Street Art & Urban Creativity | Vol. 11, No. 4, 2025 | ISSN 2183-9956
Street Art & Urban Creativity. Scientific Journal

# Prescriptive Analytics in Supply Chain Management

## Optimizing Inventory Allocation and Dynamic Pricing Under Uncertainty

MAHESH REDDY CHALLA[1]
[1]3151 SOLANA CIRCLE DENTON TEXAS ZIP CODE 76203.

## Abstract

Supply chains generate more data than ever, yet most analytics stop at description or prediction: they say what happened or what demand will probably be, and leave the decision to a planner. Prescriptive analytics goes one step further and recommends the action. This paper studies prescriptive decision making for a seasonal product sold through several stores that are supplied from a central warehouse, where two decisions are tightly coupled and both are made under demand uncertainty: how to allocate a fixed stock between stores and a central reserve, and what price to charge week by week. We propose a rolling prescriptive policy that plans with scenarios, learns about demand from early sales using a Bayesian update that accounts for stockouts, re-optimises the price each week within operational limits, and replenishes stores from the reserve. The policy is compared with four benchmarks (a static plan, a certainty-equivalent plan, a rule-based markdown policy, and a one-shot scenario-based plan) on 1,000 simulated seasons that share the same random demand. The prescriptive policy raises mean profit by 19.2% over the static plan and by 12.7% over the best one-shot plan, improves the average of the worst 10% of seasons by 33.6%, lifts the share of demand served from 82.0% to 95.3%, and cuts unsold stock from 154 to 50 units. Ablations show that dynamic pricing contributes the largest share of the gain, followed by the central reserve and by Bayesian learning, and sensitivity tests show that the advantage persists under higher demand uncertainty, forecast bias and a misjudged price sensitivity. The evidence comes from a simulation with a well-specified learning model, so the absolute gains should not be transferred to any real business without validation. The study offers a transparent template for evaluating prescriptive policies against realistic benchmarks.

**Keywords:** prescriptive analytics, supply chain management, inventory allocation, dynamic pricing, stochastic optimization, demand uncertainty, Bayesian learning, revenue management.

## 1. Introduction

Modern supply chains are among the most data-rich operations in any industry. Point-of-sale systems record every transaction, warehouse systems track every pallet, and planning systems hold forecasts for thousands of products and locations. The analytics built on top of these data are usually described along a maturity ladder with three rungs. Descriptive analytics summarises what has happened, predictive analytics estimates what is likely to happen, and prescriptive analytics recommends what to do about it [1]. Most organisations have invested heavily in the first two rungs. A forecast, however accurate, is not a decision. Someone must still decide how many units to send to each store, how many to hold back, and what price to put on the shelf, and these decisions interact with each other and with the uncertainty that the forecast describes only in part.

Two decisions deserve particular attention because they sit at the point where operations and commercial teams meet. The first is inventory allocation: when a finite quantity is available for a selling season, how much should be pushed to each store now and how much should be kept in a central reserve for later release? Pushing stock early makes it available where customers are but commits it to a location where demand might not materialise. Holding stock centrally preserves flexibility, because the same unit can later serve whichever store turns out to

be short, but delays availability and adds transport cost. The second decision is dynamic pricing: when demand is uncertain and stock is limited, what price should be charged each week? A price that is too low sells out early and leaves money on the table, while a price that is too high leaves stock unsold at the end of a season, when it can only be salvaged at a loss.

These two decisions are usually made separately. Supply planners set allocations from a demand forecast and merchants set price from a pricing calendar, often with the demand forecast treated as certain. This separation is costly for a simple reason. The best price depends on how much stock is available where, and the best allocation depends on the price that will be charged, because price changes the demand that must be served. Both depend on how much demand will eventually turn out to be, a quantity that is uncertain at the start of the season but gradually revealed by early sales. A policy that plans once and then executes blindly wastes the information that arrives every week.

This paper develops and evaluates a prescriptive approach that treats the two decisions jointly and updates them as information arrives. The approach has four ingredients. Scenario-based stochastic optimisation produces the initial allocation, the size of the central reserve and the opening price while recognising that demand is uncertain. A Bayesian learning step revises beliefs about demand each week from observed sales, taking care to treat sales in stocked-out weeks as incomplete information. A rolling re-optimisation adjusts the price each week within a practical limit on how far it may move. A replenishment rule releases reserve stock to stores that are running short. The ingredients are individually well known, but their combined value, and the contribution of each, is rarely quantified against realistic benchmarks.

The contributions of the paper are as follows.

1. A formulation of the joint allocation and pricing problem for a multi-store seasonal product, with a central reserve, one-week replenishment lead time and a discrete price ladder with limited weekly price movement.
2. A rolling prescriptive policy that combines scenario-based planning, Bayesian demand learning that handles stockout censoring, weekly re-pricing and reserve release.
3. A controlled simulation study with four benchmark policies, three ablations, paired statistical comparison on shared random seasons, and sensitivity analyses on demand uncertainty, forecast bias and price sensitivity.

Evidence on where the value comes from, expressed in profit, downside risk, service level and leftover stock, which guides the choice between simpler and more sophisticated approaches.

## 2. LITERATURE REVIEW

### *A. From descriptive to prescriptive analytics*

The distinction between descriptive, predictive and prescriptive analytics has become standard in operations management. Prescriptive analytics uses optimisation and decision models, often combined with predictive models, to recommend actions. A central insight in recent work is that predicting well and deciding well are different tasks: a forecast that minimises average error is not necessarily the one that leads to the best decision, because the costs of over- and under-forecasting are asymmetric and depend on the decision [1]. This motivates approaches that evaluate models by the quality of the decisions they induce, which is the viewpoint adopted in this paper.

### *B. Inventory allocation and risk pooling*

The classical analysis of inventory under demand uncertainty starts from the single-period problem, in which a decision maker orders before demand is known and bears the cost of too much or too little. For several locations, a well-established result is that holding stock centrally and allocating it later reduces the total stock needed for a given service level, because demand variations across locations partly cancel when pooled [2]. This risk-pooling effect is the reason that retailers keep distribution centres and regional reserves. The value of postponing allocation grows with demand uncertainty, the number of locations and the independence of their demand, and falls with the cost and delay of moving stock. Accurate-response approaches extend the idea by using early sales to refine forecasts before committing the remaining stock, and they report large reductions in the cost of demand uncertainty for seasonal fashion products [3]. The present study follows this line, with a central reserve whose release is informed by learning from early sales.

### *C. Dynamic pricing and revenue management*

Revenue management began in the airline industry and has spread to hotels, car rental, retail and e-commerce. Its central problem is selling a limited, perishable inventory over a finite horizon by adjusting prices. Foundational analysis shows how the optimal price for a given unit of remaining stock falls with the time left and rises as stock becomes scarcer, and describes the structure of optimal policies when demand is random [4]. Textbook treatments cover forecasting, capacity control and pricing in detail [5], [6]. Surveys of dynamic pricing with inventory considerations emphasise the practical constraints that real businesses face: limited price changes, price ladders, consumer perceptions of fairness and the need for pricing to be coordinated with supply decisions [7]. The present paper adopts these constraints by using a discrete price ladder and restricting price movement to one step per week.

### *D. Pricing with unknown demand*

In practice the demand function is not known, and a seller must learn it while selling. This creates a trade-off between earning and learning: a price that is good for learning may not be the best for current profit. A body of work studies near-optimal policies for pricing without knowing the demand function and provides bounds on the cost of learning [8], and surveys trace the historical origins and current directions of this literature [9]. Our setting is simpler because the structure of the demand model is assumed known up to an uncertain level, which is a common situation in retail where the shape of the response to price is estimated from history but the level for a given season is not. This simplification is a limitation to which we return in the discussion.

### *E. Stochastic programming and scenario approximation*

Stochastic programming provides a formal framework for decisions made before uncertainty is resolved, with recourse decisions taken afterwards [10], [11]. Because the exact expectation over a continuous distribution is rarely tractable, the sample average approximation method replaces it with an average over a finite set of simulated scenarios, and it has well-studied statistical properties: as the number of scenarios grows, the solution converges to the true optimum [12]. Scenario-based methods are attractive in supply chains because they accommodate non-standard demand models, discrete decisions and operating rules that make closed-form solutions impossible. Their main costs are computation and the need to specify the scenario distribution. We use a scenario approximation for both the initial plan and the weekly re-optimisation.

### *F. Risk, service and the bullwhip effect*

Expected profit is not the only criterion. Risk measures such as the average of the worst outcomes, known as conditional value at risk, quantify downside exposure [13], and service level matters because stockouts harm future demand in ways single-season profit does not capture. Poor information also amplifies demand variability upstream, the bullwhip effect, driven partly by forecasting and price fluctuation [14]. We therefore report downside risk and service level alongside mean profit..

## 3. RELATED WORK

**Joint inventory and pricing.** The literature on joint inventory and pricing decisions is extensive. Surveys of dynamic pricing with inventory emphasise that pricing and inventory should be coordinated because pricing can be used as a lever to align demand with available supply, and that ignoring the interaction can lose significant profit [7]. Studies in this area typically use analytical models with a single location and focus on structural properties of optimal policies. Our work differs in three ways: it covers several locations with a central reserve, it uses a simulation-based approach that accommodates operational constraints, and it measures the contribution of each element empirically.

**Early sales and quick response.** Fisher and Raman showed that placing orders in two stages, with the second informed by early sales, can substantially reduce the cost of demand uncertainty for products with long lead times and short selling seasons [3]. The structure of our policy shares the spirit of that idea, but the second-stage decision is the release of reserve stock and the price, and the learning model accounts explicitly for stockouts that hide demand.

**Learning-based pricing.** Learning-based pricing methods study how fast a seller can learn an unknown demand function while selling [8], [9]. We use a lighter weight Bayesian approach tailored to settings where the shape of the response is known and only the level is uncertain. The approach is consistent with the way many retailers work, using planned price elasticities from previous seasons and updating the level of demand as a season unfolds.

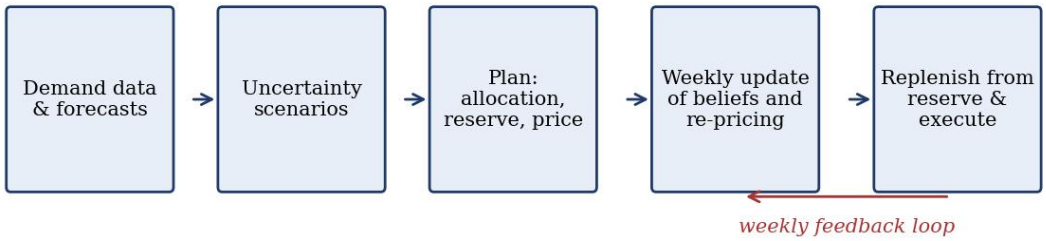
**Gap.** Few studies compare, within a single transparent protocol, the benefit of scenario-based planning, dynamic pricing, central reserves and Bayesian learning, with paired statistical comparisons and attention to downside risk. This study aims to supply such evidence in a controlled setting..

## 4. METHODOLOGY

### *A. Problem setting*

A retailer sells one seasonal product over ten weeks through five stores. A fixed quantity of 1,300 units has been bought before the season, and the purchase decision is outside the scope of the study. Before the season starts, the retailer decides how many units to place in each store and how many to keep in a central reserve, and it chooses an opening price from a ladder of nine prices from 80 to 120 in steps of five, with 100 being the reference list price. Each week, the retailer may move the price by at most one step, may release reserve stock to stores, and then observes sales. Shipments from the reserve take one week to arrive and cost three currency units per unit. Holding stock in stores costs 0.4 units per unit per week. Stock remaining at the end of the season, in stores or in the reserve, is salvaged at 20 per unit. Each unit was purchased at 45, which is a sunk cost that is the same under all policies.

Demand is uncertain in three ways. The overall level of demand for the season varies from year to year, which affects all stores together and is not known in advance. Each store also has its own deviation from its expected level. Finally, weekly demand is random around its mean, following a seasonal pattern that peaks in the middle of the season. Demand falls as price rises, with a price sensitivity such that a ten per cent rise in price reduces demand by roughly seventeen per cent.




.

***Fig. 1. Structure of the prescriptive rolling policy.***

### *B. Candidate policies*

We compare four benchmark policies and the proposed policy.

**Static policy.** The stock is divided between stores in proportion to forecast demand, no reserve is kept, and the list price of 100 is charged throughout. This represents a simple practice in which allocation is proportional and pricing is fixed.

Deterministic (certainty-equivalent) policy. The planner treats expected demand as if it were certain. It chooses the single price at which expected demand best matches the available stock, allocates stock in proportion to forecast, and keeps no reserve. Such a plan ignores the shape of the uncertainty. It represents the common practice of planning with point forecasts, even though it uses an optimisation step.

**Rule-based markdown policy.** The stock is divided proportionally and the opening price is the list price, but the retailer reduces the price by one step at fixed checkpoints in weeks four, six and eight if cumulative sales are more than ten per cent behind a plan. This is typical of retail markdown practice and responds to information, but through a simple rule and in one direction only.

**Stochastic static policy.** The planner uses scenario-based optimisation to choose the allocation to stores, the size of the central reserve and a single price for the whole season. Reserve stock is released each week with a replenishment rule. There is no price adjustment and no learning. This policy isolates the value of recognising uncertainty at the planning stage.

**Prescriptive rolling policy (proposed).** The policy starts from the stochastic plan, which sets the initial allocation, the reserve and the opening price, and then, every week, (i) updates beliefs about demand, (ii) re-optimises the price, and (iii) releases reserve stock.

### *C. Scenario-based planning*

The initial plan is found by sample average approximation. Two hundred scenarios of demand multipliers are generated from the assumed uncertainty. For each candidate price on the ladder, the planner builds an allocation incrementally: stock is added in chunks of 25 units, and each chunk goes to the store or to the reserve that raises average scenario profit the most. Average profit is computed by simulating the ten-week season in every scenario, including the replenishment rule. The best price and corresponding allocation are chosen. The method is greedy but works well here because the marginal value of additional stock at any location declines as stock increases. The same plan serves all seasons, because it depends only on forecasts and assumed uncertainty and not on realised demand.

### *D. Replenishment rule*

From the second week until the penultimate week, the retailer compares each store's stock with the expected demand over the current and next two weeks. Stores with less stock than this target receive units from the reserve, shared proportionally if the reserve is not sufficient to cover all needs. Because shipments take a week, they arrive at the start of the following week. The same rule is applied by every policy that holds a reserve, and the expected demand used in the rule is the forecast for static policies and the updated belief for the prescriptive policy.

### *E. Bayesian demand learning*

The prescriptive policy maintains a belief about each store's demand level relative to the forecast. Beliefs are expressed with a gamma distribution, which is the natural choice for Poisson demand and updates simply by adding observed sales and expected sales to its two parameters. Two features of the update matter. First, the beliefs are pooled across stores: since part of the uncertainty is a season-wide shift, the combined sales of all stores are informative about each store, and the prior for each store is centred on the common level estimated from all stores, with weight given to the amount of data seen. Second, the update uses only weeks in which a store did not run out of stock. When a store sells out, observed sales are a lower bound on demand and not a measurement, and treating them as measurements would lead the planner to underestimate demand exactly where stock is scarce. Weeks with stockouts are therefore excluded from learning. This is a conservative choice that discards some information; more sophisticated treatments of censored demand exist and could improve performance.

### *F. Weekly re-optimisation of price*

At the start of each week after the first, the policy draws forty scenarios of demand multipliers from the updated beliefs. For each price within two steps of the current price, it simulates the rest of the season with that price held constant, using the current stock levels, the reserve and the replenishment rule, and computes average profit across scenarios. If the best candidate is above or below the current price, the price moves one step in that direction. Limiting movement to one step per week reflects a practical constraint, since large and frequent price swings upset customers and are difficult to execute across channels. The use of a constant future price in the look-ahead is an approximation, but because the decision is repeated every week, the policy can correct itself as information arrives.

### *G. Evaluation metrics and statistical comparison*

The main outcome is profit per season, defined as revenue plus salvage value minus holding cost, transport cost and purchase cost. We also report the standard deviation of profit, the average profit in the worst 10% of seasons (a downside risk measure), the fill rate (the share of customer demand that was served), the sell-through (the share of the stock that was sold before the end of the season), the average realised price, and the units left unsold. To compare policies,

all of them are run on the same 1,000 simulated seasons, so that each season gives one paired difference in profit between any two policies. We report the mean paired difference and an interval of plus or minus 1.96 standard errors. Pairing removes the large season-to-season variation in demand from the comparison and gives much tighter intervals than comparing unpaired averages.

## 5. EXPERIMENTAL SETUP

### *A. Why a simulation*

A retailer observes the outcome of the decisions it actually made, not of the alternatives, so policy comparison needs a demand model. Simulation allows controlled comparison under identical conditions, at the cost of realism: the results show how policies behave in a stylised environment, not what a specific retailer would earn.

### *B. Simulated environment*

The simulated network has five stores with expected weekly demand at the reference price of 40, 32, 26, 22 and 18 units, so that expected demand over the season at the reference price is 1,380 units against 1,300 available. This near balance of supply and demand makes pricing and allocation consequential, since there is neither plentiful stock nor severe scarcity. The seasonal pattern rises from 60% of the average level in the first week to a peak of 130% in weeks five and six and falls to 50% in the last week. Price sensitivity is such that demand at a price of 80 is about 56% above demand at 100, and demand at 120 is about 31% below it. The season-wide demand level has a coefficient of variation of 30%, and each store has an additional independent deviation of 20%. Weekly sales are drawn from a Poisson distribution around the resulting mean. These settings were fixed before any experiments were run.

### *C. Planner's knowledge*

In the main experiment, the planner knows the forecast for each store, the shape of the seasonal pattern, the price sensitivity and the sizes of the two uncertainty components, but not the realised season-wide level or store deviations. This is an optimistic assumption, because it means that the scenario model and the learning model are well specified. The sensitivity analyses relax it by giving the planner a biased forecast, a wrong price sensitivity and different levels of uncertainty.

### *D. Implementation and replication*

The simulation, the policies and the analysis were implemented in Python using NumPy, and all random numbers are generated from fixed seeds. The main experiment uses 1,000 test seasons. The initial plan uses 200 scenarios, and each weekly re-optimisation uses 40 scenarios. The sensitivity analyses use 400 seasons per setting and 30 scenarios in the weekly re-optimisation, to limit computing time. Computation was light: the full main experiment, including all ablations, finished in about one minute on a single core.

## 6. RESULTS

### *A. Main comparison*

Table 1 summarises the performance of the five policies, and Fig. 2 compares mean profit and downside risk.

**Table 1. Performance over 1,000 simulated seasons (profit in currency units per season)**

| Policy | Mean profit | Std. dev. | Worst-10% average | Fill rate | Sell-through | Avg. price | Unsold units |
|---|---|---|---|---|---|---|---|
| Static | 56,534 | 14,895 | 22,964 | 82.0% | 87.9% | 100.0 | 154 |
| Deterministic | 57,659 | 17,437 | 21,238 | 86.3% | 83.9% | 105.0 | 206 |
| Rule-based markdown | 57,292 | 13,672 | 26,472 | 82.0% | 91.0% | 97.9 | 114 |
| Stochastic static | 59,788 | 17,755 | 22,732 | 87.9% | 85.5% | 105.0 | 189 |
| **Prescriptive rolling** | **67,396** | 18,422 | **30,675** | **95.3%** | **96.2%** | 101.8 | **50** |

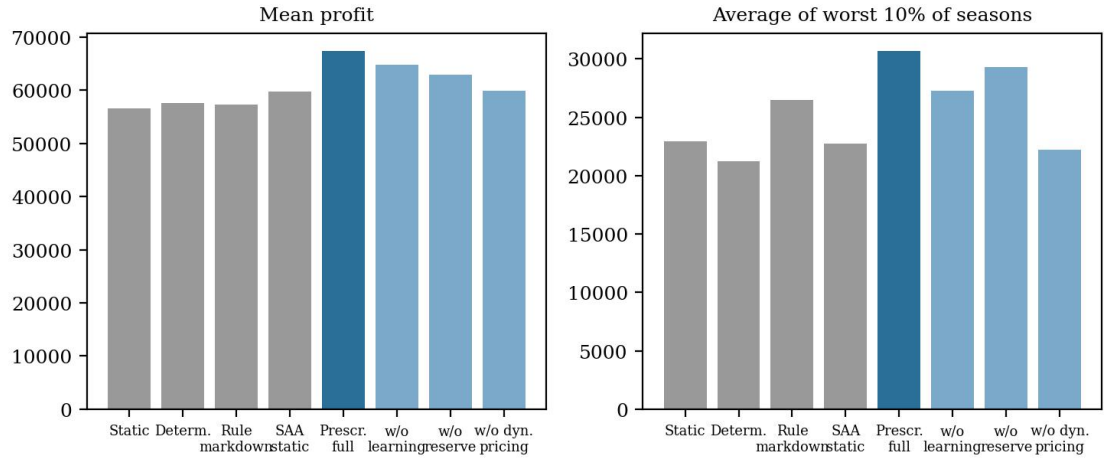



***Fig. 2. Mean profit and the average of the worst 10% of seasons for all policies and ablations.***

The prescriptive policy earns the highest mean profit, 67,396, against 56,534 for the static plan, an improvement of 19.2%. Table 2 shows that the improvement is large relative to its statistical uncertainty.

**Table 2. Paired profit differences (mean difference and 95% interval half-width, currency units per season).**

| Policy | Difference vs. static | Difference vs. stochastic static |
|---|---|---|
| Deterministic | +1,125 ± 263 | −2,129 ± 144 |
| Rule-based markdown | +757 ± 154 | −2,496 ± 340 |
| Stochastic static | +3,253 ± 289 | 0 |
| **Prescriptive rolling** | **+10,861 ± 394** | **+7,608 ± 311** |

Every policy beats the static plan, and the ordering among the benchmarks is clear. Planning with a point forecast (deterministic) adds about 2% to profit, markdown rules add about 1%, and scenario-based planning adds about 5.8%. The rolling prescriptive policy adds 12.7% on top of the best one-shot plan. The gains of the benchmarks are modest because none of them uses information from the season as it unfolds, except the markdown rule, which does so in a limited way.

The service and stock measures tell the same story. The prescriptive policy serves 95.3% of customer demand, against 82.0% for the static plan, and it leaves 50 units unsold at the end of the season, against 154. It achieves this while maintaining an average realised price of 101.8, close to the reference price, which means that the gain comes from better matching of stock to demand and not from heavy discounting or aggressive price increases.

## *B. Downside risk*

The average profit in the worst 10% of seasons improves from 22,964 under the static plan to 30,675 under the prescriptive policy, an increase of 33.6%, a larger relative gain than for the mean. A plausible reading is that when demand is low, early sales reveal the shortfall and the policy lowers prices while stock in reserve remains uncommitted, but we did not isolate this mechanism.

Two other results on risk are worth noting. First, the certainty-equivalent policy has a lower worst-10% average (21,238) than the static plan, even though its mean is higher. Its higher price is calibrated to expected demand, and the larger unsold stock (206 units) is consistent with that price hurting in bad seasons, which illustrates a danger of planning with point forecasts. Second, the rule-based markdown policy has the lowest standard deviation (13,672) and the best downside among the non-prescriptive policies (26,472), showing that a simple markdown rule is an effective guard against bad seasons, although it gives away little of the upside and leaves the average profit nearly unchanged.

The prescriptive policy has the highest standard deviation of profit (18,422). This is not a sign of greater risk in the sense of bad outcomes: its worst seasons are better than those of any other policy, and its variability comes from the upside, since it earns much more in high-demand seasons. Standard deviation penalises upside and downside alike, which is why the worst-10% average is a more appropriate risk measure here.

## C. Ablations

Table 3 shows what happens when each ingredient is removed from the prescriptive policy.

**Table 3. Ablation results.**

| Variant | Mean profit | Loss vs. full policy | Fill rate | Unsold units |
|---|---|---|---|---|
| Full prescriptive policy | 67,396 | – | 95.3% | 50 |
| Without Bayesian learning | 64,814 | −2,582 | 93.8% | 104 |
| Without central reserve | 62,901 | −4,495 | 90.6% | 107 |
| Without dynamic pricing | 59,880 | −7,516 | 88.0% | 188 |

All three ingredients contribute. Removing dynamic pricing costs the most, 7,516 per season, which brings the policy back to the level of the stochastic static plan (the difference from that plan is 92 ± 176, not distinguishable from zero). Removing the central reserve costs 4,495, and removing Bayesian learning costs 2,582. A few observations follow.

First, dynamic pricing is the dominant lever in this environment. Without it, reallocating from the reserve using updated beliefs adds essentially nothing to the one-shot plan. This is because, when price is fixed, the only way to respond to a demand shortfall or surge is to move stock, and in this setting the capacity to move stock is limited by the reserve and by the one-week lead time. With price as a second lever, the policy can shift demand to match supply, which is a much more flexible response.

Second, the pieces reinforce each other. Dynamic pricing without learning still adds 5,026 over the stochastic static plan, since inventory levels convey information about the state of the season even without a formal belief update, but with learning the gain is 7,608, about 50% larger. Similarly, the reserve is worth 4,495 in the full system, whereas learning-based release of reserve stock adds almost nothing to the one-shot plan when price is fixed. The ingredients are therefore complements, which is a reason to evaluate components in combination as well as separately.

## D. Allocation decisions

Fig. 3 compares the initial stock placement of the static and the scenario-based plan.

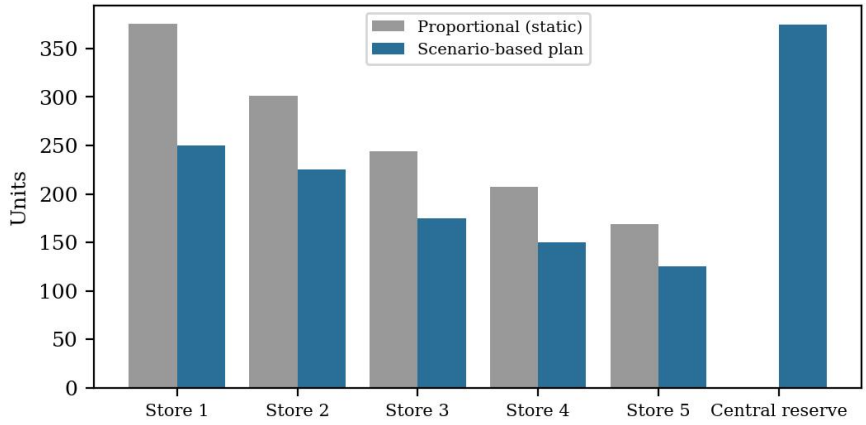



***Fig. 3. Initial stock placement: proportional allocation versus the scenario-based plan.***

The scenario-based plan places 250, 225, 175, 150 and 125 units in the five stores and keeps 375 units, or 28.8% of the total, in the central reserve. The proportional plan, in contrast, sends everything to stores in proportion to forecast, with allocations of 376, 301, 244, 207 and 169 units. The stochastic plan therefore commits only about 70% of the stock at the start and holds the rest for later, which is how it protects itself against the possibility that a store will turn out to be short while another is long. Holding back a large share for later release is the practical implementation of risk pooling.

The opening price chosen by the scenario-based plan is 105, five above the list price. Because stock is tight relative to expected demand, an opening price above the reference level is optimal, and the dynamic policy can then adjust it. The certainty-equivalent policy also selects 105 but without the reserve and without any adjustments, so it holds 206 units unsold on average.

## E. Sensitivity analyses

Fig. 4 and Table 4 show the sensitivity of the profit gain to the level of uncertainty, forecast bias and price sensitivity assumed by the planner.

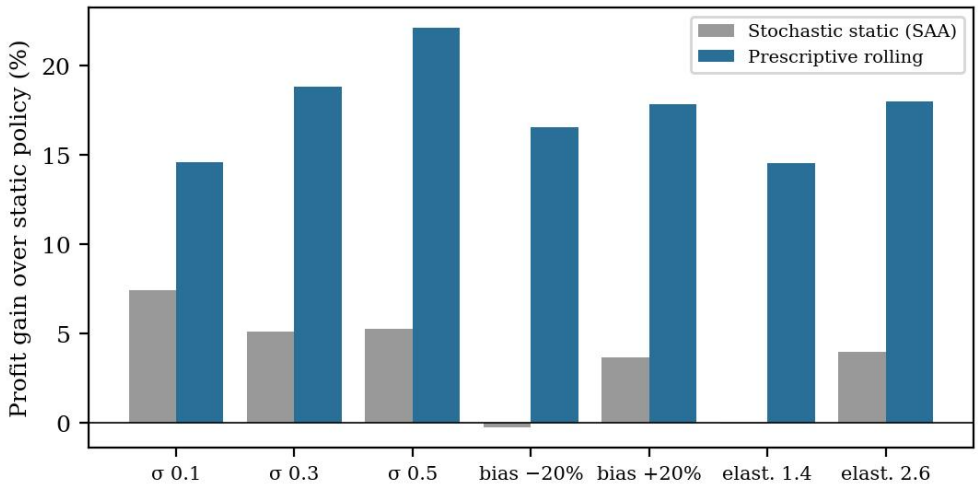



***Fig. 4. Profit gain over the static policy under different environments.***

**Table 4. Mean profit under alternative environments (currency units per season, 400 seasons each).**

| Environment | Static | Stochastic static | Prescriptive | Gain of prescriptive over static |
|---|---|---|---|---|
| Low uncertainty (10%) | 61,890 | 66,470 | 70,931 | +14.6% |
| Base uncertainty (30%) | 56,122 | 58,980 | 66,687 | +18.8% |
| High uncertainty (50%) | 48,174 | 50,699 | 58,835 | +22.1% |
| Forecast too low by 20% | 56,122 | 55,979 | 65,415 | +16.6% |
| Forecast too high by 20% | 56,122 | 58,167 | 66,122 | +17.8% |
| Price sensitivity assumed 1.4 | 56,122 | 56,097 | 64,272 | +14.5% |
| Price sensitivity assumed 2.6 | 56,122 | 58,339 | 66,226 | +18.0% |

Three findings emerge. First, the value of the prescriptive approach rises with uncertainty, from 14.6% when demand uncertainty is low to 22.1% when it is high. This is what theory suggests: the more uncertain the demand, the more there is to learn from early sales and the more valuable flexibility becomes. When uncertainty is low, even a static plan is not far off.

Second, the policy is robust to a biased forecast. When the forecast is wrong by 20% in either direction, the one-shot scenario-based plan loses some or all of its advantage over the static plan, because it commits to an allocation and price based on wrong information (its gain falls to about zero when the forecast is too low). The rolling policy, however, still gains 16.6% to 17.8%, since early sales reveal the error and the policy corrects it. This is probably the most practically important result: forecast errors are the rule and not the exception, and a policy that learns from sales is much less sensitive to them. The opening price chosen under the biased forecasts, 95 and 115 for the low and high forecast, respectively, shows how the bias distorts the plan, and the rolling policy recovers most of the loss through weekly adjustment.

Third, a misjudged price sensitivity also leaves the one-shot plan with no advantage when the planner assumes a sensitivity that is too low (1.4 against a true value of 2.0), whereas the rolling policy keeps a gain of 14.5%. The reason is similar: the weekly review reveals that the price is not clearing the stock as expected and moves it. A large part of the benefit of the policy is therefore not the quality of the initial model but the presence of feedback that corrects it. The gain is somewhat lower than in the base case, which shows that learning about the demand level cannot fully compensate for a wrong price response, since the learning step in this study does not update the price sensitivity itself.

The hybrid also sustained the highest raw event rate, about 900,000 events per second, since the edge servers shared the ingestion work, whereas the cloud-centric design saturated its broker partitions at about 310,000 events per second. The uplink traffic of the hybrid, 7.7 megabits per second, was 92% lower than that of cloud-centric streaming, which consumed 96 megabits per second, and 85% lower than that of the compressed micro-batch configuration. This reduction is the product of extracting features at the edge and forwarding summaries, with raw data around anomalies sent only when needed. The reduction also means that the architecture has headroom to grow the number of sensors on the same uplink.

# 7. DISCUSSION

## *A. Interpretation*

The results support a straightforward message: in a setting where demand is uncertain and stock is limited, the largest gains come not from a better initial plan but from using the information that arrives during the season and having levers (price, reserve) with which to act on it. A better initial plan, obtained by recognising uncertainty, is worth about 5.8% of profit, whereas the combination of learning, re-pricing and replenishment adds a further 12.7%. These figures are specific to the simulated environment, but the ordering is plausible for many seasonal products with short selling windows.

The results also show why simple rules are popular: the markdown rule improves the downside at almost no cost in complexity and may be a good first step, but it cannot exploit high-demand seasons and offers no guidance on allocation.

## *B. Implications for practice*

**Coordinate allocation and pricing.** The ablations show that the levers interact: a reserve is worth much more when price can also be adjusted, and vice versa. Organisations in which supply planning and merchandising are separate functions should therefore design joint decision processes, with shared information and shared performance measures.

**Hold back stock when uncertainty is high.** The plan keeps nearly 29% of stock in reserve; the value of a reserve should rise with uncertainty and fall if replenishment is slow or costly.

**Treat censored sales with care.** A system that records sales as demand will understate demand wherever stock runs out, which is where demand is highest. Even the simple approach used here, excluding stockout weeks from learning, avoids the worst of the problem, and better approaches exist.

**Measure policies on risk as well as average.** Policies with similar mean profit differ in their downside. The certainty-equivalent policy, for example, raises the mean but worsens the downside relative to the static plan. Decision makers should look at the distribution of outcomes and not at the mean alone.

**Keep feedback.** Much of the value arises from the feedback loop, so the speed and quality of sales data and a regular decision cadence deserve investment before more sophisticated initial models.

## *C. Threats to validity*

Simulated environment. All results come from a simulation in which the demand model is stylised. The use of a constant price sensitivity, Poisson noise and a multiplicative season-wide shock may not reflect actual demand, which can exhibit competitor reactions, customer reference prices, stockpiling, cannibalisation across products and promotions.

Well-specified learning. In the base experiment, the planner's scenario model and Bayesian learning match the process that generates demand. This favours the prescriptive policy. The sensitivity analysis relaxes this to some extent, but a fuller test would use a mis-specified demand model, in which the planner's assumptions about the structure of uncertainty are wrong.

Costs and constraints. The study ignores the costs of changing prices, such as customer reaction and operational effort, and assumes that all price changes are possible. It also assumes that the total quantity is fixed in advance and that the reserve can be shipped to any store with the same lead time and cost.

Single product. Real retailers manage thousands of products that share capacity, and substitution between products is not modelled.

Limited tuning. No policy was tuned beyond the choices described; a better-tuned markdown rule might close part of the gap.

## *D. Ethical and practical considerations*

Dynamic pricing raises questions of fairness and transparency, and pricing practices are regulated in some jurisdictions. The policy here changes one price for all customers by at most one step per week, which is mild compared with personalised pricing, but real applications should consider trust and legal compliance and keep planners able to understand and override recommendations.

### *E. Future work*

Several extensions follow naturally. Validation with real retail data, including a comparison against actual decisions, is the most important. The learning model could be extended to estimate price sensitivity alongside demand level, which would address the weakest point revealed by the sensitivity analysis. Better handling of censored demand, multiple products with substitution, supplier-side ordering decisions before the season, and an explicit risk measure in the optimisation are other directions. Finally, integrating the learning and decision steps, so that the demand model is trained for decision quality and not only for forecast accuracy, would connect this work with recent developments in prescriptive analytics [1].

## 8. CONCLUSION

This paper studied the joint problem of allocating a fixed stock between stores and a central reserve and setting prices week by week for a seasonal product under demand uncertainty. A rolling prescriptive policy that plans with scenarios, learns from early sales with a Bayesian update that treats stockouts as incomplete information, re-prices each week within a practical limit and replenishes stores from the reserve was compared with four benchmarks on 1,000 shared simulated seasons. The policy raised mean profit by 19.2% over a static plan and by 12.7% over the best one-shot scenario-based plan, improved the worst-10% average by 33.6%, raised the share of demand served from 82.0% to 95.3% and reduced leftover stock from 154 to 50 units. Dynamic pricing was the largest contributor, followed by the central reserve and by Bayesian learning, and the components were complements. The advantage rose with demand uncertainty and persisted under biased forecasts and a misjudged price sensitivity, because feedback from sales corrected the planner's errors, whereas the one-shot plan lost most of its advantage under the same errors.

The practical message is that in uncertain, short-season settings the value of prescriptive analytics lies more in closing the loop between data, beliefs and action than in the quality of an initial plan. The evidence comes from a stylised simulation with a well-specified learning model, so the size of the gains should be treated as illustrative. The next steps are validation on real data, a learning model that also estimates price sensitivity, and extensions to multiple products and to risk-averse objectives.

## References


1. Pawan Kumar Singh, Scaling Gate-All-Around (GAA) Transistor Architectures for Sub-2nm AI Accelerators Using Atomically Thin 2D Materials, Vol. 33 No. 08 (2024): JOCAAA, Journal Name: Journal of Computational Analysis and Applications, Journal's ISSN: 1521-1398 (Paper),1572-9206 (Online), https://eudoxuspress.com/index.php/pub/article/view/5713
2. Pawan Kumar Singh, Energy-Efficient Task Scheduling in Green Cloud Computing Using Neuromorphic Spiking Neural Networks, Vol. 33 No. 08 (2024): JOCAAA, Journal Name: Journal of Computational Analysis and Applications, Journal's ISSN: 1521-1398 (Paper),1572-9206 (Online), https://eudoxuspress.com/index.php/pub/article/view/5712
3. A. Mahendra Vardhan and S. Sridhar, "Determining False Positive Analysis of Software Vulnerabilities with Predefined Scan Rules using Random Forest Classifier and Decision Tree Technique," 2022 4th International Conference on Advances in Computing, Communication Control and Networking (ICAC3N), Greater Noida, India, 2022, pp. 622-625, DOI: https://doi.org/10.1109/ICAC3N56670.2022.10074458
4. Ronakkumar shah , Cognitive Process Intelligence: Integrating Generative AI and Automated Workflow Discovery, Vol. 33 No. 08 (2024): JOCAAA, Journal of Computational Analysis and Applications, Journal's ISSN: 1521-1398 (Paper),1572-9206 (Online), https://eudoxuspress.com/index.php/pub/article/view/5676
5. Ronakkumar shah, research Trends in Smart Workflow Automation: A Bibliometric Analysis of Emerging Themes, Home / Archives / Vol. 33 No. 05 (2024): JOCAAA, Journal of Computational Analysis and Applications, Journal's ISSN: 1521-1398 (Paper),1572-9206 (Online), https://eudoxuspress.com/index.php/pub/article/view/5675
6. Dip Bharatbhai Pate, Dynamic vs. Static visualizations: Understanding their use cases in big data analysis, International Journal of Science and Research Archive, 2022, 05(02), 392-395., DOI https://doi.org/10.30574/ijsra.2022.5.2.0077

7. Dip Bharatbhai Patel, Building scalable business intelligence systems in the cloud: A technical approach, International Journal of Science and Research Archive, 2021, 02(01), 212-215., Received on 03 February 2021; revised on 17 April 2021; accepted on 19 April 2021, DOI: https://doi.org/10.30574/ijsra.2021.2.1.0047
8. Dip Bharatbhai Patel, Understanding the Impact of Text Analytics on Social Media Sentiment Analysis, Vol. 5 No. 5 (2019) ,September-October, DOI: https://doi.org/10.32628/CSEIT1936413
9. Dip Bharatbhai Patel, Real-Time Data Streaming in BI Dashboards: Opportunities and Challenges, Archives / Vol. 5 No. 4 (2019) : July-August, DOI: https://doi.org/10.32628/CSEIT192142
10. Dip Bharatbhai Patel, the Role of Data Visualization in Modern Intelligence, Archives / Vol. 2 No. 2 / Research Article, Published, 2020-11-28 DOI: https://doi.org/10.32996/jcsts.2020.2.2.4
11. A. Srivastava, "Securing PHI in cloud-based healthcare systems: Challenges, frameworks, and future directions," J. Comput. Anal. Appl., vol. 33, no. 2, Aug. 2024. [Online]. Available: https://www.eudoxuspress.com/index.php/pub/article/view/4349/3190
12. A. Srivastava, "AI-assisted cloud migration of healthcare claims data from legacy systems: A metadata-driven framework," Int. J. Commun. Networks Inf. Secur., vol. 14, no. 3, Nov. 2022. [Online]. Available: https://ijcnis.org/index.php/ijcnis/article/view/8682
13. A. Srivastava, "Enhancing provider and claims data accuracy using artificial intelligence on cloud-native data platforms," J. Comput. Anal. Appl., vol. 31, no. 4, Nov. 2023. [Online]. Available: https://eudoxuspress.com/index.php/pub/article/view/4347/3189
14. A. Srivastava, "A unified framework for secure cloud data management: Integrating governance, metadata intelligence, and privacy controls," J. Inf. Syst. Eng. Manage., vol. 9, no. 2, Apr. 2024. [Online]. Available: https://www.jisem-journal.com/download/45_A_Unified_Framework.pdf
15. A. Srivastava, "Optimizing large-scale data migration for cloud-native architectures," J. Comput. Anal. Appl., vol. 31, no. 2, pp. 652–662, May 2023. [Online]. Available: https://eudoxuspress.com/index.php/pub/article/view/4603
16. Abdul Aziz. “Intelligent Health Analytics: Mastering AI in Healthcare Data”. Acta Scientiae, vol. 25, no. 4, Aug. 2024, pp. 544-55, https://www.periodicos.ulbra.org/index.php/acta/article/view/687. https://doi.org/10.5281/zenodo.21699886
17. Event-Aware Multi-Layer Storage Risk Forecasting for Oracle Database Estates Using HAPF, Raghu Gollapudi, International Journal of Computational and Experimental Science and Engineering, 10(4), 2024 ,10.22399/ijcesen.5183
18. Backup Integrity and Recovery Readiness Assessment for High-Availability Databases Raghu Gollapudi, Computer Fraud & Security, Vol. 2023, Issue 8, pp. 35-48, 10.52710/cfs.1031
19. Risk-Controlled Near-Zero-Downtime Oracle Database Migration Using GoldenGate Raghu Gollapudi, International Journal of Computational and Experimental Science and Engineering, 8(3), 2022 ,10.22399/ijcesen.5382
20. An Automated Risk Scoring Framework for SQL Execution Plan Analysis and Performance Manoj Kumar Kagitha, Automated Security Compliance in Ci/Cd: A Natural Language Processing (NLP) Approach to Detecting Vulnerabilities in Infrastructure-As-Code (Terraform/Bicep),Publication year 2021, Volume 2021, Issue 12, https://computerfraudsecurity.com/index.php/journal/article/view/970
21. Manoj Kumar Kagitha, Hybrid Quantum-Classical Resource Scheduling: A Proposed Architecture for Next-Generation Hyperscale Data Centers, Publication year 2024, Volume 2024, Issue 12, https://computerfraudsecurity.com/index.php/journal/article/view/969
22. Manoj Kumar Kagitha, Self-Healing Cloud Infrastructure: A Hybrid Reinforcement Learning Framework for Predicting Microservice Failures in Azure Kubernetes Service (AKS),Publication year 2023, Volume 2023, Issue 12,https://computerfraudsecurity.com/index.php/journal/article/view/971
23. Manoj Kumar Kagitha, LOAD BALANCING TECHNIQUES IN DISTRIBUTED CLOUD ENVIRONMENTS, (IJESMR - 2017), D.O.I:10.53555/jcr.v6:i1.13380, https://jcreview.com/archives/volume-6/issue-1/13380
24. Manoj Kumar Kagitha, AI-DRIVEN OBSERVABILITY FOR HYPERSCALE COLOCATION DATA CENTERS: A CASE STUDY OF LOG ANALYTICS AND ANOMALY DETECTION PIPELINES, Archives / Vol. 27 No. 7 (2019): JOCAAA,https://www.eudoxuspress.com/index.php/pub/article/view/5053

25. Manoj Kumar Kagitha, GREENOPS IN THE CLOUD: AN AI-DRIVEN TELEMETRY ANALYSIS MODEL FOR REDUCING CARBON FOOTPRINT IN HIGH-AVAILABILITY CLUSTERS, Vol. 28 No. 6 (2020): JOCAAA, https://www.eudoxuspress.com/index.php/pub/article/view/5054
26. Bandari, V. V., & Lekkala, H. B. (2024). AI-based operator behavior monitoring and cost optimization using digital traceability in manufacturing systems. Power System Protection and Control, 52(1), 38–45. https://pspac.info/index.php/dlbh/article/view/342
27. Venumadhav Vavilala, Shankar Balla (2024, May). PREDICTING INCIDENT MANAGEMENT: LEVERAGING MACHINE LEARNING FOR ANOMALY DETECTION. Power System Protection and Control Scopus Q1 Journal. PSPC. https://pspac.info/index.php/dlbh/article/view/270
28. Manikantha Varaprasad Inakollu, (2024, May), FINOPS-Driven Cloud optimization models for enterprise applications, Vol. 52 No. 2 (2024): April-June 2024, 99-110, Power System Protection and Control, ISSN-1674-3415, URL: https://pspac.info/index.php/dlbh/article/view/283,, DOI: https://doi.org/10.46121/pspc.52.2.10
29. Naresh Lokiny. (2022). Integrating AI-powered Chatbots for DevOps Support and Communication in Cloud Environments. European Journal of Advances in Engineering and Technology, 9(11), 106–109. https://doi.org/10.5281/zenodo.13325989
30. Naresh Lokiny, (2021), "Disaster Recovery and Business Continuity Planning in DevOps Cloud with AI", International Journal of Science and Research (IJSR), 10(3), 2024-2027. https://dx.doi.org/10.21275/SR24724151733, https://www.ijsr.net/getabstract.php?paperid=SR24724151733
31. Naresh Lokiny, & Ranganath Nandanampati. (2020). DevSecOps: Integrating Security into DevOps with AI in Cloud. Journal of Scientific and Engineering Research, 7(10), 239–242. https://doi.org/10.5281/zenodo.13348695
32. Naresh Lokiny, & Pradip Reddy. (2021). Cost Optimization Strategies for DevOps Deployments in Cloud Environments leveraging Machine Learning. European Journal of Advances in Engineering and Technology, 8(3), 69–72. https://doi.org/10.5281/zenodo.13325845
33. Jayanth Para, 6G Internet Technology Cyber Threat Notification & Alert System, Vol. 52 No. 4 (2024): October-December 2024, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/170 , DOI: https://doi.org/10.46121/pspc.52.4.9
34. Jayanth Para, AI Based Cloud Computation Observational Method & Process, Vol. 51 No. 4 (2023): October-December 2023, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/171 , DOI: https://doi.org/10.46121/pspc.51.4.3
35. Kaleshwar Aryasomayajula, PREVENTING BIAS IN AI-BASED DECISION-MAKING: ANALYZING TECHNIQUES TO REMOVE UNFAIR PREJUDICE IN ALGORITHMIC OUTCOMES, Vol. 50 No. 2 (2022): April-June 2022, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/294
36. Kaleshwar Aryasomayajula, MODERN DEVELOPMENT PRACTICES: INVESTIGATIONS INTO SERVERLESS COMPUTING, LOW-CODE/NO-CODE PLATFORMS, AND DEVELOPER PRODUCTIVITY IN REMOTE ENVIRONMENTS, Vol. 50 No. 4 (2022): October-December 2022, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/364
37. Kaleshwar Aryasomayajula, Micro services: “A Comparative Analysis of Monolithic vs. Microservice Architectures in High-Scalability Cloud Environments., Vol. 51 No. 2 (2023): April-June 2023 , Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/374
38. Vikas Suresh Bagora, KERNEL-LEVEL PERFORMANCE OPTIMIZATION FOR CARRIER-GRADE BROADBAND AND HIGH-SPEED NETWORKING SYSTEMS, Vol. 47 No. 1 (2019), Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/359
39. Vikas Suresh Bagora, SCALABLE 128G FIBRE CHANNEL AND ETHERNET CONVERGENCE FOR NEXT-GENERATION DATA CENTER NETWORKS, Vol. 50 No. 4 (2022): October-December 2022, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/362
40. Vikas Suresh Bagora, SECURE EMBEDDED NETWORKING ARCHITECTURES USING TRUSTED EXECUTION ENVIRONMENTS IN BROADBAND GATEWAYS, Vol. 52 No. 2 (2024): April-June 2024, Power System Protection and Control, ISSN-1674-3415, https://pspac.info/index.php/dlbh/article/view/363

41. Santhosh Chitraju Gopal Varma, The Evolution of Cloud-Native Architectures: Exploring the Synergy between Kubernetes and Microservices, DOI: https://doi.org/10.63282/3050-9246.IJETCSIT-V1I4P104, Vol. 1 No. 4 (2020), https://ijetcsit.org/index.php/ijetcsit/article/view/108
42. Santhosh Chitraju Gopal Varma, Artificial Intelligence in Cloud Computing: Building Intelligent, Distributed, and Fault-Tolerant Systems, Vol. 3 No. 1 (2022), DOI: https://doi.org/10.63282/3050-9416.IJAIBDCMS-V3I1P105, https://ijaibdcms.org/index.php/ijaibdcms/article/view/78
43. Sunil Karmakar, AI driven Identity & Access Management with Oracle Cloud security, Vol. 30 No. 2 (2022): JOCAAA, https://eudoxuspress.com/index.php/pub/article/view/4118
44. Rakshitha Kiran P., Samitha Khaiyum, Aparna R. Palandye, Apoorva S. D., Arjun R., Akshaya S. A., Annapoorna S., Leveraging LLaMA3 and LangChain for Rapid AI Application Development, Vol. 20 No. 10s (2024), https://doi.org/10.52783/jes.5539, https://journal.esrgroups.org/jes/article/view/5539.
45. Kancharla Vijayalakshmi, Bethapudi Samuel Anand Andrews, Estimation of Cyproheptadine Hydrochloride and Tricholine Citrate Simultaneously in Syrup Dose by Applying Stability Indicating HPLC Methodology, December 31, 2020, Vol. 11, Issue 4, Pages 1813–1819, DOI: 10.26452/ijrps.v11iSPL4.4383, https://ijrps.com/v11si4/3683
46. Vijayalakshmi Kancharla ,Kumar Abbaraju V D N, Andrews B S A, Nagaraju Dasari, Nagu Korupolu, Melphan as Anticancer and New Insights of Drug as Parenteral Dosage form against SARS-CoV-2 by Using Development and Validation , November 18, 2020, Vol. 11, Issue 1, Pages 1437–1446
DOI: 10.26452/ijrps.v11iSPL1.3686, https://ijrps.com/v11si1/2911