (p_7). However, this conservative strategy struggles to meet demand for certain components, particularly the drone casing (p_5), which exhibits high expected demand in t_2 and can be efficiently produced only by m_1, m_3, or m_4. Since m_1 is acquired later in the planning horizon, only three m_3 machines are active in t_2 and must also serve multiple other parts. To avoid exceeding capacity — especially given the high supervision requirements of m_6 — the planner underproduces p_5 (produced = 476, demand $\sim$ Gamma($k = 625,\ \theta = 0.8$)), resulting in a low value of α_{demand}.
- The solution selected from region (C) aims to maximize demand satisfaction while keeping acquisition costs relatively low (Fig. 5(c)). This is achieved by concentrating production on inexpensive technologies such as m_1, which performs efficiently on simple parts like functional supports (p_1) and p_7, but is less suitable for more complex components such as the aesthetic part (p_3). Although a wide range of technologies is activated to improve demand coverage, the resulting plans frequently exceed feasible capacity. For instance, in period t_5 the required human supervision (15.254 h) far exceeds the available 9.000 h. This solution is further characterized by the acquisition of m_5 (DMLS), an expensive technology that efficiently produces only p_8 and requires substantial infrastructure and operator supervision. As a result, total costs increase and the solution appears close to an outlier within region (C).
- Finally, the solution from region (B) follows an aggressive strategy that prioritizes the maximization of both probabilistic objectives at the expense of financial efficiency (Fig. 5(b)). This is achieved through systematic over-provisioning: a large number of machines are acquired to ensure high demand coverage and robust compliance with capacity constraints across all periods. The technology selection mirrors the logic of region (A) but intensifies acquisitions to improve reliability. Parts are assigned to the fastest available technologies; for example, p_1 is split between m_1 and m_3, the gear (p_4) is produced using m_3 and m_7 (Material Jetting), and p_8 is manufactured exclusively with m_6. Since demand variability is higher in the final periods, the strategy increases machine activations accordingly, while favoring low-supervision technologies such as m_3 to limit human resource usage.

If viewed from the perspective of chance-constrained optimization, our approach does not fully conform to either the conjunctive or disjunctive formulations traditionally adopted in the literature. Instead, it can be interpreted as a hybrid structure: stochastic constraints are categorized into two distinct families — one for demand satisfaction and another for capacity feasibility — and within each group, a conjunctive approach is applied. Notably, unlike conventional chance-constrained programming, our model does not require predefined probability thresholds; instead, the acceptance levels are treated as objectives (refer to Section 4). This allows the Pareto front to serve as a built-in sensitivity analysis, explicitly revealing how variations in budget allocation affect the likelihood of satisfying the stochastic constraints. However, the model can be reformulated into a fully conjunctive structure by aggregating the two probabilistic indicators into a single robustness metric, defined as $\alpha^* = \min\{\alpha_{\text{demand}}, \alpha_{\text{capacity}}\}$. This conservative interpretation assumes that a solution is only as reliable as its weakest probabilistic guarantee. By redefining the objective space as the set of points (f_1, α^*), the resulting Pareto front, shown in Fig. 6, captures the trade-off between acquisition cost and overall reliability against extreme events. This bi-objective view provides valuable insights: it quantifies how increasing investment in equipment can reduce the risk of operational failure under uncertainty, for instance, in scenarios characterized by high demand fluctuations or reduced machine availability. In this sense, the curve in Fig. 6 can be interpreted as a protection-level frontier, guiding risk-aware decision-making.

In our results, the curve increases almost linearly between $\alpha^* = 0.1$ and $\alpha^* = 0.9$, while further increases in acquisition cost beyond this range yield only marginal gains in probabilistic guarantees. This indicates that a robust solution can be identified at approximately $\alpha^* = 0.88$, with a corresponding total cost of 61.993.

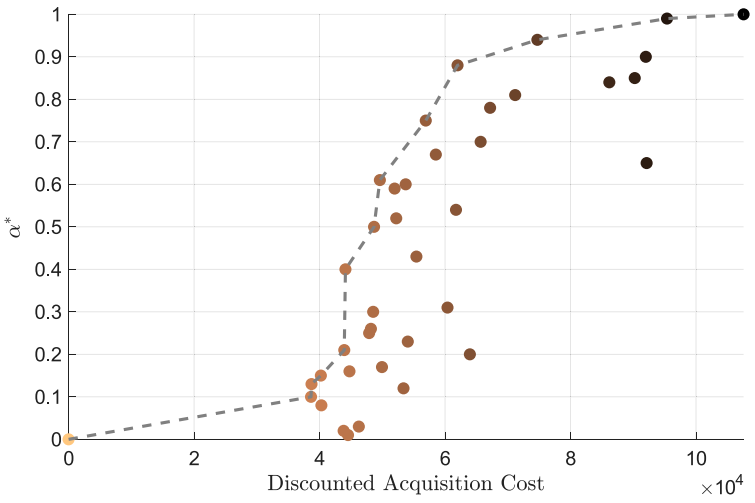

Fig. 6. Reformulated perspective of the Pareto front, showing the trade-off between discounted acquisition cost (O1) and overall robustness ($\alpha^* = \min\{\alpha_{\text{demand}}, \alpha_{\text{capacity}}\}$).

7. Discussion

The proposed framework is designed for implementation within a rolling horizon strategy, which has been successfully applied in various manufacturing and supply chain contexts. For example, Alimian et al. (2022) and Zhang et al. (2025) show that periodic re-optimization improves schedule quality and robustness in dynamic production and maintenance settings, while Herding and Mönch (2024) reports superior delivery performance and stability in semiconductor supply chains. These results suggest that a similar approach can enhance robustness in AM investment planning. In practice, long-term demand forecasts for specific part families are highly uncertain. For this reason, our setup adopts six-month planning periods rather than annual intervals, as shorter horizons yield more reliable forecasts. Accordingly, the model is intended to be re-solved periodically (e.g., at the end of each period $t \in \mathcal{T}$) to incorporate updated demand information and budget adjustments, enabling adaptive and responsive long-term planning.

A potential challenge in adopting the framework is the requirement for uncertain input data. However, this is well aligned with the digital nature of AM technologies, as modern 3D printers are typically equipped with monitoring systems that collect detailed operational data such as execution times, material usage, and process parameters. These data can be used to estimate probability distributions for key uncertain inputs. When considering the acquisition of new machines, distributions can be inferred by analogy with similar technologies already in operation. For example, if a facility has historical production data for a polymer extrusion system and evaluates a higher-throughput model from the same manufacturer, the processing-time distribution can be scaled according to relative deposition rates. This enables the framework to remain applicable even in exploratory investment scenarios with limited direct data. This perspective naturally supports integration with digital twin technologies (Jin et al., 2024). By combining real-time printer monitoring with rolling-horizon re-optimization, uncertainty distributions can be dynamically updated. For instance, if a polymer extrusion system is modeled with a Gamma distribution with mean 15 min but monitoring indicates an average closer to 12 min, the parameters can be recalibrated and incorporated into the next optimization cycle. In a scenario where 500 units of a standard connector are planned, this update may lead the model to recommend reallocating part of the workload to a compatible technology (e.g., MJF) or advancing the acquisition of an additional extrusion unit to improve demand fulfillment and capacity feasibility.

To operationalize the perspective discussed above, Fig. 7 presents a conceptual architecture that embeds the proposed optimization framework into a closed-loop decision support workflow. The scheme illustrates how real-world observations are transformed into model inputs, how optimization results are translated into actionable plans, and how execution outcomes feed back into subsequent planning cycles. Phase 1 represents the interface between the physical AM facility and its digital counterpart. Operational data — such as job logs, order inflows and backlog, and machine status — are collected and used to calibrate the probabilistic representations of demand $\tilde{d}_{pt}$, processing times $\tilde{\tau}_{mp}$, and available machine capacity $\tilde{H}_m$. In parallel, deterministic parameters describing the system configuration and decision space — such as machine lifetimes L_m, pre-existing inventory s_{mt}, workforce availability W_t, and economic coefficients (c_m, B_t, r) — are maintained and updated when contractual, technological, or financial conditions change. The digital twin thus maintains an up-to-date structural representation of system behavior. Phase 2 corresponds to the analytical core of the framework. The updated stochastic parameters, together with budget, compatibility, and brand constraints, define a rolling planning problem over the horizon $(t, \dots, T)$. Monte Carlo scenario generation and NSGA-II evolution produce a Pareto front capturing the trade-offs between discounted investment cost and probabilistic performance indicators. Phase 3 translates the set of non-dominated solutions into managerial action. Decision-makers filter and visualize the Pareto front — e.g., by focusing on Region B, bounding the robustness metric α^*, or enforcing economic targets — and implement the selected plan in the operational information systems, such as Enterprise Resource Planning (ERP) and Manufacturing Execution Systems (MES). The *Real-world operations* layer closes the loop: execution outcomes (e.g., realized processing times, machine utilization, and demand fulfillment) feed back into Phase 1, enabling periodic re-optimization in a rolling-horizon fashion. The computational cost associated with data processing is relatively low; aside from the optimization step, the only additional task is identifying the distributions that best fit the available data, which can be completed in a matter of seconds. The framework can therefore be interpreted as a continuously operating decision support system that integrates uncertainty modeling, optimization, and real-time operational data into an adaptive planning cycle.

The model was developed with the operational requirements of a publicly supported AM hub in mind. Such facilities often act as innovation enablers for local enterprises, providing consulting services and prototyping capabilities rather than operating for direct profit. Accordingly, the current formulation adopts a non-profit perspective, in which the budget for each planning period is specified exogenously. Nevertheless, the framework can be readily extended to profit-oriented contexts. This would require modifying the objective function (O1) to explicitly account for revenues and operating costs, and potentially allowing the available budget in period $t + 1$ to depend on the profit

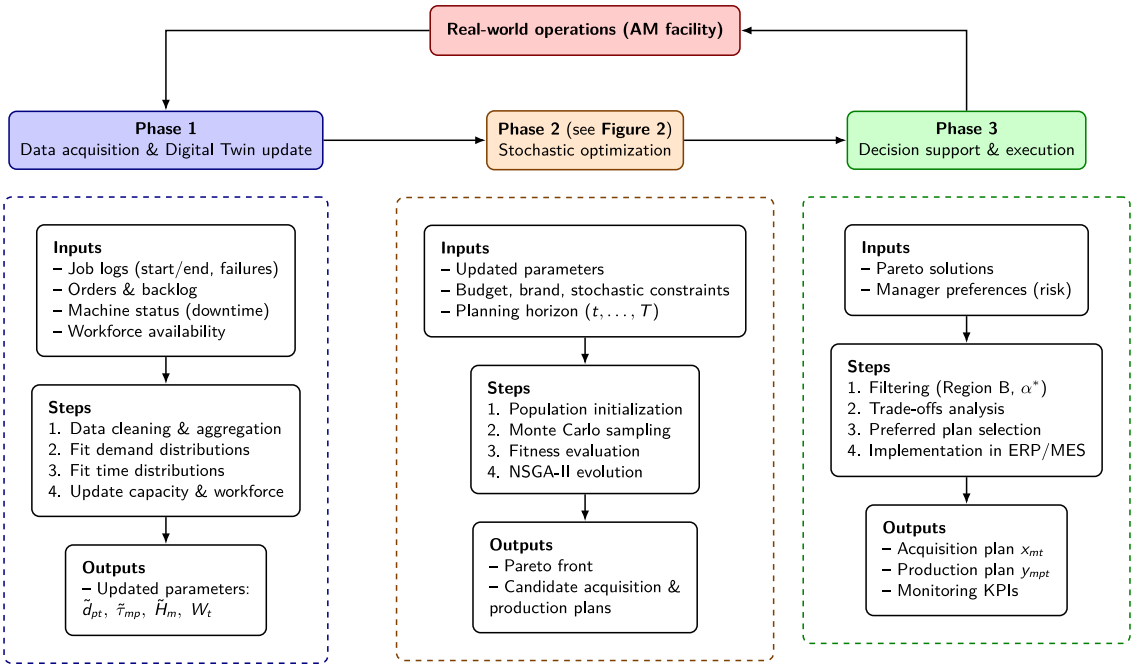

Fig. 7. Conceptual architecture of the proposed decision support workflow.

realized in period t. However, such an extension introduces additional modeling challenges. In the current formulation, only acquisition costs are explicitly considered, while operational expenses (e.g., energy consumption, materials, and maintenance) are not modeled. If future budgets were linked to estimated revenues without accounting for these costs, the optimization could favor unrealistically high production levels, leading to inefficient or infeasible solutions. Addressing this limitation would require incorporating a more detailed cost structure or introducing additional constraints to prevent overproduction.

Among the various operating costs that could be incorporated into the model, labor expenses are particularly significant. In the current formulation, workforce availability is represented by a fixed parameter W_t, which denotes the total number of operator hours available in period t. However, this parameter implicitly depends on both the number of technicians and their individual working capacities (refer to Section 3, constraint C7). By making this relationship explicit — e.g., defining $W_t = w \cdot z_t$, where $w > 0$ represents the number of hours provided per technician and $z_t \in \mathbb{Z}_>$ denotes the number of technicians employed in t — the model can be extended to include workforce sizing as a decision variable. The values of z_t can then be incorporated into (O1) through a per-operator cost term, thereby enabling the evaluation of trade-offs between investing in human resources and expanding machine capacity. This generalization underscores the modular nature of the proposed framework and its applicability to a broad class of resource allocation problems.

Overall, the optimization problem (1) is highly customizable and can be extended to accommodate additional decision variables and constraints. However, its computational complexity increases with the dimensionality of the problem. In scenarios involving larger sets of parts, technologies, or planning periods, the time required to reach convergence may become substantial. While this was not a limiting factor in the presented case study, more complex configurations — such as those featuring a greater number of (m, p) incompatibilities or a larger initial machine fleet s_{mt} — may benefit from a warm-start strategy to accelerate convergence. Preliminary tests indicate that the optimization process can be effectively initialized using a small set of diverse and feasible solutions. Specifically, we conducted a scalability experiment by simultaneously doubling the number of machine types (from 7 to 14) and part types (from 8 to 16). Three warm-start configurations were tested: (i) solving a relaxed deterministic version of the problem that considers only the cost objective (O1); (ii) solving a bi-objective variant that simultaneously optimizes (O1) and (O2), while temporarily omitting the probabilistic capacity constraint; and (iii) including a trivial solution in which all decision variables are set to zero, which satisfies (O1) and (O3) by construction. Although the warm start does not reduce the computational cost of a single fitness evaluation (as each evaluation still requires $N_{\text{sim}} = 10.000$), it significantly reduces the number of generations required to obtain a high-quality approximation of the Pareto front (about 100 instead of 200), as evidenced by earlier stabilization of the hypervolume indicator. Under the same computational conditions reported in Section 6, the total runtime increased from approximately 47 min to 2 h and 21 min for the larger-scale instance, confirming that the proposed warm-start strategy provides a simple and effective computational enhancement for higher-dimensional variants of the optimization framework.

A major computational bottleneck in the simulation-based fitness evaluation is the sampling of stochastic processing times. In the baseline formulation, and consistently with the probabilistic constraints in Section 4 (see (C6-s) and (C7-s)), a separate processing time $\tilde{\tau}_{mp}^{(q)}$ is sampled for each produced unit. While this approach accurately captures unit-level variability, it becomes computationally expensive for large production volumes. However, in many AM processes, machines operate on build jobs that include multiple parts produced simultaneously within a single build cycle (Framinan et al., 2023; Pinto et al., 2024). To improve computational efficiency, we therefore tested a batch sampling strategy using the same computational setup (population size = 200, maximum generations = 200, and $N_{\text{sim}} = 10.000$). Produced units are grouped into fixed-size batches of 10 units, and a single processing time is sampled per batch and assigned to all units within it. This reduces the number of random draws for the processing-time component by approximately one order of magnitude. In our experiments, this approach reduced the total runtime from 47 min to 13 min. The batch size thus acts as a tunable parameter that allows practitioners to balance computational efficiency and the level of detail in representing processing-time uncertainty, while approximating real-world production practices.

A further option to reduce computational time is the use of parallel computing. In MATLAB, the NSGA-II implementation in `gamultiobj` supports parallel execution through the `UseParallel` option, which distributes the evaluation of the fitness function across multiple workers (CPU cores) so that different individuals in the population are evaluated simultaneously. Compared to serial execution, parallel evaluation can yield a substantial reduction in runtime, approaching a factor proportional to the number of workers in ideal conditions, although the actual speed-up is typically lower because of scheduling and communication overhead. The benefits become more pronounced for larger problem instances or higher values of N_{sim}.

8. Conclusion

This study presented a stochastic multi-objective optimization framework for supporting long-term planning decisions in AM facilities. The proposed model addresses the challenge of selecting 3D printing technologies and allocating their production capacity under uncertainty, explicitly accounting for variability in both demand and processing capabilities. By combining a realistic representation of the AM decision-making context with Monte Carlo simulation and a tailored NSGA-II evolutionary algorithm, the framework enables the identification of robust and cost-efficient investment strategies. The modular structure of the model allows for a broad range of extensions. These include the incorporation of profit-driven objectives, endogenous budgeting mechanisms, and workforce sizing decisions, among others. Numerical experiments based on realistic data demonstrated the model's ability to capture the trade-offs between economic investment, service level robustness, and resource feasibility, providing valuable insights for decision-makers. Although the model was developed to support the optimal growth of AM hub facilities, its flexible structure makes it applicable to a wider range of strategic planning problems. In particular, it can be used as a general decision support tool for managing technology transitions and guiding the decommissioning of obsolete equipment.

In addition to providing high-quality solutions, the proposed optimization model can support managers in making more transparent and defensible decisions. At the same time, the current analysis is based on controlled computational experiments, and further work is required to bridge the gap toward real-world deployment, including data calibration and validation through case studies.

It is also worth noting that the proposed formulation includes original modeling features that preclude direct comparison with existing optimization frameworks. Notably, the use of integer decision variables as upper limits in the probabilistic constraints (C6-s) and (C7-s) introduces a level of structural complexity that requires a simulation-based evaluation. This unique aspect opens new algorithmic challenges that merit further investigation. One promising avenue for future research is the replacement of probability distributions with fuzzy numbers. Since the addition of fuzzy numbers preserves linearity, this approach could enable the formulation of fuzzy linear programming models that approximate uncertainty without requiring computationally intensive sampling procedures (Figueroa-García et al., 2022).

In addition to this, future work may explore several other directions. The assumption of independence among uncertain parameters adopted in this work is commonly used in stochastic optimization, particularly to ensure tractability in exploratory analyses, where modeling dependence would significantly increase complexity (Charnes & Cooper, 1959; Luedtke & Ahmed, 2008; Pagnoncelli et al., 2009). Nevertheless, refinements in the modeling of uncertainty could be introduced, including correlations across time periods or between demand and capacity variables. Copula-based methods (Joe, 2014) or multivariate time-series models (Lütkepohl, 2005) could be used to preserve temporal and cross-variable dependence, for example capturing situations in which high-demand periods coincide with reduced machine availability. This would allow the probabilistic indicators α_{demand} and α_{capacity} to reflect systemic risk rather than isolated sources of variability. Finally, the proposed approach could be embedded into a broader decision support system, enabling real-time updates and dynamic re-optimization in response to new information.

CRediT authorship contribution statement

Federica Tomelleri: Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Matteo Brunelli:** Writing – original draft, Visualization, Supervision, Resources, Project administration, Methodology, Formal analysis, Conceptualization.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

We gratefully thank the ProM Facility of Trentino Sviluppo3 and Weerg4 for sharing their technical expertise on additive manufacturing technologies and production constraints. Their feedback helped shape the modeling assumptions and ensure that the parameters adopted in this study are consistent with industrial practice.

Appendix A. Notation

The following notation is used throughout the mathematical formulation.

Sets.

- $\mathcal{M} = \{1, \dots, M\}$: Set of machine types (3D printer models)
- $\mathcal{P} = \{1, \dots, P\}$: Set of part families (product types)
- $\mathcal{T} = \{1, \dots, T\}$: Set of planning periods
- $\mathcal{I} \subseteq \mathcal{M} \times \mathcal{P}$: Set of feasible machine–part combinations
- $S_h \subseteq \mathcal{M}$: Subset of machine types corresponding to brand h

Variables.

- $x_{mt} \in \mathbb{Z}_{\geq}$: Number of machines of type m acquired at period t
- $y_{mpt} \in \mathbb{Z}_{\geq}$: Number of units of part p produced with machine m in period t
- $a_{mt} \in \mathbb{Z}_{\geq}$: Number of active machines of type m in period t
- $b_{ht} \in \{0, 1\}$: Binary variable; 1 if brand h is active in period t
- α_{demand}: Probability level for satisfying demand
- α_{capacity}: Probability level for satisfying capacity constraints

Parameters (deterministic).

- $s_{mt} \in \mathbb{Z}_{\geq}$: Number of pre-existing machines of type m still active at time t
- $L_m \in \mathbb{Z}_{>}$: Lifetime (in periods) of machine type m
- $c_m \in \mathbb{R}_{>}$: Acquisition cost of one unit of machine m
- $r \in \mathbb{R}_{>}$: Discount rate
- $B_t \in \mathbb{R}_{>}$: Budget available in period t
- $d_{pt} \in \mathbb{Z}_{\geq}$: Demand of part p in period t
- $\tau_{mp} \in \mathbb{R}_{>}$: Time to process one unit of part p with machine m
- $H_m \in \mathbb{R}_{>}$: Available production hours per machine of type m
- $\rho_m \in [0, 1]$: Fraction of processing time requiring human supervision for machine m
- $W_t \in \mathbb{R}_{>}$: Available human hours in period t
- $u \in \mathbb{Z}_{>}$: Maximum number of brands allowed per period
- $\underline{\alpha}_{\text{demand}} \in [0, 1]$: Minimum acceptable probability of satisfying demand
- $\underline{\alpha}_{\text{capacity}} \in [0, 1]$: Minimum acceptable probability of satisfying capacity constraints

Table B.5
Representative industrial datasheet characteristics for the 3D printers considered in this study. Values are indicative and reflect typical ranges reported in manufacturer specifications.

Technology	Build Volume [m^3]	Throughput[a]	Layer Thickness [μm]	Material Class	Acquisition Cost [€]
FDM	0008–0451	50–500 mm/s	100–400	Polymers, composites	2.000–25.000
SLA	0,0037–0309	20–700 mm/h	20–100	Photopolymers	20.000–80.000
SLS	0,0082–0154	300–1.000 cm^3/h	100–120	Polymers	100.000–200.000
MJF	~0,041	4.000–5.000 cm^3/h	~80	Polymers	150.000–250.000
DMLS	0,0008–0049	50–300 cm^3/h	20–80	Metals	300.000–500.000
Binder Jetting	~0,025	1.000–5.000 cm^3/min	35–200	Metals, ceramics	250.000–450.000
Material Jetting	0019–0400	100–500 mm/s	16–32	Polymers	120.000–220.000

[a] Throughput refers to the nominal build speed reported in machine datasheets and represents the rate at which material is deposited or consolidated during printing. Depending on the technology, it is expressed either as a linear deposition speed (mm/s or mm/h), a volumetric build rate (cm^3/h), or, for high-throughput powder-based systems, as a volumetric rate per minute (cm^3/min). In the proposed model, throughput serves as a proxy for the unit processing time parameter τ_{mp}, which is adjusted across machine–part combinations to reflect relative production efficiency.

Parameters (uncertain).

- $\tilde{d}_{pt}$: Random demand of part p in period t
- $\tilde{\tau}_{mp}^{(q)}$: Random processing time for the qth unit of part p on machine m
- $\tilde{H}_m$: Random total capacity (in hours) available for machine type m

Appendix B. Datasheet features of AM technology

See Table B.5.

Data availability

Data will be made available on request.

References

Aazami, A., & Saidi-Mehrabad, M. (2021). A production and distribution planning of perishable products with a fixed lifetime under vertical competition in the seller-buyer systems: A real-world application. *Journal of Manufacturing Systems*, *58*, 223–247.

Ahmed, R., Heese, H. S., & Kay, M. (2023). Designing a manufacturing network with additive manufacturing using stochastic optimisation. *International Journal of Production Research*, *61*, 2267–2287. http://dx.doi.org/10.1080/00207543.2022.2056723.

Akbari, M. (2023). Data-driven review of additive manufacturing on supply chains: Regionalization, key research themes and future directions. *Computers & Industrial Engineering*, *184*, Article 109600.

Algunaid, K. M. A., & Liu, J. (2022). Decision support system to select a 3D printing process/machine and material from a large-scale options pool. *International Journal of Advanced Manufacturing Technology*, *121*, 7643–7659. http://dx.doi.org/10.1007/s00170-022-09362-2.

Alimian, M., Ghezavati, V., Tavakkoli-Moghaddam, R., & Ramezanian, R. (2022). Solving a parallel-line capacitated lot-sizing and scheduling problem with sequence-dependent setup time/cost and preventive maintenance by a rolling horizon method. *Computers & Industrial Engineering*, *168*, Article 108041.

Aloui, A., & Hadj-Hamou, K. (2021). A heuristic approach for a scheduling problem in additive manufacturing under technological constraints. *Computers & Industrial Engineering*, *154*, Article 107115.

Altekin, F. T., & Bukchin, Y. (2022). A multi-objective optimization approach for exploring the cost and makespan trade-off in additive manufacturing. *European Journal of Operational Research*, *301*, 235–253. http://dx.doi.org/10.1016/j.ejor.2021.10.020.

Amorim, M., Antunes, F., Ferreira, S., & Couto, A. (2019). An integrated approach for strategic and tactical decisions for the emergency medical service: Exploring optimization and metamodel-based simulation for vehicle location. *Computers & Industrial Engineering*, *137*, Article 106057.

Charnes, A., & Cooper, W. W. (1959). Chance-constrained programming. *Management Science*, *6*, 73–79. http://dx.doi.org/10.1287/mnsc.6.1.73.

Deb, K., Pratap, A., Agarwal, S., & Meyarivan, T. (2002). A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, *6*, 182–197. http://dx.doi.org/10.1109/4235.996017.

Figueroa-García, J. C., Hernández, G., & Franco, C. (2022). A review on history, trends and perspectives of fuzzy linear programming. *Operations Research Perspectives*, *9*, Article 100247. http://dx.doi.org/10.1016/j.orp.2022.100247.

Framinan, J. M., Perez-Gonzalez, P., & Fernandez-Viagas, V. (2023). An overview on the use of operations research in additive manufacturing. *Annals of Operations Research*, *322*, 5–40. http://dx.doi.org/10.1007/s10479-022-05040-4.

Franco, D., Ganga, G. M. D., de Santa-Eulalia, L. A., & Godinho Filho, M. (2020). Consolidated and inconclusive effects of additive manufacturing adoption: A systematic literature review. *Computers & Industrial Engineering*, *148*, Article 106713.

Gao, Y., Yuan, B., & Cui, W. (2024). A math-heuristic approach for scheduling the production and delivery of a mobile additive manufacturing hub. *Computers & Industrial Engineering*, *188*, Article 109929.

Gibson, I., Rosen, D., Stucker, B., Khorasani, M., Rosen, D., Stucker, B., & Khorasani, M. (2021). *Additive manufacturing technologies: vol. 17*, Springer.

Guerreiro, A. P., Fonseca, C. M., & Paquete, L. (2021). The hypervolume indicator: Computational problems and algorithms. *ACM Computing Surveys*, *54*, 1–42. http://dx.doi.org/10.1145/3453474.

Herding, R., & Mönch, L. (2024). A rolling horizon planning approach for short-term demand supply matching. *Central European Journal of Operations Research*, *32*, 865–896.

Howard, R. A., & Abbas, A. E. (2015). Foundations of decision analysis. Pearson.

Jimo, A., Braziotis, C., & Pawar, K. (2025). Beyond prototyping: mapping the relative advantages of adopting additive manufacturing for industrial production. *International Journal of Production Research*, 1–30. http://dx.doi.org/10.1080/00207543.2025.2504165.

Jin, L., Zhai, X., Wang, K., Zhang, K., Wu, D., Nazir, A., Jiang, J., & Liao, W. H. (2024). Big data, machine learning, and digital twin assisted additive manufacturing: A review. *Materials & Design*, *244*, Article 113086.

Joe, H. (2014). *Dependence modeling with copulas*. CRC Press.

Kabiri, N. N., Emami, S., & Safaei, A. S. (2022). Simulation–optimization approach for the multi-objective production and distribution planning problem in the supply chain: using NSGA-II and monte carlo simulation. *Soft Computing*, *26*, 8661–8687.

Kanishka, K., & Acherjee, B. (2023). Revolutionizing manufacturing: A comprehensive overview of additive manufacturing processes, materials, developments, and challenges. *Journal of Manufacturing Processes*, *107*, 574–619. http://dx.doi.org/10.1016/j.jmapro.2023.10.024.

Lu, Z., Hu, K., & Ng, T. S. (2023). Improving additive manufacturing production planning: A sub-second pixel-based packing algorithm. *Computers & Industrial Engineering*, *181*, Article 109318.

Luedtke, J., & Ahmed, S. (2008). A sample approximation approach for optimization with probabilistic constraints. *SIAM Journal on Optimization*, *19*, 674–699. http://dx.doi.org/10.1137/070702928.

Luo, L., O'Hehir, J., Regan, C. M., Meng, L., Connor, J. D., & Chow, C. W. (2021). An integrated strategic and tactical optimization model for forest supply chain planning. *Forest Policy and Economics*, *131*, Article 102571.

Lütkepohl, H. (2005). *New introduction to multiple time series analysis*. Springer, http://dx.doi.org/10.1007/978-3-540-27752-1.

Ma, H., Zhang, Y., Sun, S., Liu, T., & Shan, Y. (2023). A comprehensive survey on NSGA-II for multi-objective optimization and applications. *Artificial Intelligence Review*, *56*, 15217–15270.

Makui, A., Heydari, M., Aazami, A., & Dehghani, E. (2016). Accelerating Benders decomposition approach for robust aggregate production planning of products with a very limited expiration date. *Computers & Industrial Engineering*, *100*, 34–51.

Marmarelis, M. G., & Ghanem, R. G. (2020). Data-driven stochastic optimization on manifolds for additive manufacturing. *Computational Materials Science*, *181*, Article 109750. http://dx.doi.org/10.1016/j.commatsci.2020.109750.

Homem-de Mello, T., & Bayraksan, G. (2014). Monte Carlo sampling-based methods for stochastic optimization. *Surveys in Operations Research and Management Science*, *19*, 56–85. http://dx.doi.org/10.1016/j.sorms.2014.05.001.

Pagnoncelli, B. K., Ahmed, S., & Shapiro, A. (2009). Sample average approximation method for chance constrained programming: theory and applications. *Journal of Optimization Theory and Applications*, *142*, 399–416. http://dx.doi.org/10.1007/s10957-009-9523-6.

Pinto, M., Silva, C., Thürer, M., & Moniz, S. (2024). Nesting and scheduling optimization of additive manufacturing systems: Mapping the territory. *Computers & Operations Research*, *165*, Article 106592. http://dx.doi.org/10.1016/j.cor.2024.106592.

Pourhejazy, P., Kravetc, T., & Sarkis, J. (2025). Performance evaluation of 3D print farms in additive manufacturing-based supply chains. *International Journal of Production Research*, 1–21. http://dx.doi.org/10.1080/00207543.2025.2532755.

Ransikarbum, K., Ha, S., Ma, J., & Kim, N. (2017). Multi-objective optimization analysis for part-to-printer assignment in a network of 3d fused deposition modeling. *Journal of Manufacturing Systems*, *43*, 35–46. http://dx.doi.org/10.1016/j.jmsy.2017.02.012.

Rüßmann, M., Lorenz, M., Gerbert, P., Waldner, M., Justus, J., Engel, P., & Harnisch, M. (2015). Industry 4.0: The future of productivity and growth in manufacturing industries. *Boston Consulting Group*, *9*, 54–89.

Sgarbossa, F., Peron, M., Lolli, F., & Balugani, E. (2021). Conventional or additive manufacturing for spare parts management: An extensive comparison for Poisson demand. *International Journal of Production Economics*, *233*, Article 107993. http://dx.doi.org/10.1016/j.ijpe.2020.107993.

Shi, S., & Xiong, H. (2024). Solving the multi-objective job shop scheduling problems with overtime consideration by an enhanced NSGA-II. *Computers & Industrial Engineering*, *190*, Article 110001.

Stittgen, T., & Schleifenbaum, J. H. (2021). Simulation of utilization for LPBF manufacturing systems. *Production Engineering*, *15*, 45–56. http://dx.doi.org/10.1007/s11740-020-00998-1.

Strong, D., Kay, M., Conner, B., Wakefield, T., & Manogharan, G. (2019). Hybrid manufacturing—locating am hubs using a two-stage facility location approach. *Additive Manufacturing*, *25*, 469–476. http://dx.doi.org/10.1016/j.addma.2018.11.027.

Tomelleri, F., Bosetti, P., & Brunelli, M. (2025). Optimizing 3D printer selection through multi-criteria decision analysis. *International Journal of Advanced Manufacturing Technology*, *139*, 3871–3890. http://dx.doi.org/10.1007/s00170-025-16148-9.

Yılmaz, Ö. F. (2020). Examining additive manufacturing in supply chain context through an optimization model. *Computers & Industrial Engineering*, *142*, Article 106335.

Zhai, C., Wang, J., Xu, J., Wang, B., & Tu, Y. (2025). Quality improvement and evaluation for profile responses in cloud-based additive manufacturing processes. *International Journal of Production Research*, *63*, 6411–6429. http://dx.doi.org/10.1080/00207543.2025.2472416.

Zhang, Y., Feng, Q., Fan, D., Ren, Y., Song, Y., Liu, M., & Wang, Z. (2025). Predictive control for operation and maintenance in smart manufacturing systems with multiple operating modes. *Computers & Industrial Engineering*, *207*, Article 111196.

Zheng, T., Ardolino, M., Bacchetti, A., & Perona, M. (2021). The applications of industry 4.0 technologies in manufacturing context: a systematic literature review. *International Journal of Production Research*, *59*, 1922–1954. http://dx.doi.org/10.1080/00207543.2020.1824085.