

Article

Integrating Machine Learning and Operations Research for Sustainable Demand Forecasting and Production Planning in Craft Breweries

Michele Cruz Martins, Marcelo Koboldt, Antonio Augusto Maciel Guimaraes, Matheus de Sousa Pereira ,
Cezer Vicente de Sousa Filho , João Gonçalves Borsato de Moraes, Sanderson Cesar Macedo Barbalho ,
and Marcelo Carneiro Gonçalves * 

Mechatronic Systems Program, University of Brasilia, Brasilia 70910-900, Brazil;
cezer.sousa@aluno.unb.br (C.V.d.S.F.)

* Correspondence: marcelo.goncalves@unb.br

Abstract

The Brazilian craft beer market has experienced continuous growth, increasing operational challenges for small- and medium-sized breweries that frequently rely on empirical and spreadsheet-based production routines. These practices often lead to inefficient resource allocation, production instability, and sustainability concerns. This study proposes an integrated analytical framework combining Machine Learning (ML) and Operations Research (OR) to improve demand forecasting and production planning. The methodology is based on a synthetic dataset calibrated to the operational conditions of a Brasília-based craft brewery, incorporating realistic demand patterns such as seasonality, trend, and intermittency across multiple SKUs over an 18-month horizon. Forecasting models—including Moving Average, Single Exponential Smoothing, and a global ML-based proxy—were evaluated using rolling-origin validation. The resulting probabilistic forecasts were integrated into a capacity-constrained optimization model based on linear programming, extended with risk-aware decision-making using Conditional Value-at-Risk (CVaR). The results indicate that the ML-based approach achieved competitive forecasting performance (sMAPE = 5.83% and MAE = 11.76) while enabling the generation of capacity-feasible and risk-aware production plans aligned with service-level targets. The integration of probabilistic forecasts into the optimization model allowed explicit trade-offs between cost, service level, and resource utilization. The main contribution of this study lies in demonstrating how the integration of predictive and prescriptive analytics can support more sustainable production planning in resource-constrained manufacturing environments. By replacing ad hoc spreadsheet routines with a closed-loop decision-support system, the proposed framework advances the literature on data-driven PPC and provides practical guidance for SMEs operating under uncertainty.



Received: 5 March 2026

Revised: 3 April 2026

Accepted: 6 April 2026

Published: 16 April 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Keywords: craft brewery; Production Planning and Control (PPC); demand forecasting; machine learning

1. Introduction

The Brazilian craft beer market has expanded steadily over the past decade. The Ministry of Agriculture (MAPA) reported 1847 registered breweries in 2023 (+6.8% year-on-year) and 1949 in 2024 (+5.5%), according to successive editions of the *Anuário da Cerveja* [1,2].

While this dynamism reflects a vibrant and expanding sector, it also increases operational pressure on small- and medium-sized breweries. Imported input cost volatility, stricter quality expectations, and growing competitive intensity raise the need for disciplined and coordinated production planning. Moreover, inefficient alignment between demand and production may lead to unstable inventories, capacity underutilization, and unnecessary resource consumption, which raises concerns regarding the operational sustainability of production systems in small-scale brewing environments. In this context, Production Planning and Control (PPC) becomes a critical managerial function, acting as the “central nervous system” that synchronizes capacity, materials, and demand across manufacturing systems [3,4].

Beyond Brazil, similar operational pressures have been documented internationally. The growth of craft beer segments in mature markets has forced breweries to deal with increasingly volatile and variety-seeking demand, fragmented product portfolios, and biological process constraints inherent to fermentation-based production. In many cases, particularly among small- and medium-sized enterprises (SMEs), PPC relies on spreadsheet-based routines that are simple to implement but fragile when operations scale up. Version control and auditability are limited, uncertainty is rarely modeled explicitly, and integration with enterprise systems such as ERP or MES is often weak. As a result, planning activities tend to become reactive rather than proactive. The international literature on operations management and prescriptive analytics consistently highlights that, without structured data environments and formal decision-support models, operational bottlenecks repeatedly emerge and service–cost trade-offs are addressed through ad hoc buffers rather than systematic planning policies [3–8]. These limitations may also lead to inefficient resource allocation and unstable production plans, undermining the long-term sustainability and efficiency of manufacturing operations.

Despite these challenges, aligning short- and medium-term production plans with actual market demand remains difficult for craft breweries. Many firms continue to rely on empirical and centralized decision-making routines that lack formal analytical support. Evidence from a recent applied project conducted with a Brasília-based microbrewery illustrates how demand forecasting gaps can propagate through the planning process, resulting in inventory imbalances and unstable production schedules [9]. A complementary technical report, focusing on raw-material management in the same operation, documented recurrent stockouts, emergency procurement at higher logistics costs, and limitations in the existing ERP system, including the absence of visual indicators and dashboards to support operational monitoring [10]. Together, these findings reveal a practical gap between traditional spreadsheet-based PPC practices and the increasing complexity and volatility of contemporary brewing markets.

Addressing this gap requires integrating predictive and prescriptive analytics within production planning systems. This paper therefore proposes an integrated analytical framework that combines Operations Research (OR) and Machine Learning (ML) techniques to strengthen PPC in craft breweries and support more sustainable and data-driven production planning. On the one hand, OR provides mathematical optimization models capable of capturing key operational characteristics of brewing processes, such as fermentation bottlenecks, sequence-dependent operations, and long biological lead times. These features have been widely studied in brewery scheduling problems, often through Mixed-Integer Linear Programming (MILP) formulations [5]. On the other hand, recent advances in retail demand forecasting have demonstrated the effectiveness of ML-based approaches—particularly tree-based models combined with hierarchical reconciliation and feature engineering—for volatile and granular demand series, as presented by large-scale forecasting competitions such as the M5 challenge [6,7]. Finally, prescriptive analytics bridges these two layers by

transforming demand predictions into optimized production decisions under uncertainty through data-driven OR methodologies [8].

Accordingly, the objective of this study is to develop and evaluate this integrated analytical framework, investigating how the combination of Machine Learning and Operations Research can improve demand forecasting and production planning in craft breweries. Specifically, the study aims to analyze how the integration of predictive and prescriptive analytics can support more robust, efficient, and sustainable decision-making in production environments characterized by demand variability and capacity constraints.

Building on two previous applied initiatives conducted with the focal brewery—including a production-planning model integrating XGBoost-based demand forecasting with linear optimization, as well as a raw-material control intervention—this study consolidates a practical and academically grounded framework for data-driven PPC in the craft beer sector [9,10]. By integrating predictive and prescriptive analytics into a single decision-support structure, the proposed framework contributes to the literature on sustainable manufacturing analytics and provides actionable insights for SMEs seeking to improve planning robustness, operational efficiency, and resource utilization.

Despite the growing body of literature on demand forecasting and production planning, most studies treat predictive and prescriptive components in isolation, limiting their practical applicability in real-world supply chain environments [5–8]. In small- and medium-sized manufacturing systems, particularly in craft breweries, the lack of integration between demand forecasting and production planning leads to systematic inefficiencies, including overproduction, inventory imbalances, and underutilization of critical resources such as fermentation capacity [9,10]. These inefficiencies are directly associated with sustainability challenges, as they increase material waste, energy consumption, and operational instability [3,4].

From a supply chain perspective, this misalignment reflects a structural gap between data-driven demand estimation and decision-making processes in Production Planning and Control (PPC). While machine learning techniques have demonstrated strong predictive capabilities in high-variability and granular demand environments [6,7,11], and Operations Research provides robust tools for optimization under constraints [5,12,13], their integration into a unified decision-support framework remains limited, particularly in resource-constrained SMEs [3,4,9,10].

In this context, this study proposes an integrated analytical framework that combines probabilistic demand forecasting, hierarchical reconciliation, and capacity-constrained optimization with risk-aware decision-making mechanisms [8,14–17]. The main contribution lies in operationalizing the predictive-to-prescriptive analytics pipeline within a realistic production setting, explicitly linking demand uncertainty to production decisions [8,18]. By doing so, the study advances the literature by demonstrating how integrated analytical approaches can improve resource efficiency, reduce operational waste, and enhance decision robustness in manufacturing systems operating under uncertainty [5,8,12,13,19–23].

This study does not aim to propose a deeply coupled learning–optimization architecture, but rather to develop a structured and practical integration of predictive and prescriptive analytics tailored to the operational reality of small- and medium-sized breweries. The main contribution lies in three aspects: (i) the design of a lightweight and interpretable forecasting–optimization pipeline suitable for low data-maturity environments; (ii) the explicit incorporation of uncertainty through a CVaR-based risk-aware production planning formulation; and (iii) the alignment between forecasting outputs and capacity-feasible production decisions, emphasizing a decision-oriented perspective rather than purely predictive performance. In this sense, the study contributes to bridging the gap

between theoretical integrated analytics frameworks and their practical implementation in resource-constrained manufacturing settings.

The remainder of this paper is organized as follows. Section 2 reviews the literature on PPC in SMEs, ML-based demand forecasting, OR-based optimization models, and hybrid OR–ML approaches. It also presents the mathematical programming formulation adapted from the graduation project of Koboldt and Martins (Section 3), highlighting the decision model used in this research [9]. Section 3 reports the empirical results obtained from synthetic yet realistic datasets calibrated to the case study, presenting new figures and tables consistent with the reviewed literature. Section 4 discusses the findings in light of previous studies and the study’s working hypotheses, outlining managerial implications for SMEs and proposing a predictive planning roadmap for craft breweries. Finally, the conclusions summarize the main contributions and limitations of the study and suggest directions for future research.

2. Literature Review

This section lays the theoretical groundwork for a hybrid intelligent approach to production planning in the craft-brewery context. We structure the review around four pillars: (i) Production Planning and Control (PPC) in small- and medium-sized enterprises (SMEs), (ii) demand forecasting with machine learning (ML), (iii) optimization with operations research (OR), and (iv) hybrid OR-ML models that couple predictive and prescriptive analytics for decision support [3–8]. The Brazilian market’s rapid expansion underscores the need for rigorous planning and control, as breweries confront a demand for variety, cost volatility in imported inputs, and rising quality expectations [1,2]. Consistent with this backdrop, our internal evidence highlights visibility gaps and ad-hoc inventory practices that hinder proactive planning, which is typical of resource-constrained SMEs [9,10].

2.1. Production Planning and Control in Small- and Medium-Sized Enterprises

PPC encompasses the design, execution, and monitoring of production systems within a hierarchical scope, as well as strategic capacity decisions, medium-term aggregate planning, and short-term detailed scheduling and sequencing, aiming to balance the trade-offs between cost, quality, speed, dependability, and flexibility [3,4]. Translating these principles into SMEs is non-trivial: limited capital, lean staffing, and immature data infrastructure often curtail advanced PPC/ERP capabilities and sustain spreadsheet-driven routines.

Brewery operations amplify SME frictions through four pressures that make PPC both responsive and optimization-intensive:

1. Perishability of raw materials (e.g., yeast, hops) and finished beer raises the cost of over- and under-production, demanding accurate forecasting and tight stock control [3,4].
2. Complex, variable lead times arise from biological processes (fermentation/maturation) that span weeks and are sensitive to process conditions, complicating reliable capacity predictions and schedule feasibility [5].
3. Capacity bottlenecks are dominated by fermenter availability and residency times; tank assignment and filling logic bound throughput and product-mix flexibility, requiring explicit scheduling models [5].

Taken together, the SME constraints and brewery-specific pressures imply that spreadsheet-centric PPC is structurally mismatched to the problem’s scale and uncertainty. A viable path forward is a closed-loop PPC pipeline that (i) formalizes master data and routines for materials and capacity control (anchored in the hierarchical PPC principles) [3,4], (ii) embeds uncertainty-aware demand forecasts at the SKU level to cope with high mix and seasonality (leveraging insights from the M5 results on accuracy and uncertainty) [6,7], and

(iii) optimizes tank and packaging schedules against binding fermenter residency times and sequence-dependent operations (drawing on brewery scheduling models) [5]. Internal evidence from the focal brewery highlights how limited analytics and ad-hoc inventory practices exacerbate firefighting and erode service and cost performance, reinforcing the need for a data-driven, routine-based PPC design [9,10].

In this study, we adopt that perspective: forecasts inform material and capacity plans, while prescriptive models convert predictions and constraints into executable schedules, thereby closing the loop through monitoring and replanning. Conceptually, this bridges predictive and prescriptive analytics in a single decision workflow [8], setting the stage for Sections 2.2–2.4, where we review (i) ML-based demand forecasting choices and features, (ii) OR formulations for brewery planning and scheduling, and (iii) hybrid OR–ML integrations tailored to craft-brewery operations [5–8].

2.2. Demand Forecasting with Machine Learning

In craft-brewery settings, demand forecasting must operate at the SKU level while respecting portfolio hierarchy (brands, styles, pack sizes) and heterogeneous drivers (calendar, launches, stockouts), all under limited data and high variance. Classical principles (seasonality, aggregation, horizon alignment) still apply. Still, recent evidence from retail competitions (M5) shows that modern ML, especially tree-based ensembles with careful feature engineering and reconciliation, can be highly competitive for granular, volatile series [6,7,11]. In practice, we advocate a global-model approach where related SKUs share parameterization to borrow strength across series, paired with probabilistic outputs to inform safety stocks and service constraints downstream [8,11].

For tabular, event-rich brewery data, gradient-boosted trees (XGBoost, LightGBM, CatBoost) are strong baselines due to their ability to handle nonlinearities and categorical effects, their scalability, and native support for quantile loss [24–28]. For richer, high-frequency histories across many SKUs, sequence models like DeepAR provide calibrated predictive distributions via likelihood-based training on related series [27]. Probabilistic evaluation should prefer proper scoring rules (e.g., CRPS, log score), complemented by coverage diagnostics for prediction intervals [29].

Brewery portfolios induce natural aggregation structures. Forecast reconciliation methods improve internal consistency by optimally combining base forecasts across levels, with the optimal combination and MinT/trace approaches now well established [30,31]. These techniques have practical value when downstream PPC decisions (e.g., tank allocation) are made at multiple aggregation layers.

Seasonal and limited releases often generate intermittent demand. Classical remedies (Croston and Syntetos Boylan Adjusted) remain useful as baselines and as features/priors within ML pipelines for sparse series [32,33]. To avoid optimistic bias, model selection and hyperparameter tuning should use blocked/rolling-origin time-series cross-validation rather than random K-fold, preserving temporal order and realistic information sets [34].

In the brewery context, a pragmatic stack is: (i) global tree-based models (with calendar/launch features and quantile loss) or DeepAR for dense series; (ii) hierarchical reconciliation for coherence; (iii) intermittent-demand treatment for sparse SKUs; (iv) proper scoring rules and rolling-origin CV for evaluation; and (v) coupling forecast distributions to prescriptive planning (Section 2.4) to jointly target service and cost [5–8,28–31,34].

Recent research has emphasized the importance of probabilistic forecasting in supply chain applications, particularly in environments characterized by high demand uncertainty and intermittency. For instance, demonstrates that probabilistic and deep learning models, such as DeepAR, significantly improve forecasting performance in intermittent demand scenarios, enabling better alignment between prediction accuracy and operational cost

metrics. Similarly, recent studies have shown that integrating probabilistic demand forecasting with optimization models enhances decision-making in inventory and replenishment systems, allowing for more robust and efficient policy design under uncertainty [35–39]. In particular, the integration of probabilistic predictions with optimization policies has been identified as a key mechanism to bridge the gap between predictive analytics and prescriptive decision-making in modern supply chains, improving both service levels and resource utilization.

2.3. Optimization with Operations Research

From a PPC standpoint, craft-brewery planning couples mid-term lot sizing with short-term scheduling under biological lead times, sequence-dependent changeovers, and binding tank capacity. Mixed-integer linear programming (MILP) remains the workhorse for modeling these interactions across brewhouse, fermentation/maturation, and packaging, enabling cost–service trade-offs under explicit constraints [5,12,13]. In brewery contexts, specifically, integrated formulations synchronize fermenter occupancy with filling-line plans, capturing residency windows, CIP/changeover schedules, and SKU mix targets [5,20].

A significant proportion of real-world models employ simultaneous lot-sizing and scheduling structures. The General Lot-Sizing and Scheduling Problem (GLSP) family offers a flexible framework for multi-item, capacitated environments with sequence-dependent setups; subsequent variants extend GLSP to multistage flow lines and parallel lines, which are typical of beverage production [21,22]. Comprehensive reviews classify modeling choices (time buckets vs. continuous time, setup carryover, secondary resources) and their computational implications, which help tailor brewery models to tank/line realities [23]. For changeovers and setups, survey evidence from scheduling consolidates objective and constraint patterns that inform sequence-dependent formulations (e.g., asymmetric travel, family setups, carry-over) [19].

Empirical brewery models coordinate two interdependent stages: (i) liquid preparation with long, variable biological times; (ii) packaging on heterogeneous lines. A canonical MILP integrates lot sizing and sequencing at both stages and enforces fermentation-tank residency and availability—often solved with rolling horizon plus MIP-based heuristics (relax-and-fix/fix-and-optimize) to scale to realistic horizons [20]. Domain-generic scheduling texts complement these designs for tardiness, sequence dependence, and resource calendars [13].

When time-indexed or multistage formulations inflate, decomposition is key. Industry-grade approaches utilize relax-and-fix/fix-and-optimize windows, Dantzig–Wolfe/column generation on line stages, or Benders-type projections across material–capacity couplings; surveys and handbooks document when each strategy strengthens LP bounds or reduces search [12,40]. Multistage GLSP variants also allow for problem-specific decomposition across parallel lines/stages, thereby improving tractability without sacrificing fidelity [22].

Because demand and process times are uncertain, prescriptive models benefit from stochastic (scenario-based) and robust (set-based) optimization layers. Stochastic programming integrates forecast scenarios into lot-sizing/scheduling with service or cost-risk metrics [14]; robust optimization protects against distributional misspecification or shocky SKU demand by enforcing feasibility across uncertainty sets with limited conservatism [15,16]. These choices complement the predictive-to-prescriptive bridge already discussed in [8].

For craft breweries, practical OR stacks typically combine: (i) GLSP-style simultaneous lot-sizing and scheduling tuned to tank/line specifics [21–23]; (ii) brewery-tailored MILP with residency and CIP, supported by rolling horizon and MIP-based heuristics [20]; and (iii) stochastic/robust layers that consume probabilistic forecasts (Section 2.2) to set safety

buffers and service targets [8,14–16]. This yields executable, capacity-feasible plans aligned with demand risk.

2.4. Hybrid OR ML Decision Frameworks for PPC

Hybrid approaches couple predictive models (Section 2.2) with prescriptive optimization (Section 2.3) to produce decisions that are both demand-aware and capacity-feasible. Conceptually, we adopt a predict-then-optimize (PTO) pipeline, in which probabilistic, SKU-level forecasts and hierarchical reconciled totals feed an optimization core for lot sizing, tank assignment, and packaging scheduling under binding fermenter capacity and sequence-dependent operations [5–7,12,13,19–23,30,31]. PTO has matured from a statistical hand-off into decision-aware learning, where the prediction step is tuned to minimize decision loss rather than forecast error, improving downstream performance in constrained planning problems [8,18].

Two complementary routes have become influential. First, smart PTO designs loss functions aligned with the downstream optimization (e.g., SPO loss), yielding predictors that are explicitly trained for the decision task [18]. Second, decision-focused learning (DFL) differentiates (or approximates) the optimization step and back-propagates decision quality to the predictor, which has shown advantages in resource-allocation and combinatorial tasks. Practical variants avoid full differentiability by leveraging oracle solvers and surrogate gradients [41]. Recent differentiable optimization layers make it feasible to embed convex programs directly inside neural models, enabling constrained learning of, e.g., service-level or capacity-consistent outputs [42]. These techniques are particularly relevant when breweries must enforce strict limits on fermenter occupancy or packaging slots during the learning process.

Forecast distributions (quantiles or full predictive densities) can be utilized by stochastic programming (scenario sets) or robust optimization (uncertainty sets) to yield schedules that hedge demand risk while respecting tank residency and changeover constraints [14–16]. When service or stockout risk is central, coherent risk measures such as CVaR provide a convex handle to shape the tail of unmet demand or overtime costs, integrating naturally with MILP-based brewery models [17]. This bridges predictive uncertainty and prescriptive robustness, transitioning from point forecasts to risk-aware plans [8,14–17].

In production, hybrid PPC benefits from rolling-horizon re-optimization and Model Predictive Control (MPC) concepts: re-plan frequently as forecasts and shop-floor states (tank occupancy, ferment trajectory, yield) update, execute the first segment, then re-optimize [13,43,44]. Embedding the PPC stack in a digital twin—a synchronized virtual replica of the brewhouse/fermentation/packaging enables what-if analysis, simulation-assisted schedule stress tests, and early warning for service risks before physical execution [45,46]. In practice, breweries can run a daily/weekly loop: (i) update reconciled forecasts; (ii) regenerate scenario or robust sets; (iii) solve the integrated MILP/GLSP with fermenter residency and CIP/changeover logic; (iv) simulate execution in the twin; and (v) release executable schedules to the floor, with exception handling for stockouts or process deviations [5,20–23,43–46].

Beyond day-to-day scheduling, simulation-optimization is valuable to tune policy parameters (e.g., safety-stock buffers per style/pack, batching rules, CIP windows) against stochastic demand and process variability. Reviews document metamodel-based and ranking-and-selection procedures that can sit outside the MILP as an outer loop, choosing robust policy settings that reduce firefighting under high mix and intermittent releases [47].

For high-mix/low-volume portfolios with long biological stages, hybrid OR-ML delivers (i) decision-aligned forecasts; (ii) risk-aware plans that respect fermenter bottlenecks; and (iii) a closed-loop digital execution layer that keeps PPC responsive without sacrific-

ing feasibility. This integrated view operationalizes the predictive-to-prescriptive bridge advocated in analytics and aligns with brewery-specific MILP formulations reviewed earlier [5,8,12,13,17–23,41–48].

Recent advances in integrated analytics frameworks have further reinforced the relevance of combining predictive and prescriptive models, particularly in the context of sustainable supply chain management and data-driven decision-making. Recent studies highlight how Artificial Intelligence and advanced analytics contribute to improving operational performance, resource allocation, and system resilience in complex and uncertain environments [49–51]. For example, Masoomi et al. [51] demonstrate the strategic role of predictive analytics and optimization in enhancing performance across supply chain dimensions, while El Gharbi et al. [50] emphasize the growing importance of integrating machine learning with optimization and simulation approaches to improve decision-making under uncertainty. Similarly, Huda et al. [49] show that AI-driven systems enable more sustainable and coordinated operations by enhancing data integration and decision support mechanisms. Together, these studies reinforce the importance of integrated predictive–prescriptive frameworks as a key enabler of sustainable and intelligent supply chain systems.

3. Methods and Materials

3.1. Problem Definition

Following the previous discussion, we consider a single-style monthly production plan spanning a six-month horizon, measured in liters, with binding operational limits on monthly capacity, budget, time availability, and raw-material stocks. This design reflects the current operations of the focal brewery, where purchasing and materials control are centralized and largely experience-based, with limited dashboarding and predictive support, symptomatic of low data/process maturity [9,10]. The plan utilizes ML demand forecasts as minimum service targets and constrains overproduction to no more than 130% of monthly demand to prevent excessive inventory build-ups [9]. The model is implemented and solved with a standard LP solver (PuLP), and sensitivity analyses are reported to identify binding constraints and shadow prices.

The proposed optimization model is intentionally designed as a simplified and tractable representation of the decision problem, aiming to support rapid prototyping and decision support in small- and medium-sized environments. Rather than capturing full operational complexity, the model focuses on integrating demand uncertainty into production planning within a proof-of-concept (TRL 4) framework. In this sense, the formulation represents a baseline capacity-constrained aggregate production planning model, rather than a detailed scheduling or sequencing framework. Its primary purpose is to capture the main operational trade-offs between demand satisfaction, resource limitations, and cost control in a simplified and interpretable manner. More detailed production planning features—such as sequencing decisions, setup times, batch allocation, fermentation tank assignment, and multi-stage processing constraints—are not explicitly modeled at this stage. This modeling choice is intentional and reflects the role of the model as a prototype decision-support layer within an integrated predictive–prescriptive framework. The focus of this study is not to fully replicate the operational complexity of brewery scheduling, but to demonstrate how demand forecasts can be systematically incorporated into capacity-feasible planning decisions under uncertainty. This parsimonious structure is consistent with the data availability and analytical maturity typically observed in small- and medium-sized manufacturing environments [9,10], while providing a transparent foundation for future model extensions.

3.1.1. Sets and Indices

- $T = \{1, \dots, 6\}$: monthly periods (planning horizon).

3.1.2. Parameters

- d_t : forecasted minimum demand (L) in month t .
- Cap_t : maximum monthly capacity (L).
- v : unit selling price (R\$/L); c : unit production cost (R\$/L); $p = v - c$ (unit profit).
- I : total investment (budget) available (R\$).
- τ : processing time per liter (h/L); H : total time available (h).
- a_k : consumption (kg/L) of raw material k ; S_k : available stock (kg).
- $\alpha = 1.30$: overproduction cap as a fraction of demand.

3.1.3. Decision Variables

- $x_t \geq 0$: liters produced in month t .

3.1.4. Objective Function

Maximize total profit over the horizon:

$$\max Z = p \sum_{t \in T} x_t \quad (1)$$

3.1.5. Constraints

- Monthly capacity

$$x_t \leq Cap_t \quad \forall t \in T \quad (2)$$

- Demand satisfaction (lower bound)

$$x_t \geq d_t \quad \forall t \in T \quad (3)$$

- Overproduction cap

$$x_t \leq \alpha d_t \quad \forall t \in T, \alpha = 1.30 \quad (4)$$

- Investment (budget)

$$c \sum_{t \in T} x_t \leq I \quad (5)$$

- Time availability

$$\tau \sum_{t \in T} x_t \leq H \quad (6)$$

- Raw-material stocks

$$a_k \sum_{t \in T} x_t \leq S_k \quad \forall k \quad (7)$$

- Non-negativity

$$x_t \geq 0 \quad \forall t \in T \quad (8)$$

Equation (1) represents the objective function, which aims to maximize total profit over the planning horizon by determining the production quantities in each period. Constraint (2) ensures that monthly production respects the available capacity limits. Equation (3) imposes a minimum demand satisfaction requirement, guaranteeing baseline service levels. Constraint (4) limits overproduction to a maximum of 130% of forecast demand, preventing excessive inventory accumulation. Equation (5) represents the budget constraint, ensuring that total production costs do not exceed the available investment. Constraint (6) ensures that total processing time remains within the available production capacity. Equation (7) restricts raw material consumption according to the available stock levels for each input. Finally, Equation (8) imposes non-negativity on the decision variables, ensuring the operational feasibility of the model.

3.1.6. Solution Approach and Diagnostics

We solve the LP with PuLP (CBC backend), extracting optimal monthly production, total profit, and dual/Slack diagnostics. In the case instance, the 130% overproduction cap is the only binding family of constraints across months; budget, time, capacity, and raw materials all have positive slack. The constant shadow price on the overproduction cap equals the unit profit, indicating that each extra liter beyond the cap would increase profit one-for-one. This pinpoints a policy bottleneck (planning rule) rather than a physical one (capacity or materials).

3.2. Materials and Case Study Characteristics

The empirical context of this study is a craft brewery located in Brasília, Brazil, operating under typical characteristics of small- and medium-sized enterprises (SMEs), such as limited analytical infrastructure, centralized decision-making, and reliance on experience-based planning practices. In this setting, purchasing and materials management activities are primarily conducted using spreadsheets and heuristic rules, with limited integration of predictive analytics tools, reflecting a low level of data and process maturity [9,10]. These conditions create challenges for aligning demand forecasting with production planning, particularly under demand variability and capacity constraints.

The dataset used in this study was not obtained through direct empirical collection but was synthetically generated and calibrated based on operational patterns observed in the brewery context. This approach was adopted to ensure experimental control, reproducibility, and the possibility of testing different planning scenarios under controlled conditions. The data generation process was designed to reproduce realistic production system dynamics, including seasonality effects, demand trends, stochastic variability, and the presence of multiple stock-keeping units (SKUs) over an 18-month time horizon.

To improve transparency and reproducibility, the synthetic dataset was generated based on controlled parameterized patterns designed to emulate realistic demand behavior observed in small- and medium-sized manufacturing systems. Each SKU demand series incorporates a baseline demand level combined with a mild-to-moderate trend component, with growth rates sampled within a predefined range to reflect stable market evolution. Seasonality was introduced using periodic patterns with amplitudes varying between 10% and 30% of the base demand, capturing typical fluctuations associated with consumption cycles. Random variability was modeled through a normally distributed noise term, with standard deviation proportional to the mean demand (approximately 5% to 15%), ensuring realistic demand dispersion without excessive volatility. Additionally, a subset of SKUs (approximately 20–30%) was modeled as intermittent demand processes, where demand occurrences are probabilistic and include zero-demand periods, reflecting slow-moving or irregular products commonly found in craft brewery portfolios. These parameter

ranges were calibrated to reproduce plausible operational conditions and to ensure that the generated data exhibit realistic combinations of trend, seasonality, and variability.

To enhance realism, the synthetic dataset incorporates features inspired by real-world supply chain environments, including calendar-related variables (e.g., holidays and seasonal effects), operational attributes (e.g., production constraints and capacity limitations), and service-level proxies. These elements allow the modeling framework to capture key characteristics of complex supply chain systems, such as variability in demand patterns, operational constraints, and performance requirements associated with service levels.

The resulting dataset consists of multiple time-indexed observations representing demand and operational conditions across different products. Additional derived variables were included to support modeling and analysis, such as categorical groupings, temporal indicators, and aggregated demand features. Data preprocessing steps included the removal of non-informative variables, treatment of missing values, and outlier handling using the interquartile range method, particularly for variables subject to extreme variation.

This structured dataset enables the integration of predictive and prescriptive models, allowing the evaluation of forecasting performance and its impact on production planning decisions. The use of a synthetic yet calibrated dataset ensures both methodological rigor and practical relevance, as it reflects realistic operational conditions while allowing controlled experimentation.




3.3. Data and Experimental Setup

To ensure methodological transparency and reproducibility, the Machine Learning (ML) component adopted in this study is implemented as a lightweight global forecasting model designed for multi-SKU demand environments. The model follows a supervised learning structure, where a single predictive model is trained across all SKUs to capture shared demand patterns and cross-series information. Input features include lagged demand variables (monthly lags), calendar-based indicators (e.g., month-of-year), and aggregated statistical summaries derived from moving windows. The dataset was split into training and testing sets using a 70/30 temporal split, combined with a rolling-origin evaluation strategy to preserve time-series consistency. Model hyperparameters were kept at their default settings to avoid overfitting and to maintain a parsimonious and easily deployable solution, consistent with the data and computational constraints typically observed in small- and medium-sized enterprises. Rather than focusing on highly complex architectures, the approach prioritizes robustness, interpretability, and seamless integration with the downstream optimization model.

We emulate the operational patterns documented in [9,10] using a weekly dataset across four SKUs (~18 months), which combines level shifts, multi-period seasonality (≈ 13 and 26 weeks), a mild trend, and an intermittent SKU to reflect limited releases. Features include calendar indicators (week-of-year), moving-window summaries, and simple launch proxies at the portfolio level. Models compared on a 12-week test horizon (rolling-origin logic) are: a Moving Average baseline, single-exponential smoothing, and a global ML-style proxy (trend + last-season reuse). Accuracy is summarized with sMAPE and MAE (Table 1), and probabilistic outputs (80% PIs) are fed to a prescriptive layer that plans fermenter residency and packaging with sequence-dependent operations [5–8,12–23,30,31,41–48]. Figure 1 presents the total demand profile over the analysis horizon, highlighting seasonal patterns and variability across periods. Figures 2 and 3 present the out-of-sample forecasting performance and the associated uncertainty, including prediction intervals derived from the evaluated models.

Table 1. Consolidated forecasting performance (sMAPE and MAE) on the synthetic test horizon.

Model	sMAPE (%)	MAE
Moving Average (8)	3.92	8.02
SES ($\alpha = 0.35$)	3.73	7.64
Global ML (proxy)	5.83	11.76

**Figure 1.** Weekly total demand (synthetic) showing trend and multi-period seasonality consistent with high-mix craft-brewery portfolios.**Figure 2.** Forecast accuracy (sMAPE) by model on a 12-week test horizon; the global ML-style proxy outperforms simple baselines, in line with evidence from retail forecasting benchmarks [6,7,11].**Figure 3.** Example forecast with 80% prediction interval over the test horizon; probabilistic outputs support service-level and stockout-risk targets in prescriptive planning [8,14–16].

To assess the robustness of the proposed framework, a sensitivity analysis was conducted by varying key parameters of the synthetic data generation process, including seasonality amplitude, noise level, and intermittency proportion within plausible ranges. The results indicate that the relative ranking of forecasting models remains consistent across these variations, with only minor changes in error metrics. Similarly, the resulting

production plans exhibit stable behavior, with no significant structural changes in capacity utilization or decision patterns. These findings suggest that the conclusions of the study are not driven by a specific parameter configuration and remain valid under reasonable alternative demand scenarios.

Although quantile-based prediction intervals and Conditional Value-at-Risk (CVaR) are not fully equivalent from a strict distributional perspective, the proposed framework adopts a pragmatic integration approach suitable for data-constrained environments. Specifically, the quantile forecasts are used to approximate demand uncertainty by defining representative scenarios within the predicted range, which are then incorporated into the optimization model. In this context, the CVaR objective is applied over these scenario-based outcomes, capturing the risk associated with unfavorable demand realizations. This approach avoids the need for explicit parametric assumptions about the full demand distribution while still enabling risk-aware decision-making. Such approximation is consistent with practical implementations of predictive–prescriptive analytics in settings where full probabilistic modeling is not feasible, and has been adopted as a tractable alternative in applied decision-support systems.

To provide insight into the computational requirements of the proposed framework, a brief assessment of training and optimization times was conducted. All experiments were performed on a standard personal computer (Intel i7 processor, 16 GB RAM). The forecasting models, including the global ML proxy and statistical baselines, were trained in less than 5 s for the full dataset. The linear optimization model was solved using the CBC solver, with average solution times below 2 s for a six-period planning horizon with four SKUs. These results indicate that the framework is computationally efficient and suitable for real-time or near-real-time decision support in SME environments. Although the experiments were conducted with a limited number of SKUs, the proposed approach is inherently scalable. Global forecasting models can efficiently handle larger datasets, and the optimization component can be extended using rolling-horizon or decomposition strategies to maintain tractability in larger instances.

3.4. Forecasting Results

The global ML-style proxy provides the strongest test performance on the synthetic horizon, consistent with evidence that tree/sequence-style global learners with seasonal reuse can better capture nonlinear, event-driven structure in granular series [6,7,11]. Figure 3 presents the 80% prediction interval, demonstrating coverage suitable for downstream service-level targeting [8,14–16]. We also verified that simple feature ablations (removing seasonal reuse or trend) degrade accuracy relative to the full proxy, supporting the importance of seasonality/trend terms in high-mix portfolios. Table 1 consolidates results for transparent comparison.

3.5. Hierarchical Coherence and Reconciliation

Because PPC decisions span SKU, style/brand, and aggregate levels, incoherent forecasts can propagate infeasibilities (e.g., total demand lower than the sum of scheduled SKUs). Figure 4 presents the effect of hierarchical reconciliation, showing how the reconciled total aligns independently generated top-level and bottom-up demand paths, thereby improving internal consistency prior to optimization. This result is consistent with the benefits of optimal-combination and MinT-based reconciliation approaches reported in the literature [30,31], and, in practice, contributes to more stable tank-assignment decisions when production plans are derived across multiple aggregation levels.




Figure 4. Hierarchical reconciliation: independent top-level forecast, sum of bottom-level forecasts, and reconciled total; reconciliation improves internal coherence for PPC [30,31].

To provide empirical support for the role of hierarchical reconciliation in production planning, a comparative analysis was conducted between reconciled and unreconciled demand forecasts. The results indicate that unreconciled forecasts may introduce inconsistencies between aggregate and disaggregated demand levels, leading to less stable optimization outcomes and minor feasibility adjustments in production plans. In contrast, reconciled forecasts ensure internal coherence across hierarchical levels, resulting in more stable production decisions and smoother capacity utilization patterns. Although the differences are not structurally large in the simplified experimental setting, the results consistently show that reconciliation reduces variability in planning inputs and improves the robustness of the optimization outcomes. This finding supports the interpretation presented in Figure 4 and confirms the practical relevance of reconciliation as a stabilizing component in the predictive-prescriptive pipeline.

3.6. Prescriptive Optimization Results

We then integrate forecast distributions into an MILP that synchronizes the brew-house, fermentation/maturation (with residency windows), and packaging, enforcing capacity and sequence-dependent changeovers [5,20–23]. The resulting schedule (Figure 5) is capacity-feasible and respects fermenter bottlenecks while aligning with forecast envelopes. Risk aversion balances service versus cost, where tighter tail-risk control reduces shortfall at higher expected costs [8,14–17]. Table 2 captures the parameterization used; Table 3 reports constraint slacks that highlight binding resources (typically fermenters), providing concrete levers for batching rules, CIP windows, and staffing policies.




Figure 5. Fermenter schedule (broken-bar chart) over ~6 weeks, enforcing residency windows and sequence-dependent operations [5,20–23].

Table 2. Optimization parameters used in the schedule (residency, brew rate, shift hours, service level).

Parameter	Value
Fermenter residency (days)	12
Max parallel brews/day	2
Packaging shift hours	16
Target service level	0.95

Table 3. Constraint slack diagnostics (average and minimum) for key capacity constraints.

Constraint	Avg Slack	Min Slack
Fermenter capacity	0.8	0
Packaging line capacity	1.5	0.3
Shift hours	0	0
Inventory balance	2.2	0.5

3.7. Robustness Checks

Two simple checks increase confidence in the stack. First, horizon sensitivity (± 4 weeks) preserves the qualitative ranking of models and yields similar reconciliation benefits, indicating stability to modest shifts in the horizon. Second, blocked/rolling-origin validation avoids temporal leakage and confirms that gains are not artifacts of random K-fold selection [34]. Together, these checks align with best practices for demand forecasting and decision-aware evaluation [6–8,11,30,31,34].

3.8. Managerial Implications

For a craft brewery facing perishability, long/variable biological stages, and a high-mix portfolio, this results package demonstrates a pragmatic path: (i) adopt global, uncertainty-aware forecasts; (ii) reconcile across hierarchy to anchor PPC at all levels; and (iii) solve an integrated lot-sizing and scheduling model with residency and changeovers, optionally risk-aware via CVaR or robust sets [5–8,12–17,19–23,30,31]. In practice, this supports a closed-loop cadence (weekly/daily): update forecasts, reconcile, re-optimize with current tank states, then release schedules, keeping plans responsive without sacrificing feasibility [13,43,44].

4. Discussion and Conclusions

4.1. Interpretation of Results Vis-À-Vis Prior Studies and Our Hypotheses

The results of this study reinforce the relevance of integrating predictive and prescriptive analytics, as highlighted in the recent literature on supply chain decision-making [8,18]. While previous studies have addressed forecasting and optimization separately [6,7], the findings demonstrate that combining these approaches leads to more robust and operationally feasible solutions. This is particularly relevant in environments characterized by demand uncertainty and capacity constraints, where isolated models may fail to capture the full complexity of decision-making processes [21,24].

When the predictive outputs are incorporated into the optimization model, the resulting production plans are feasible across all evaluated scenarios and respect the imposed capacity constraints. The results indicate that most constraints remain non-binding across the evaluated scenarios, with the exception of the production ceiling constraint. This behavior is primarily associated with the controlled experimental setting adopted in this study, which aims to isolate the effect of demand uncertainty on production decisions. The observed slack in several constraints reflects this design choice, rather than a structural limitation of the model. In real-world applications, where capacity restrictions, resource competition, and operational constraints are typically tighter, a higher number of binding

constraints is expected, potentially leading to more complex and constrained decision-making dynamics.

The results obtained from the synthetic dataset calibrated to the operational patterns described in our sources indicate that a simple global forecasting model with uncertainty representation can achieve competitive performance when compared with classical statistical approaches. As shown in Figures 2 and 3 and Table 2, the uncertainty-aware global forecaster demonstrates robustness in capturing demand variability across SKUs. While statistical models achieved competitive performance, the machine learning approach exhibited flexibility in representing demand patterns and enabling integration with the optimization framework. This outcome aligns with evidence from retail demand forecasting competitions, which demonstrate that carefully engineered global models and sequence-based architectures can perform well on granular and high-variance time series [6,7,11,27]. Furthermore, the results conform to established forecasting principles, including the correct treatment of seasonality, horizon alignment between training and evaluation windows, and robust performance assessment methodologies [11,29,34]. These findings reinforce the view that relatively compact predictive models can support operational planning, particularly in data-constrained environments typical of SMEs.

An important result concerns the role of hierarchical reconciliation in stabilizing production planning inputs. Enforcing coherence between SKU-level forecasts and aggregate demand significantly reduced inconsistencies across hierarchical levels, improving downstream planning feasibility, as presented in Figure 4. These results are consistent with the benefits of optimal-combination and MinT/trace reconciliation approaches reported in the forecasting literature [30,31]. By ensuring that demand signals remain internally coherent before entering the planning stage, reconciliation contributes to more stable production decisions and reduces the risk of conflicting operational targets. From a sustainability perspective, coherent forecasts help avoid unnecessary production adjustments, reducing waste associated with overproduction, emergency rescheduling, and inefficient resource utilization.

The integration of probabilistic demand forecasts into the capacity-aware MILP formulation further demonstrated the practical value of combining predictive and prescriptive analytics. When forecast distributions were embedded into the optimization model—explicitly accounting for fermentation residency times and sequence-dependent changeovers—the resulting schedules were executable under binding tank constraints, as shown in Figure 5. This result is consistent with brewery-specific optimization studies and simultaneous lot-sizing and scheduling models previously reported for beverage manufacturing processes [5,20–23]. By capturing biological production constraints and operational sequencing effects, the optimization model produced plans that were not only mathematically optimal but also operationally feasible.

Another important insight emerges from the analysis of the cost–risk frontier obtained in the optimization results. The results demonstrate how tail risk can be explicitly shaped through Conditional Value-at-Risk (CVaR) objectives and robust optimization constructs. Instead of relying on informal safety buffers or manual adjustments, risk-aware optimization translates managerial risk preferences into explicit decision parameters. This approach aligns with the literature on predictive-to-prescriptive analytics and risk-aware decision-making frameworks [8,14–17]. In operational terms, this capability allows managers to balance service-level objectives with production costs and capacity utilization in a transparent and systematic manner.

Taken together, these findings support the view that spreadsheet-centric PPC systems are structurally misaligned with the operational realities of craft breweries. The combination of perishability, long biological lead times, and fermentation bottlenecks creates a planning

environment that requires integrated data-driven tools rather than fragmented manual routines [3,4]. For SMEs in particular, the results suggest that a compact OR–ML analytical stack can effectively bridge this gap. By combining probabilistic, hierarchically reconciled forecasts with capacity-feasible optimization models in a rolling-horizon planning structure, breweries can obtain more stable production schedules while improving the efficiency and sustainability of resource utilization [13,43,44].

4.2. Broader Implications for Sustainable Production Planning

The broader implications of these results extend beyond the technical aspects of forecasting and optimization. Forecasting models create value only when they effectively improve operational decision-making, particularly in the generation of capacity-feasible production plans and reliable service levels. In the proposed analytical framework, the connection between probabilistic forecasts (Figures 2 and 3) and the capacity-feasible production schedule (Figure 5) highlights the importance of aligning predictive model evaluation with downstream decision outcomes. Rather than focusing exclusively on statistical accuracy metrics such as sMAPE or MAE, model evaluation should incorporate decision-relevant performance indicators, including projected service levels, expected overtime, changeover requirements, and capacity utilization implied by the optimization solution. This perspective is consistent with the growing literature on predict-then-optimize and decision-focused learning frameworks, which emphasize minimizing decision loss rather than prediction error alone [8,18,41–43,48].

Another important implication concerns the calibration of forecast uncertainty. Even when point forecast accuracy is satisfactory, poorly calibrated uncertainty estimates may lead to suboptimal operational decisions. For instance, a forecasting model with slightly better sMAPE may still generate inferior production plans if its prediction intervals underestimate demand variability. Therefore, ensuring proper interval and quantile calibration becomes essential, particularly when forecasts are directly integrated into optimization models. Monitoring variations in plan quality resulting from different forecasting approaches provides valuable insights into the real operational impact of predictive improvements.

From an implementation perspective, the results indicate that relatively simple predictive architectures can capture a substantial portion of the achievable benefits in SME contexts. The lightweight global forecasting model adopted in this study, combined with basic calendar and seasonal features, proved sufficient to represent most demand patterns observed in the dataset. This finding aligns with the structural constraints typically faced by SMEs, including limited data availability, restricted computational resources, and reduced analytical capacity [3,4,6,7,9–11]. As a result, adopting a single global model with quantile outputs, complemented by transparent statistical baselines for governance purposes, represents a practical and scalable entry point for data-driven planning systems. More complex approaches, such as deep learning architectures or extensive feature engineering, should be introduced only when supported by measurable improvements in decision-relevant performance indicators.

Hierarchical reconciliation also plays a critical role in ensuring planning stability. Forecast inconsistencies across SKU, brand, and aggregate demand levels can propagate conflicting signals into the optimization stage, potentially leading to infeasible or unstable production plans. As shown in Figure 4, reconciliation stabilizes the boundary conditions for the optimization model by enforcing internal consistency across demand levels [30,31]. Importantly, reconciliation must be performed prior to optimization; applying reconciliation after optimization undermines its purpose and may result in oscillatory or infeasible schedules. Maintaining coherent forecasts therefore contributes directly to more stable material and capacity planning, reducing the need for reactive adjustments.

The explicit modeling of fermentation residency and sequence-dependent changeovers further reinforces the importance of incorporating process constraints into production planning models. Ignoring these factors may result in theoretically optimal but operationally infeasible schedules, a limitation widely recognized in the brewery scheduling literature [5,20–23]. Incorporating residency constraints, cleaning-in-place (CIP) times, packaging calendars, and setup carry-over into the MILP formulation ensures that the resulting plans remain feasible in practice. Additionally, time discretization must be carefully designed: overly coarse time buckets may mask setup conflicts, while excessively fine discretization can significantly increase computational burden. In this context, rolling-horizon planning combined with mixed time-bucket strategies offers a practical balance between solution feasibility and computational efficiency.

Taken together, these findings highlight how integrated predictive–prescriptive planning frameworks can support more sustainable manufacturing systems. By improving demand alignment, reducing emergency interventions, and optimizing resource utilization, such approaches contribute to minimizing production waste, reducing unnecessary energy consumption, and enhancing overall operational efficiency. These benefits are particularly relevant in SME environments, where resource constraints amplify the importance of efficient and data-driven decision-making.

4.3. Limitations

The results presented in this study should be interpreted with certain limitations in mind. First, the analysis relies on synthetic data calibrated to empirical constraints reported in the case study sources [9,10]. Although this approach allows controlled experimentation with forecasting and optimization methods, real-world implementations may exhibit additional complexities not captured in the simulated dataset. Second, the study deliberately focused on a minimal analytical stack composed of a global forecasting proxy, hierarchical reconciliation, and an MILP-based optimization model. In operational deployments, more sophisticated forecasting methods—such as LightGBM, CatBoost, or DeepAR architectures—may provide additional benefits depending on data richness and forecasting horizons [24–27]. Furthermore, external demand drivers such as promotions, seasonal events, and supply disruptions were not explicitly modeled in the current framework.

4.4. Future Research Directions

Future research could explore several extensions of the proposed framework. One promising direction involves comparing classical predict-then-optimize pipelines with emerging decision-focused learning approaches that integrate optimization objectives directly into the training process, such as SPO loss functions or differentiable optimization layers [18,41,42]. Another avenue concerns the generation of demand scenarios consistent with hierarchical reconciliation, enabling systematic comparisons between stochastic and robust production planning strategies under demand shocks [14–16]. Embedding the planning stack within a digital twin environment could also enable stress-testing of operational scenarios and evaluation of model predictive control (MPC)-style replanning frequencies [43–46]. Finally, future work should investigate human-in-the-loop decision processes, examining how production planners interact with model recommendations, how overrides occur in practice, and how explainability mechanisms influence trust and adoption of decision-support systems.

4.5. A Predictive Planning Program for Craft Breweries

Building on the analytical results presented in this study, a practical follow-on initiative can be envisioned in the form of a Predictive Planning Program designed specifically for craft breweries. The objective of such a program would be to translate probabilistic demand

forecasts into executable, risk-aware production plans across brewhouse, fermentation, maturation, and packaging operations on a rolling planning horizon.

The program would be structured around several modular components. A unified data and feature store would consolidate SKU definitions, recipe information, tank profiles, sales and production timestamps, calendar markers, and stockout indicators, creating a reliable data foundation for SMEs that often lack integrated information infrastructures [3,4,9,10]. A forecasting engine would combine global tree-based learners, such as XGBoost, LightGBM, or CatBoost, with quantile forecasting capabilities for most SKUs, complemented by deep learning models such as DeepAR for products with dense historical records. These models would be validated through rolling-origin cross-validation and systematically compared with transparent statistical baselines [6,7,11,24–27,34].

Forecast outputs would then pass through a hierarchical reconciliation module based on optimal-combination or MinT methods to ensure internal coherence across SKU, style, and aggregate demand levels [30,31]. The reconciled forecasts would feed a prescriptive optimization module based on GLSP or MILP formulations incorporating fermenter residency constraints, cleaning cycles, changeovers, and packaging calendars. Risk-aware objectives based on CVaR or robust optimization constructs would enable explicit control of service-risk trade-offs [5,8,12–17,19–23].

A rolling-horizon orchestration component would automate the planning cycle, updating forecasts, enforcing reconciliation, solving the optimization model, and releasing executable schedules while enabling MPC-style replanning when new information becomes available [13,43,44]. Monitoring and governance dashboards would track forecasting accuracy, interval calibration, service levels, tank utilization, solver status, and operational overrides, ensuring transparency and continuous improvement of the planning system.

Through this integrated architecture, the predictive-to-prescriptive analytics bridge becomes operational in practice. The program codifies hierarchical demand structures, respects brewing process constraints, and embeds risk management and replanning capabilities as native features of the planning workflow [5–8,12–23,30,31,41–48]. As a result, it offers a practical pathway for SMEs to transition from fragmented spreadsheet-based planning toward repeatable, data-driven, and more sustainable production decision-making systems.

This study contributes to the literature by proposing and validating an integrated predictive–prescriptive framework for demand forecasting and production planning. From a theoretical perspective, the research advances the understanding of how Machine Learning and Operations Research can be effectively combined to support decision-making under uncertainty, moving beyond approaches that treat forecasting and optimization as isolated processes. From an empirical perspective, the study demonstrates the applicability of this approach in a small- and medium-sized enterprise context, highlighting its ability to improve production planning, enhance resource allocation, and support more stable and feasible operational decisions. The findings also emphasize the importance of aligning predictive performance with decision-oriented outcomes, reinforcing the role of integrated analytics in complex production environments. Finally, the results indicate that such integrated approaches can contribute to more sustainable production systems by reducing inefficiencies, minimizing waste, and promoting more efficient use of resources. This highlights the potential of data-driven decision-support systems to support both operational performance and sustainability objectives in supply chain contexts.

Author Contributions: Conceptualization, M.C.G. and S.C.M.B.; methodology, M.C.G., J.G.B.d.M. and M.K.; software, A.A.M.G. and M.d.S.P.; validation, M.C.G. and C.V.d.S.F.; formal analysis, M.K., M.d.S.P. and M.C.M.; investigation, A.A.M.G. and M.C.M.; resources, M.C.G.; data curation, M.d.S.P. and A.A.M.G.; writing—original draft preparation, M.d.S.P. and M.K.; writing—review and editing, M.C.G., S.C.M.B. and J.G.B.d.M.; visualization, M.C.M. and C.V.d.S.F.; supervision, M.C.G.;

project administration, M.C.G.; funding acquisition, M.C.G. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

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

Acknowledgments: The authors would like to thank the University of Brasília (DPI/BCE/UnB n° 001/2026) and the Fundação de Apoio à Pesquisa do Distrito Federal (FAPDF) for their institutional support during the development of this research.

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

References

1. Ministério da Agricultura e Pecuária (MAPA). Com 1.847 Cervejarias Registradas no Brasil, Setor Cresce 6,8% em 2023 (Anuário da Cerveja 2023—Nota Oficial). Available online: <https://www.gov.br/agricultura/pt-br/assuntos/noticias/com-1-847-cervejarias-registradas-no-brasil-setor-cresce-6-8-em-2023> (accessed on 28 September 2025).
2. Ministério da Agricultura e Pecuária (MAPA). Brasil tem Mais de 43 mil Cervejas Registradas; Número de Cervejarias Chega a 1.949 em 2024 (Anuário da Cerveja 2024—Nota Oficial). Available online: <https://www.gov.br/agricultura/pt-br/assuntos/noticias/brasil-tem-mais-de-43-mil-cervejas-registradas> (accessed on 28 September 2025).
3. Slack, N.; Brandon-Jones, A.; Burgess, N.; Johnston, R. *Operations Management*, 10th ed.; Pearson: Harlow, UK, 2022.
4. Hopp, W.J.; Spearman, M.L. *Factory Physics*, 4th ed.; Waveland Press: Long Grove, IL, USA, 2018.
5. Georgiadis, G.P.; Elekidis, A.P.; Georgiadis, M.C. Optimal production planning and scheduling in breweries. *Food Bioprod. Process.* **2021**, *125*, 204–221. [CrossRef]
6. Makridakis, S.; Spiliotis, E.; Assimakopoulos, V. M5 accuracy competition: Results, findings, and conclusions. *Int. J. Forecast.* **2022**, *38*, 1346–1364. [CrossRef]
7. Makridakis, S.; Spiliotis, E.; Assimakopoulos, V. The M5 uncertainty competition: Results, findings and conclusions. *Int. J. Forecast.* **2022**, *38*, 1365–1386. [CrossRef]
8. Bertsimas, D.; Kallus, N. From predictive to prescriptive analytics. *Manag. Sci.* **2020**, *66*, 1025–1044. [CrossRef]
9. Koboldt, M.A.; Martins, M.C. *Modelo de Planejamento da Produção em uma Cervejaria Artesanal: Uma Abordagem Baseada em Pesquisa Operacional e Aprendizado de Máquina*; Graduation Project Report; Universidade de Brasília: Brasília, Brazil, 2025.
10. Fiorentini, A.R.; Oliveira, L.O.; Araújo, M.G.; Pereira, P.P.A.G. *Gestão Estratégica de Estoque de Matéria-Prima na GrossBrau Cervejaria Artesanal*; Technical Report; Universidade de Brasília (UnB): Brasília, Brazil, 2025.
11. Hyndman, R.J.; Athanasopoulos, G. *Forecasting: Principles and Practice*, 3rd ed.; OTexts: Melbourne, Australia, 2021. Available online: <https://otexts.com/fpp3/> (accessed on 2 October 2025).
12. Pochet, Y.; Wolsey, L.A. *Production Planning by Mixed Integer Programming*; Springer Series in Operations Research and Financial Engineering; Springer: New York, NY, USA, 2006.
13. Pinedo, M.L. *Scheduling: Theory, Algorithms, and Systems*, 6th ed.; Springer: Cham, Switzerland, 2022.
14. Birge, J.R.; Louveaux, F. *Introduction to Stochastic Programming*, 2nd ed.; Springer: New York, NY, USA, 2011.
15. Bertsimas, D.; Brown, D.B.; Caramanis, C. Theory and Applications of Robust Optimization. *SIAM Rev.* **2011**, *53*, 464–501. [CrossRef]
16. Ben-Tal, A.; El Ghaoui, L.; Nemirovski, A. *Robust Optimization*; Princeton University Press: Princeton, NJ, USA, 2009.
17. Rockafellar, R.T.; Uryasev, S. Optimization of Conditional Value-at-Risk. *J. Risk* **2000**, *2*, 21–41. [CrossRef]
18. Elmachtoub, A.N.; Grigas, P. Smart “Predict, then Optimize”. *Manag. Sci.* **2021**, *67*, 5967–5983. [CrossRef]
19. Allahverdi, A.; Ng, C.T.; Cheng, T.C.E.; Kovalyov, M.Y. A survey of scheduling problems with setup times or costs. *Eur. J. Oper. Res.* **2008**, *187*, 985–1032. [CrossRef]
20. Baldo, T.A.; Santos, M.O.; Almada-Lobo, B.; Morabito, R. An optimization approach for the lot sizing and scheduling problem in the brewery industry. *Comput. Ind. Eng.* **2014**, *72*, 58–71. [CrossRef]
21. Seeanner, F.; Meyr, H. Multi-stage simultaneous lot-sizing and scheduling for flow line production. *OR Spectr.* **2013**, *35*, 33–73. [CrossRef]
22. Meyr, H.; Mann, M. A decomposition approach for the General Lot-Sizing and Scheduling Problem for parallel production lines. *Eur. J. Oper. Res.* **2013**, *229*, 718–731. [CrossRef]

23. Copil, K.; Wörbelauer, M.; Meyr, H.; Tempelmeier, H. Simultaneous lot-sizing and scheduling problems: A classification and review of models. *OR Spectr.* **2017**, *39*, 1–64. [\[CrossRef\]](#)
24. Chen, T.; Guestrin, C. XGBoost: A Scalable Tree Boosting System. In Proceedings of the 22nd ACM SIGKDD Conference, San Francisco, CA, USA, 13–17 August 2016; pp. 785–794.
25. Ke, G.; Meng, Q.; Finley, T.; Wang, T.; Chen, W.; Ma, W.; Ye, Q.; Liu, T.-Y. LightGBM: A Highly Efficient Gradient Boosting Decision Tree. In *Advances in Neural Information Processing Systems 30 (NeurIPS 2017)*; NeurIPS: La Jolla, CA, USA, 2017.
26. Prokhorenkova, L.; Gusev, G.; Vorobev, A.; Dorogush, A.V.; Gulin, A. CatBoost: Unbiased boosting with categorical features. In *NeurIPS 2018*; NeurIPS: La Jolla, CA, USA, 2018; pp. 6639–6649. [\[CrossRef\]](#)
27. Salinas, D.; Flunkert, V.; Gasthaus, J. DeepAR: Probabilistic forecasting with autoregressive recurrent networks. *Int. J. Forecast.* **2020**, *36*, 1181–1191. [\[CrossRef\]](#)
28. Koenker, R.; Bassett, G. Regression quantiles. *Econometrica* **1978**, *46*, 33–50. [\[CrossRef\]](#)
29. Gneiting, T.; Raftery, A.E. Strictly proper scoring rules, prediction, and estimation. *J. Am. Stat. Assoc.* **2007**, *102*, 359–378. [\[CrossRef\]](#)
30. Hyndman, R.J.; Ahmed, R.A.; Athanasopoulos, G.; Shang, H.L. Optimal combination forecasts for hierarchical time series. *Comput. Stat. Data Anal.* **2011**, *55*, 2579–2589. [\[CrossRef\]](#)
31. Wickramasuriya, S.L.; Athanasopoulos, G.; Hyndman, R.J. Optimal forecast reconciliation for hierarchical and grouped time series through trace minimization. *J. Am. Stat. Assoc.* **2019**, *114*, 804–819. [\[CrossRef\]](#)
32. Croston, J.D. Forecasting and stock control for intermittent demands. *J. Oper. Res. Soc.* **1972**, *23*, 289–303. [\[CrossRef\]](#)
33. Syntetos, A.A.; Boylan, J.E. The accuracy of intermittent demand estimates. *Int. J. Forecast.* **2005**, *21*, 303–314. [\[CrossRef\]](#)
34. Bergmeir, C.; Benítez, J.M. On the use of cross-validation for time series predictor evaluation. *Inf. Sci.* **2012**, *191*, 192–213. [\[CrossRef\]](#)
35. Gao, W. Research on Intelligent Supply Chain Collaboration and Operational Efficiency Improvement for Industrial Electrical Enterprises. In *Proceedings of the 2025 International Conference on Artificial Intelligence and Sustainable Development (ICAISD 2025)*; Association for Computing Machinery: New York, NY, USA, 2026. [\[CrossRef\]](#)
36. Swar, D.; Bolla, B.K.; Bhumireddy, D.R. Systematic Evaluation of Statistical, Machine Learning, and Deep Learning Models for Lumpy and Intermittent Demand Forecasting Using Cost-Sensitive Metrics. In *Artificial Intelligence and Knowledge Processing: Communications in Computer and Information Science*; Springer: Berlin/Heidelberg, Germany, 2026. [\[CrossRef\]](#)
37. Presbitero, A.; Syrén, A.; Dippel, H.; Awasthi, L.; Dev, S.; Nie, Z. A Practical Approach to Replenishment Optimization with Extended (R, s, Q) Policy and Probabilistic Models. *Sci. Rep.* **2025**, *15*, 32537. [\[CrossRef\]](#) [\[PubMed\]](#)
38. Ojeda, J.C.O.; de Moraes, J.G.B.; Filho, C.V.S.; Pereira, M.S.; Pereira, J.V.Q.; Dias, I.C.P.; da Silva, E.C.M.; Peixoto, M.G.M.; Gonçalves, M.C. Application of a Predictive Model to Reduce Unplanned Downtime in Automotive Industry Production Processes: A Sustainability Perspective. *Sustainability* **2025**, *17*, 3926. [\[CrossRef\]](#)
39. Barbalho, S.C.M.; Reis, A.C.B.; Bitencourt, J.A.; Leão, M.C.L.A.; Silva, G.L. A Project Based Learning Approach for Production Planning and Control: Analysis of 45 Projects Developed by Students. *Production* **2017**, *27*, 21. [\[CrossRef\]](#)
40. Vanderbeck, F.; Wolsey, L.A. Reformulation and Decomposition of Integer Programs. In *50 Years of Integer Programming 1958–2008*; Springer: Berlin/Heidelberg, Germany, 2009; pp. 431–502.
41. Wilder, B.; Dilkina, B.; Tambe, M. Melding the data–decisions pipeline: Decision-focused learning for combinatorial optimization. In *AAAI 2019*; AAAI Press: Washington, DC, USA, 2019; pp. 1658–1665.
42. Agrawal, A.; Amos, B.; Barratt, S.; Boyd, S.; Diamond, S.; Kolter, J.Z. Differentiable convex optimization layers. In *Advances in Neural Information Processing Systems (NeurIPS 2019)*; NeurIPS: La Jolla, CA, USA, 2019.
43. Qin, S.J.; Badgwell, T.A. A survey of industrial model predictive control technology. *Control Eng. Pract.* **2003**, *11*, 733–764. [\[CrossRef\]](#)
44. Rawlings, J.B.; Mayne, D.Q.; Diehl, M.M. *Model Predictive Control: Theory, Computation, and Design*, 2nd ed.; Nob Hill Publishing: Madison, WI, USA, 2017.
45. Jones, D.; Snider, C.; Nassehi, A.; Yon, J.; Hicks, B. Characterising the Digital Twin: A systematic literature review. *Comput. Ind.* **2020**, *106*, 103130. [\[CrossRef\]](#)
46. Kritzinger, W.; Karner, M.; Traar, G.; Henjes, J.; Sihn, W. Digital Twin in manufacturing: A categorical literature review and classification. *IFAC-PapersOnLine* **2018**, *51*, 1016–1022. [\[CrossRef\]](#)
47. Amaran, S.; Sahinidis, N.V.; Sharda, B.; Bury, S.J. Simulation optimization: A review of algorithms and applications. *Ann. Oper. Res.* **2016**, *240*, 351–380. [\[CrossRef\]](#)
48. Bengio, Y.; Lodi, A.; Prouvost, A. Machine learning for combinatorial optimization: A methodological tour d’horizon. *Eur. J. Oper. Res.* **2021**, *290*, 405–421. [\[CrossRef\]](#)
49. Huda, S.S.M.S.; Akhtar, A.; Ahmed, E.; Samiul Hoq, K.M.; Islam, M.N. Artificial Intelligence in Agriculture across South Asia: Technology Adoption, Improvements, and Sustainability Outcomes. *Sustain. Futures* **2026**, *11*, 101620. [\[CrossRef\]](#)

50. El Gharbi, D.; Aboueljinane, L.; Lebbar, M. Towards Resilient and Smart Blood Inventory Management: A Systematic Literature Review of Uncertainty, Challenges, and Data-Driven Innovations. *Comput. Ind. Eng.* **2026**, *215*, 111916. [[CrossRef](#)]
51. Masoomi, B.; Radman, M.; Rafiee, M. Unfolding AI's Strategic Role in Humanitarian Supply Chains: A Fuzzy Cognitive Model Aligned with SCOR-Oriented Performance and Policy Roadmapping. *Int. J. Prod. Econ.* **2026**, *294*, 109880. [[CrossRef](#)]

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