ISBN: 978-628-96613-3-0. ISSN: 2414-6390. Digital Object Identifier: https://dx.doi.org/10.18687/LACCEI2026.1.1.1410

Simulation-Based Feasibility Assessment Under Demand Uncertainty: Linking Forecasting Outputs to Discounted-Cash-Flow Robustness

Enrique Valiente[1]; Mariana Rojas[2]; Eduardo Redondo[3]
[1,2,3]Facultad de Ingeniería, Universidad Nacional de Asunción, Paraguay
evaliente@fiuna.edu.py, mrojas@fiuna.edu.py, eredondo@ing.una.py

Abstract*– This paper presents a scenario-based decision-support workflow to assess the feasibility of vertically integrating high-density polyethylene (HDPE) cap production (PCO-type beverage closures) into an existing PET bottle manufacturing operation in Paraguay under demand uncertainty. The approach links (i) monthly demand forecasting, (ii) technical and operational feasibility assessment through a shift-level stochastic production–inventory simulation that represents cycle-time variability and cycle-based preventive maintenance, and (iii) discounted-cash-flow appraisal derived from replication-level operational outputs. Demand projections instantiate discrete market-coverage scenarios, which are evaluated via repeated replications to obtain distributions of throughput, utilization, inventory/service indicators, and cost drivers. Financial performance is computed from scenario-specific cash flows using Net Present Value (NPV), Internal Rate of Return (IRR), and discounted payback, and complemented with a bisection-based critical-value analysis that estimates selling-price floors and resin-cost ceilings consistent with $\overline{NPV} \approx 0$. The case study shows that higher coverage increases average financial metrics over the evaluated range, while distributional results highlight increased variability and downside exposure as operation approaches effective capacity, underscoring the value of risk-aware feasibility assessment.

***Keywords*-- Demand forecasting; discrete-event simulation; injection molding; capacity planning; financial feasibility; sensitivity analysis.**

I. Introduction

Global demand for plastic products continues to grow, with packaging remaining one of the largest end-use sectors and a major driver of polymer conversion activities worldwide [1]. Within packaging, closures (caps) are high-volume components that must meet strict requirements in dimensional stability, sealing performance, and process consistency. High-density polyethylene (HDPE) is widely used for closures due to its mechanical properties and processability, although production performance and quality remain sensitive to injection-molding conditions and shop-floor variability [2 – 4].

For small and mid-sized manufacturers — particularly in emerging economies — decisions to vertically integrate cap manufacturing instead of sourcing externally are capital-intensive and difficult to reverse. Such decisions require not only equipment sizing, but also a realistic assessment of demand evolution, capacity utilization, and operational risk under alternative scenarios. While rolling-horizon approaches combining forecasting with production planning have been studied [7,8], integrating demand uncertainty, operational variability, and financial evaluation into a single decision-support framework remains a methodological challenge.

Demand forecasting plays a critical role in this context, as forecast errors propagate into capacity requirements, inventory exposure, and unit-cost variability. The forecasting literature highlights the importance of combining multiple time-series approaches (e.g., ARIMA and exponential-smoothing models) with transparent model selection and out-of-sample validation using appropriate accuracy metrics [9 – 12]. However, feasibility studies often rely on point forecasts without explicitly analyzing how uncertainty affects operational and financial performance.

From an operational perspective, discrete-event simulation (DES) provides a suitable framework for modeling manufacturing systems under stochastic conditions, including variability in processing times and maintenance interruptions [15]. When combined with cost modeling, DES enables evaluation of both operational performance and economic outcomes under realistic conditions [16 – 18]. Nevertheless, many feasibility assessments still treat operations deterministically, limiting their ability to capture variability-induced risks.

Motivated by these gaps, this work proposes a structured, scenario-based feasibility workflow that integrates demand forecasting, stochastic production – inventory simulation, and discounted-cash-flow analysis through a single decision variable: market coverage. This approach allows tracing how demand assumptions propagate through operational performance and ultimately affect financial outcomes, providing a risk-aware and decision-oriented evaluation framework. The remainder of the paper is organized as follows: Section II presents the methodology, Section III describes the case study and data, Section IV reports the results, and Section V discusses implications and limitations.

II. Methodology

A. Workflow Overview and Scenario Definition

This study follows a scenario-based decision-support workflow that links (i) market-demand projection, (ii) technical feasibility assessment through stochastic production–inventory simulation, and (iii) financial feasibility appraisal using discounted cash-flow analysis. The decision variable connecting all modules is the market-coverage level, denoted by α, defined as the fraction of the projected national demand supplied by the proposed HDPE cap production line.

A discrete set of coverage scenarios is evaluated: $\alpha \in \{0.25, 0.27, 0.29, 0.31, 0.33, 0.35\}$. For each α, the simulator is executed over multiple independent replications ($N = 30$ in the case study) to obtain empirical distributions—rather than point estimates—of operational and financial outcomes. The workflow consists of: (i) forecasting monthly demand over the planning horizon, (ii) constructing scenario-specific effective demand under each α, (iii) simulating shift-level production under stochastic cycle times and preventive-maintenance interruptions, (iv) translating operational outputs into cost drivers and scenario-specific cash flows, and (v) computing profitability indicators and critical economic thresholds through a structured sensitivity procedure. Algorithm 1 formalizes the end-to-end computational pipeline used in this study—from baseline demand forecasting to scenario scaling, stochastic production–inventory simulation, replication-level post-processing, and discounted-cash-flow evaluation—highlighting how the coverage parameter α links all modules.

Algorithm 1. Overall workflow (end-to-end pipeline)

```
Inputs: demand history {D_m}; coverage set 𝒜; replications N; parameters Θ_sim, Θ_fin; horizon H
Outputs: scenario summaries S(α), financial distributions F(α), thresholds T(α)
Fit candidate forecasting models on training data
Select baseline model f* using validation on test data
Compute baseline forecast: D̂_m ← f*(m),  m = 1..H
FOR each α ∈ 𝒜 DO
|  Construct scenario demand: D_m^(α) ← α · D̂_m
|  Ω_α ← ∅
|  FOR r = 1..N DO
|  |  Draw stochastic inputs ξ_r from Θ_sim
|  |  y_r ← SimulateShiftSystem(D^(α), ξ_r)
|  |  Ω_α ← Ω_α ∪ {y_r}
|  END FOR
|  S(α) ← Summarize(Ω_α; CI=95%)
|  F(α) ← FinancialMetrics(Ω_α, Θ_fin)   # NPV, IRR, payback
|  T(α) ← BreakEvenThresholds(Ω_α, Θ_fin) # bisection, NPV ≈ 0
END FOR
RETURN {S(α), F(α), T(α)}
```

B. Demand Forecasting and Effective-Demand Construction

Monthly demand is projected over a medium-term horizon of $H = 60$ months. Multiple time-series forecasting candidates, including ARIMA and exponential smoothing models, are fitted on the training data and compared on a holdout set using out-of-sample accuracy metrics (e.g., RMSE and sMAPE). The selected model provides the baseline forecast trajectory $\widehat{D}_m$, representing the projected national demand for month m.

For each coverage level α, the scenario-specific effective demand is defined by scaling the baseline forecast:

$$D_m^{(\alpha)} = \alpha \widehat{D}_m, m = 1, \dots, H. \qquad (1)$$

Sensitivity variants of the demand trajectory may be used as stress tests (e.g., systematic upward/downward perturbations around $\widehat{D}_m$); however, the core scenario analysis is driven by the baseline forecast combined with alternative values of α.

C. Shift-Level Production Simulation Model

A discrete-event simulation model is implemented at the shift level, where each shift represents an 8-hour production block. For each scenario α, every replication simulates the full sequence of shifts over the planning horizon while updating production output, inventory balances, and unmet demand. Outputs are recorded per shift and subsequently aggregated to monthly periods to align with the demand-forecasting granularity and the financial accounting intervals.

1) Stochastic cycle-time model

Cycle-time variability is modeled using a triangular distribution with parameters (a, c, b), where a is the minimum, b the maximum, and c the most-likely value. For each shift, a pseudo-random draw $u \sim \mathcal{U}(0,1)$ is generated and the cycle time is obtained via inverse transform:

$$t_{\text{cycle}} = \begin{cases} a + \sqrt{u(b-a)(c-a)}, & \text{if } u < \dfrac{c-a}{b-a}, \\ b - \sqrt{(1-u)(b-a)(b-c)}, & \text{otherwise}. \end{cases} \qquad (2)$$

This representation preserves bounded support and concentrates probability mass around the most-likely value, which is appropriate when empirical samples are limited but plausible operating ranges are available.

2) Preventive maintenance driven by accumulated cycles

Preventive maintenance is modeled as a deterministic interruption rule driven by accumulated mold cycles. The simulator tracks cumulative cycles and triggers a maintenance stop once a predefined threshold is reached. The maintenance event occupies one full shift (i.e., production is suspended during that shift), after which operations resume and the cycle counter is reset (or updated according to the maintenance logic). This mechanism explicitly couples realized throughput and downtime exposure, allowing maintenance-driven capacity losses to emerge endogenously from the simulated production trajectory.

3) State variables, flow balance, and outputs

At minimum, the simulator updates the following state variables at each shift: raw-material inventory, finished-goods inventory, and unmet demand (backlog). Each replication yields shift-level time series and monthly aggregates of produced units, throughput, capacity utilization, inventory trajectories for raw material and finished goods, unmet demand, and service-risk indicators (e.g., stockout occurrences and/or backlog magnitude). The simulation logic and inventory–demand balance updates are summarized in Algorithm 2.

Algorithm 2. Shift-level production, maintenance, and inventory dynamics

```
Inputs: monthly demand D_m^(α); triangular cycle-time θ_ct=(a,c,b); shift time T_shift
        PM threshold C_max and duration T_PM; inventory policy Π_inv
State:  I_FG, I_RM, B, C   # finished goods, raw material, backlog, cumulative cycles
Initialize I_FG, I_RM, B ← 0;  C ← 0
FOR each shift s in horizon DO
|  IF C ≥ C_max THEN
|  |  Production_s ← 0
```

```
| | Downtime_s ← 1 shift
| | C ← ResetCycles(C)
| ELSE
| | u ← Uniform(0,1)
| | t_ct ← TriangularInv(u; a,c,b)
| | Cap_s ← floor(T_shift / t_ct)
| | Production_s ← ApplyMaterialConstraint(Cap_s, I_RM)
| | C ← C + Production_s
| END IF
| I_RM ← UpdateRawMaterial(I_RM, Production_s, Π_inv)
| (I_FG, B, stockout_s) ← ServeDemand(I_FG, B, Production_s,
demand_at_shift(s))
END FOR
Aggregate shift outputs to monthly metrics; RETURN replication outputs
```

D. Statistical Post-Processing Across Replications

For each scenario α, replication outputs are summarized using descriptive statistics and 95% confidence intervals to quantify dispersion and compare scenario performance robustly. Distributional views (e.g., density plots and boxplots) are used to assess whether scenario differences are explained primarily by central tendency or by tail behavior under variability. The number of replications is selected to balance computational effort and confidence-interval stability.

E. Financial Appraisal and Automated Sensitivity Analysis

1) Computational protocol for scenario execution

To ensure repeatability and consistent treatment across scenarios, all experiments follow a fixed computational protocol. For each α, the simulation model is evaluated over R independent replications with controlled random seeds; replication-level operational outputs are then mapped into scenario-specific cash flows and financial indicators. This protocol ensures that differences across scenarios are attributable to α and simulated operational variability rather than to manual processing.

2) Conditional critical-value analysis via bisection ($NPV \approx 0$)

Economic robustness is assessed through a conditional critical-value analysis that identifies break-even thresholds for selected financial drivers. In the case study, two drivers are evaluated: the unit selling price of caps and the unit cost of resin. This procedure should not be interpreted as a Monte Carlo simulation over market prices or resin costs; instead, it searches for deterministic threshold values conditional on the operational uncertainty represented in the replications.

For a generic driver x, the procedure searches within a predefined interval $[x_L, x_U]$ and iteratively updates the bounds based on the sign of $\text{NPV}(x)$ until:

$$|\,\text{NPV}(x^{\star})\,| \leq \varepsilon \quad (3)$$

where ε is a tolerance and $x^{\star}$ is the estimated threshold. The resulting thresholds provide compact managerial interpretations: a minimum selling price (price floor) and a maximum resin cost (cost ceiling) under which the investment remains financially viable, conditional on the simulated operational variability. The overall protocol for threshold identification is summarized in Algorithm 3.

Algorithm 3: Automated critical-threshold search (bisection, NPV ≈ 0).

```
Inputs: scenario α; driver x ∈ {unit_price, resin_cost}; bounds [x_L, x_U];
tolerance ε
Output: threshold x* such that |NPV(x*)| ≤ ε (approx.)
WHILE (x_U − x_L) > δ DO
| x_mid ← (x_L + x_U)/2
| Set driver x ← x_mid in workbook model
| Run N replications; compute NPV_mid (e.g., mean or median over
replications)
| IF |NPV_mid| ≤ ε THEN RETURN x_mid END IF
| IF NPV_mid > 0 THEN
| | Update bounds toward less favorable side (depends on driver
monotonicity)
| ELSE
| | Update bounds toward more favorable side
| END IF
END WHILE
RETURN x_mid
```

III. Case Study and Data

A. Industrial setting and decision context

The case study considers a Paraguayan plastics manufacturer operating in the Greater Asunción industrial area, with additional production footprint in neighboring South American markets. The firm's core business is the manufacturing of PET preforms and bottles primarily for beverage packaging (carbonated and non-carbonated), with additional sales to cleaning, pharmaceutical, and cosmetics segments. The company also commercializes complementary packaging components, including small PET containers as well as caps and handles for jugs.

The strategic decision assessed in this work is whether to vertically integrate the production of HDPE caps to complement the existing PET bottle offering. The integration is evaluated as a means to reduce dependency on external suppliers, improve coordination across the packaging bundle offered to current customers, and expand the firm's product portfolio within its existing commercial channels.

From a product standpoint, the analysis focuses on PCO-type beverage closures. As technical references, the study considers two widely adopted 28-mm finishes—PCO 1810 and PCO 1881—which differ primarily in neck height and material usage. The PCO 1881 design is lighter and is commonly associated with resin-reduction objectives while maintaining sealing performance. These standards are used to define a consistent target product family for the feasibility assessment and to ensure internal consistency of operational and financial calculations across scenarios.

B. Demand data and market-coverage scenarios

Demand for caps is derived from internal commercial records under an operational linkage assumption: each unit of finished bottle sold requires one cap. Accordingly, monthly cap demand is proxied by the monthly volume of bottles sold over the same period. Historical records are consolidated into a monthly time series to obtain a regular demand signal suitable for forecasting and scenario construction. Prior to modeling, the dataset is cleaned and standardized to support reproducible downstream analysis, including removal of incomplete or inconsistent entries and harmonization of units and time stamps.

Market-coverage levels are evaluated as defined in Section II, reflecting that the firm currently serves this segment through procurement and commercialization rather than internal production. For each coverage level, the resulting scenario demand series is treated as the target demand to be served. When operational constraints generate stockouts, realized sales are endogenously capped by finished-goods availability, ensuring that revenue and service outcomes remain consistent with feasible operations.

For forecasting model selection, the monthly series is partitioned into training and test subsets using a 70% / 30% split. This enables out-of-sample validation of candidate models and avoids basing scenario construction solely on in-sample fit.

C. Production configuration and operational data used in the model

The production system is modeled as a single injection-molding line dedicated to HDPE cap manufacturing. The baseline configuration includes an injection molding machine equipped with a 96-cavity mold, supported by auxiliary feeding and dosing systems, material handling equipment, and downstream conveying elements.

Operational parameters are defined based on technical specifications and expert input. Cycle times are modeled as stochastic variables following a triangular distribution, reflecting minimum, most-likely, and maximum processing conditions observed in practice. Preventive maintenance is implemented as a cycle-based rule, where production is interrupted after a predefined number of accumulated cycles.

Production is simulated at the shift level, assuming 8-hour shifts, with output constrained by cycle time realizations, raw material availability, and maintenance interruptions. Inventory dynamics include both raw material and finished goods, with demand fulfillment subject to availability and resulting in potential backlog when production is insufficient.

All operational parameters used in the simulation model are derived from internal engineering estimates, supplier specifications, and structured discussions with technical personnel, ensuring consistency between modeled assumptions and realistic operating conditions.

D. Investment and financial input data

Investment is modeled at the equipment level and integrated into scenario-specific cash flows. The acquisition cost structure includes equipment purchase costs and import-related expenses (e.g., customs dispatch, port fees, brokerage services, and international and domestic freight), as well as an allowance for supplier support during installation and commissioning. For monetary consistency across technical and financial calculations, all figures are expressed in local currency when required using a reference exchange rate of 7,850 Gs per USD.

The baseline machinery configuration comprises the injection molding machine, 96-cavity mold, auxiliary feeding and dosing equipment, conveyor, forklift, and installation/commissioning items, resulting in a total initial machinery investment of 12,787,900,736 Gs. Alternative supplier quotations for key components (notably the injection machine and multi-cavity molds) are also documented to support procurement sensitivity checks and to assess exposure to capital-cost variability.

Financial feasibility is assessed over a five-year evaluation horizon, consistent with standard internal screening practices for capital projects. The discount rate is defined through a sector-representative cost of capital (WACC) used as the benchmark for interpreting IRR results. Depreciation is modeled using a straight-line method, with a ten-year useful life and a residual value specified as a fraction of the acquisition cost, consistent with common accounting/tax treatments for industrial machinery.

E. Data provenance, integration, and reproducibility notes

Operational and financial inputs were compiled from multiple internal functions, primarily sales, engineering, maintenance, and finance, and complemented with structured interviews with relevant personnel and with supplier/market information used to parameterize the equipment and cost assumptions. All inputs were organized into structured worksheets and standardized to maintain traceability across the forecasting, simulation, and financial modules.

To support reproducibility and auditability, the workflow is designed so that each market-coverage scenario is evaluated under the same data and modeling conventions: the same cleaned demand series feeds the forecasting and scenario-construction steps; the same parameter set governs the injection system and maintenance policy; and the same cash-flow construction logic is applied to derive financial indicators from replication-level operational outputs. This integration enables consistent comparisons across scenarios and reduces the risk of manual handling errors.

IV. RESULTS

A. Forecast-driven coverage scenarios (inputs to the operational–financial evaluation)

The empirical evaluation is organized around discrete market-coverage levels that scale the baseline monthly demand forecast into scenario-specific effective demand streams. Each coverage level represents a managerial commitment regarding the fraction of the target market that the proposed line intends to supply, and therefore acts as the linking decision parameter between demand uncertainty, operational feasibility, and discounted-cash-flow performance. Consistent with the workflow, scenario demand is constructed by multiplying the baseline forecast by the selected coverage factor, so that differences across scenarios are attributable to the intended market-coverage decision rather than to changes in modeling assumptions. This scenario design supports transparent comparisons because all scenarios share the same forecasting backbone, time aggregation, and accounting conventions, while their operational realizations differ through stochastic cycle times and maintenance-driven capacity losses.

Beyond interpretability, the discrete coverage design also enables a direct "stress gradient" over which the system's effective slack decreases. As coverage increases, the simulated line is asked to deliver a progressively larger fraction of monthly demand with the same installed capacity and preventive-maintenance logic. This framing is crucial for interpreting the downstream results: coverage increases may improve revenue potential and fixed-cost absorption, but they also reduce buffers against unfavorable operational realizations, making service deterioration and cash-flow volatility more visible.

B. Operational performance under stochastic cycle times and cycle-based preventive maintenance

For each coverage level, the shift-level simulator is executed across multiple independent replications, producing a distribution of outcomes rather than a single deterministic trajectory. The replication design is essential for feasibility decisions because it reveals not only expected performance but also dispersion and downside exposure driven by the interaction between (i) bounded-yet-variable cycle times and (ii) discrete maintenance interruptions triggered by accumulated cycles. Operational results are therefore interpreted through two complementary lenses: (a) central tendency (typical performance), and (b) variability (risk of unfavorable realizations that can compromise service, throughput stability, and ultimately revenues).

Throughput and effective utilization.

As market coverage increases, throughput requirements rise and the system is pushed closer to its effective capacity. Under low-to-mid coverage, replications tend to exhibit stable production trajectories: utilization remains moderate and fluctuations induced by cycle-time variability are typically absorbed without persistent service consequences. As coverage grows, the system operates with reduced slack, and replication-to-replication differences widen because the same bounded cycle-time uncertainty translates into larger absolute production shortfalls when utilization is already high. In practical terms, this means that, even if the mean throughput remains aligned with demand over the horizon, the probability of "tight months" increases: months in which realized output plus available inventory is insufficient to serve demand without stockouts or backlog accumulation.

Inventory dynamics, service exposure, and operational tipping behavior.

Inventory buffers play a central role in feasibility under uncertainty: they absorb temporary mismatches between stochastic production and deterministic scenario demand. Under lower coverage, inventory tends to act as a stabilizer; production variability and planned maintenance events are largely offset by available finished-goods stock, leading to a low incidence of unmet demand. As coverage increases, the system increasingly relies on just-in-time behavior: inventories become thinner, and adverse realizations (e.g., multiple slower-cycle shifts coinciding with maintenance) are less likely to be absorbed. In such regimes, service exposure grows nonlinearly: a modest reduction in realized output can translate into stockouts, lost sales, or backlog carryover, which then affects both operational indicators and financial performance.

A useful way to operationalize the notion of an "operational tipping point" in this replication-driven setting is not as a single deterministic threshold, but as the coverage range where service-risk indicators accelerate—e.g., a visible increase in the frequency or magnitude of unmet demand across replications. This risk-aware definition is more defensible for reviewers because it aligns with the logic of stochastic feasibility: feasibility should be framed as a probability of meeting demand, not as a binary property.

Maintenance-driven capacity losses and their interaction with variability.

Cycle-based preventive maintenance introduces discrete downtime events that reduce available production time and create intermittent capacity drops. Under scenarios with slack, these interruptions are absorbed with limited service impact; production catches up in subsequent shifts or months. Under high-coverage scenarios, the same downtime events become structurally more consequential because there is less residual capacity to recover. Importantly, the risk is not merely the average downtime fraction, but its timing relative to stochastic cycle-time realizations and demand peaks: unfavorable timing clusters can generate tail outcomes in which inventories deplete and unmet demand increases, thereby amplifying revenue variability.

C. Cost drivers and unit economics implied by operational realizations

Operational outcomes are mapped into cost drivers to connect physical feasibility with economic consequences. In this framework, unit economics are shaped by (i) resin consumption proportional to realized output, (ii) operating costs linked to machine time and shift activity, and (iii) feasibility-related penalties such as lost sales or demand not served when inventories are insufficient. Two mechanisms dominate the cost behavior across coverage levels.

First, as coverage increases from low levels, average unit costs can improve due to higher throughput and better absorption of quasi-fixed operating components. This is the conventional scale effect: producing closer to the intended operating point reduces the sensitivity of average costs to fixed overheads and partially smooths the impact of variability.

Second, once the system approaches effective capacity, service losses and throughput instability can offset these gains. In particular, stockouts (or persistent backlog) reduce realized sales relative to scenario demand, weakening fixed-cost absorption and increasing the effective cost per sold unit. This feasibility–economics coupling is one reason the replication-based approach is valuable: it prevents over-interpreting a purely deterministic cost calculation that assumes all demand is served.

D. Financial feasibility across market-coverage scenarios

Financial indicators are computed from scenario-specific cash flows derived from replication-level operational outputs, ensuring that the financial evaluation reflects operational feasibility constraints rather than assuming perfect service. Three indicators are used as primary screening metrics: Net Present Value (NPV), Internal Rate of Return (IRR), and discounted payback.

Across the evaluated coverage levels (25% to 35%), the results indicate strong financial attractiveness under the case-study assumptions. NPV increases from USD 1.69M at 25% coverage to USD 2.64M at 35% coverage, while IRR rises from 36% to 48% over the same range. Discounted payback shortens overall, from 3.26 years at 25% to 2.47 years at 35%. While the pattern is broadly monotonic in NPV and IRR, the payback metric exhibits mild non-monotonicity (e.g., small reversals across intermediate coverages), which is common when discounting interacts with the timing of cash flows and with feasibility-driven revenue variability. This reinforces that payback should be interpreted as a complementary indicator rather than as the sole decision criterion.

TABLE I
FINANCIAL PERFORMANCE INDICATORS ACROSS MARKET-COVERAGE SCENARIOS

Scenario (coverage, %)	NPV (USD)	IRR (%)	Discounted payback (years)
25	1,692,766	36	3.26
27	1,891,402	38	2.94
29	2,012,350	39	2.98
31	2,209,019	43	2.66
33	2,574,237	46	2.55
35	2,641,239	48	2.47

E. Economic robustness via critical thresholds (break-even margins)

Beyond point profitability metrics, the study reports economic robustness using a conditional critical-value analysis that searches for break-even thresholds consistent with the simulated operational uncertainty. Two thresholds are especially decision-relevant: (i) a selling-price floor (minimum unit selling price that keeps NPV approximately zero), and (ii) a resin-cost ceiling (maximum HDPE resin cost that keeps NPV approximately zero). These thresholds translate simulation-informed feasibility into compact managerial margins that can inform pricing decisions, supply negotiations, and risk screening.

The estimated price floors decrease from USD 6.72 per 1,000 caps at 25% coverage to approximately USD 6.38–6.40 at higher coverage levels, indicating that higher coverage improves tolerance to price pressure under the case-study conditions. In parallel, resin-cost ceilings increase from USD 1,837/ton at 25% to a peak around USD 1,946/ton near 33% coverage (with a small decline thereafter), suggesting improved resilience to input-cost increases over much of the evaluated range. The slight non-monotonicity at the top end is consistent with the operational tightening effect: when the system approaches effective capacity, feasibility-driven losses and variability can erode the marginal robustness benefits of expanding coverage.

TABLE II
CRITICAL ECONOMIC THRESHOLDS FOR FINANCIAL FEASIBILITY (NPV ≈ 0)

Scenario (coverage, %)	Critical selling price (USD per 1,000 caps)	Critical HDPE resin cost (USD/ton)
25	6.72	1,837.49
27	6.65	1,858.21
29	6.64	1,858.13
31	6.56	1,880.93
33	6.38	1,945.99
35	6.40	1,934.85

V. DISCUSSION

A. Interpretation of the main findings and decision relevance

The results support two linked interpretations. First, under the parameterization and financial assumptions of the case study, increasing market coverage within the evaluated range improves average profitability (NPV and IRR) and generally shortens discounted payback. This is consistent with higher revenue potential and improved fixed-cost absorption when demand is scaled upward without a parallel increase in installed capacity. Second, the same coverage increases reduce operational slack, which—under stochastic cycle times and discrete maintenance interruptions—can increase dispersion and downside exposure through service shortfalls. A reviewer-facing strength of the workflow is that it does not treat these as separate analyses: operational feasibility and finance are computed coherently, so that financial improvements are not "free," but conditional on the operational realizations that actually occur.

From a decision-making standpoint, this directly motivates a risk-aware interpretation of feasibility: the preferable coverage level is not determined only by the largest mean NPV, but by a balance between (i) expected value, (ii) tolerance to volatility and service deterioration, and (iii) managerial flexibility to respond (e.g., inventory buffers, maintenance policy adjustments, or contingency sourcing). This is precisely where a replication-driven analysis is more defensible than a single-run deterministic capacity check, because it shows how close-to-capacity operation can generate tail outcomes that matter for management, even when average profitability remains strong.

B. Traceability and Reproducibility of the Forecast – Simulation – Finance Pipeline

The workflow aligns with established principles across forecasting, manufacturing simulation, and investment appraisal research. In forecasting, the emphasis on out-of-sample validation and transparent model-selection logic is consistent with evidence that no single method dominates across contexts and that rigorous evaluation protocols are necessary before forecasts justify strategic investments. In operations modeling, the use of discrete-event simulation is appropriate when performance depends on queues, stochastic processing times, and policy-driven downtime—features that are difficult to capture faithfully in closed-form capacity formulas. In financial appraisal, translating replication-level operational outputs into cash flows addresses a common weakness of feasibility studies: treating operations as deterministic while interpreting financial metrics as certain.

A key methodological contribution, therefore, is not any single module, but the explicit coupling between modules through a managerial decision variable (coverage). This makes the analysis auditable: reviewers can trace how a change in coverage modifies scenario demand, how operational slack and maintenance interact with cycle-time uncertainty, and how these propagate into realized sales, costs, and discounted cash flows. In applied industrial papers, this traceability often matters as much as algorithmic novelty because it reduces ambiguity and increases credibility.

C. Managerial implications and practical levers

Several practical levers emerge from the results and their mechanisms:

1. **Coverage should be chosen jointly with a service-risk stance.** If the firm's commercial strategy can tolerate occasional service losses (e.g., through backordering or short-term procurement top-ups), higher coverage may remain acceptable even under increased variability. If service continuity is critical, the decision should incorporate explicit service constraints or risk limits (e.g., maximum acceptable stockout probability).
2. **Inventory policy can be treated as an additional hedge.** The results suggest that inventory buffers are a primary stabilizer under uncertainty. A firm can partially "buy" robustness by increasing safety stock or by adjusting replenishment triggers for raw material and finished goods. This is often cheaper than capacity expansion, but it carries working-capital costs that should be integrated into the financial module if policy changes are considered.
3. **Maintenance policy is not just an engineering detail—it is a capacity lever.** Cycle-based preventive maintenance affects not only average uptime but the timing of capacity drops. If higher coverage is desired, managerial attention to maintenance scheduling (e.g., aligning maintenance with low-demand periods when possible) can reduce service risk without changing installed capacity. Incorporating alternative maintenance policies would strengthen the paper by demonstrating that feasibility is policy-dependent rather than purely equipment-limited.
4. **Critical thresholds are actionable for contracting.** Price floors and resin-cost ceilings provide compact negotiation targets. They can be used to evaluate whether current commercial contracts and resin procurement conditions leave adequate margin under operational uncertainty, and to define escalation clauses or hedging policies.

D. Limitations, validity conditions, and what reviewers may challenge

To preempt typical reviewer concerns, the discussion should state clearly what the model does and does not claim:

- **Demand proxy and market representation.** Demand is proxied through bottle sales under a one-cap-per-bottle linkage. This is reasonable for the product structure but may miss shifts in cap mix, customer behavior, or substitution effects.
- **Operational scope.** The simulation captures stochastic cycle times and preventive maintenance but does not explicitly model unplanned downtime, scrap/yield losses, labor constraints, start-up learning curves, or multi-product scheduling complexities. These factors could widen uncertainty bounds and shift the effective tipping range.
- **Parameter uncertainty.** Triangular cycle-time ranges and maintenance thresholds are plausible but would benefit from empirical calibration using shop-floor measurements. A reviewer may request sensitivity checks that demonstrate robustness of conclusions to wider variability or to alternative downtime assumptions.
- **Financial assumptions.** Results depend on discount rate, depreciation choices, tax/accounting conventions, and exchange-rate consistency. Sensitivity to the discount rate (e.g., ±2–3%) and key cost drivers (in addition to resin cost and selling price) may affect the magnitude of financial indicators and should be considered when interpreting results.
- **Generalizability.** The contribution is best positioned as a reproducible evaluation workflow for make-or-buy/vertical-integration feasibility under uncertainty, illustrated through an HDPE closure line. Claims should avoid implying universal profitability and instead emphasize decision transparency and risk-aware screening.

E. Extensions that strengthen the paper without changing its core identity

Several extensions can be framed as future work (or brief robustness checks, if feasible within space limits):

- Calibrate cycle-time variability and downtime using empirical logs; add unplanned downtime and losses.
- Evaluate alternate inventory and maintenance policies as controllable levers, not fixed assumptions.
- Introduce an explicit service-level objective/constraint and report feasibility in probabilistic terms (e.g., probability of meeting monthly demand).
- Combine the threshold analysis with multi-parameter sensitivity (e.g., resin cost and selling price) to map a feasible region rather than single break-even points.

VI. CONCLUSION

This paper presents a reproducible, decision-oriented pipeline for industrial feasibility assessment under uncertainty by integrating demand forecasting, stochastic production – inventory simulation, and discounted-cash-flow analysis. The main contribution lies in the traceable coupling between these components through a managerial decision variable—market coverage — allowing operational feasibility and financial performance to be evaluated coherently under variability. This structure enables a risk-aware assessment by explicitly linking processing variability, downtime, inventory dynamics, and financial outcomes.

Applied to the HDPE cap injection-molding integration case study, the results indicate favorable financial performance across the evaluated market-coverage range (25% – 35%) under the assumptions. NPV increases from USD 1.69M to USD 2.64M, IRR from 36% to 48%, and discounted payback shortens from 3.26 to 2.47 years. In addition, critical thresholds provide actionable robustness measures: the selling-price floor decreases from USD 6.72 to approximately USD 6.38 – 6.40 per 1,000 caps, while the resin-cost ceiling increases from USD 1,837/ton to approximately USD 1,946/ton.

From a practical perspective, the proposed pipeline supports repeatable evaluation of make-or-buy and vertical-integration decisions under operational uncertainty, producing interpretable and decision-oriented outputs that can be used by both engineering and financial stakeholders. Further work may include empirical calibration of operational parameters, incorporation of unplanned downtime and yield losses, and evaluation of alternative inventory and maintenance policies within the same framework.

REFERENCES

[1] M. Dokl, A. Copot, D. Krajnc, Y. V. Fan, A. Vujanović, K. B. Aviso, R. R. Tan, Z. Kravanja, and L. Čuček, "Global projections of plastic use, end-of-life fate and potential changes in consumption, reduction, recycling and replacement with bioplastics to 2050," *Sustainable Production and Consumption*, vol. 51, pp. 498–518, 2024. https://doi.org/10.1016/j.spc.2024.09.025

[2] P. Matarrese, A. Fontana, M. Sorlini, L. Diviani, I. Specht, and A. Maggi, "Estimating energy consumption of injection moulding for environmental-driven mould design," Journal of Cleaner Production, vol. 168, pp. 1505–1512, 2017, doi: 10.1016/j.jclepro.2017.07.144.

[3] A. Elduque, D. Elduque, C. Pina, I. Clavería, and C. Javierre, "Electricity Consumption Estimation of the Polymer Material Injection-Molding Manufacturing Process: Empirical Model and Application," Materials, vol. 11, no. 9, art. 1740, 2018, doi: 10.3390/ma11091740.

[4] B. Agarski, D. Vukelić, M. Ilić Micunović, I. Budak, A. M. Pavlović, and D. Vučurović, "Evaluation of the environmental impact of plastic cap production, packaging, and disposal," Journal of Environmental Management, vol. 245, pp. 55–65, 2019, doi: 10.1016/j.jenvman.2019.05.078.

[5] W. Rungswang, W. Nantasetphong, and C. Jarumaneeroj, "Influences of processing conditions on gradient structures in injection-molded polyethylene caps and their property correlations," Journal of Applied Polymer Science, vol. 140, no. 35, e54354, 2023, doi: 10.1002/app.54354.

[6] N. Dulal, K. L. F. J. Looijmans, J. A. B. M. Janssen, G. L. Drisko, and G. C. Rutledge, "Slip-additive migration, surface morphology, and performance on injection moulded high-density polyethylene closures," Journal of Colloid and Interface Science, vol. 505, pp. 537–545, 2017, doi: 10.1016/j.jcis.2017.06.040.

[7] L. Tiacci and S. Saetta, "Demand forecasting, lot sizing and scheduling on a rolling horizon basis," International Journal of Production Economics, vol. 140, no. 2, pp. 803–814, 2012, doi: 10.1016/j.ijpe.2012.02.007.

[8] Y. R. Sagaert, N. Kourentzes, S. De Vuyst, E.-H. Aghezzaf, and B. Desmet, "Incorporating macroeconomic leading indicators in tactical capacity planning," International Journal of Production Economics, vol. 209, pp. 12–19, 2019, doi: 10.1016/j.ijpe.2018.06.016.

[9] R. J. Hyndman and A. B. Koehler, "Another look at measures of forecast accuracy," *International Journal of Forecasting*, vol. 22, no. 4, pp. 679–688, 2006.

[10] R. J. Hyndman, A. B. Koehler, R. D. Snyder, and S. Grose, "A state space framework for automatic forecasting using exponential smoothing methods," *International Journal of Forecasting*, vol. 18, no. 3, pp. 439–454, 2002. https://doi.org/10.1016/S0169-2070(01)00110-8

[11] R. J. Hyndman and Y. Khandakar, "Automatic time series forecasting: the forecast package for R," *Journal of Statistical Software*, vol. 27, no. 3, pp. 1–22, 2008. https://doi.org/10.18637/jss.v027.i03

[12] S. Kim and H. Kim, "A new metric of absolute percentage error for intermittent demand forecasts," International Journal of Forecasting, vol. 32, no. 3, pp. 669–679, 2016, doi: 10.1016/j.ijforecast.2015.12.003.

[13] S. Makridakis, E. Spiliotis, and V. Assimakopoulos, "The M4 Competition: 100,000 time series and 61 forecasting methods," *International Journal of Forecasting*, 2020.

[14] S. Makridakis, E. Spiliotis, and V. Assimakopoulos, "Statistical and Machine Learning forecasting methods: Concerns and ways forward," *PLOS ONE, 13(3)*, 2018.

[15] A. Negahban and J. S. Smith, "Simulation for manufacturing system design and operation: Literature review and analysis," *Journal of Manufacturing Systems*, vol. 33, no. 2, pp. 241–261, 2014.

[16] K. Kendall, C. Mangin, and E. Ortiz, "Discrete event simulation and cost analysis for manufacturing optimisation of an automotive LCM component," *Composites Part A: Applied Science and Manufacturing*, vol. 29, no. 7, pp. 711–720, 1998. https://doi.org/10.1016/S1359-835X(98)00003-7

[17] A. Cigolini, M. Cozzi, and M. Perona, "Configuration and performance analysis of supply chains: a DES-based approach," *Simulation Modelling Practice and Theory*, 2014.

[18] T. A. Spedding and G. Q. Sun, "Application of discrete event simulation to the activity based costing of manufacturing systems," *International Journal of Production Economics*, 1999.

[19] R. Pozzi, C. Noè, V. Lazzarotti, and T. Rossi, "Using simulation for reliable investment appraisal: evidence from a case study," *International Journal of Operational Research*, vol. 23, no. 1, pp. 45–62, 2015.

[20] A. Koutsandreas, E. Spiliotis, F. Petropoulos, and V. Assimakopoulos, "On the selection of forecasting accuracy measures," *Journal of the Operational Research Society*, 2021.