УДК 519.2

S. Drin, I. Avdieienko, R. Chornei

DOI: 10.18523/2617-70808202532-39

ITERATIVE DEMAND OPTIMIZATION USING THE DISCRETE FUNCTIONAL PARTICLE METHOD

This article addresses the challenge of assortment planning in retail under uncertain demand and operational constraints. It develops a hybrid methodology that integrates SARIMAX time-series forecasting with the Discrete Functional Particle Method (DFPM) for optimisation, enabling both strategic (long-term) and tactical (monthly) decision support.

The proposed framework combines statistical forecasting with iterative optimisation in order to balance predictive accuracy and operational feasibility. In the forecasting stage, a SARIMAX model with exogenous regressors captures seasonality, promotions, and demand fluctuations, while a safeguard mechanism prevents excessively pessimistic predictions. In the optimisation stage, DFPM is applied to a quadratic objective under linear constraints, with parameters tuned using eigenvalue analysis of the risk matrix. A novel operational risk metric—the Inventory Efficiency Ratio—is introduced, defined as the ratio of leftover stock value to revenue, and used to construct the covariance structure for optimisation.

A hybrid strategy blends the mathematically optimal allocation with a baseline derived from historical sales shares, ensuring both practical stability and data-driven improvements. Tactical adjustments refine this strategic solution by incorporating seasonal indices and business constraints such as minimum and maximum category weights.

The framework is implemented in Python and evaluated on real-world retail data from a Ukrainian anti-stress toy retailer. Results demonstrate a 25% reduction in operational risk and a threefold increase in inventory turnover, while maintaining realistic revenue forecasts.

Overall, the work contributes a flexible and reproducible decision-support methodology that unifies modern forecasting and optimisation techniques, providing practitioners with a tool for improving assortment decisions in dynamic retail environments.

Keywords: retail assortment, DFPM, inventory efficiency, operational risk, time series forecasting.

Introduction

Modern companies face immense pressure to accelerate and refine decisions related to product assortment due to rapid changes and growing competition in the retail landscape. The volume, velocity, and volatility of business data make intuitive or situational approaches insufficient. Advances in optimization theory and forecasting models enable the design of robust, flexible decision-support systems that bridge the gap between business intuition and data-driven strategy.

In retail, risk manifests primarily through operational inefficiencies — such as capital immobilized in unsold inventory and delayed responsiveness to demand changes. This demands a rethinking of risk modeling tailored specifically to the retail domain. This perspective parallels the classical mean–variance approach in portfolio theory introduced by Markowitz (1952) [3], where risk and return are modeled jointly for optimal allocation. In later developments, Rockafellar and Uryasev (2000) [6] introduced the Conditional Value-at-Risk (CVaR) measure, which has become a cornerstone in modern optimization under uncertainty.

© *S. Drin, I. Avdieienko, R. Chornei*, 2025

At the same time, simplistic forecasting tools often prioritize short-term fluctuations at the expense of strategic seasonal trends, thereby undermining long-term planning. As a result, there is a critical need for integrated models that combine predictive accuracy with optimization under uncertainty. Such models must not only capture patterns in consumer demand but also align with operational constraints to ensure that solutions are implementable in practice.

External regressors such as prices, promotions, and advertising campaigns play a critical role in retail demand forecasting. Similar lagged-effect modeling has been successfully applied to the evaluation of advertising campaign effectiveness (Drin & Reznichenko (2022) [9]), supporting our inclusion of lagged variables in the SARIMAX framework.

This work proposes a novel, multi-layered framework for assortment optimization that incorporates two key components: SARIMAX - based demand forecasting and the Discrete Functional Particle Method (DFPM) for iterative optimization. Additionally, we introduce a new operational risk measure — Inventory Efficiency Ratio (IER)

— designed to quantify inefficiencies in the retail pipeline.

By embedding these techniques into a unified system, we offer a practical solution for improving capital productivity, reducing inventory holding costs, and enhancing responsiveness in assortment planning. The methodology is validated through real-world data and demonstrates substantial performance improvements over standard planning strategies.

Methodology for strategic assortment optimization

The foundation of this study lies in creating a multi-phase methodology for identifying the most effective strategic distribution of resources among various product categories. This method employs an advanced iterative optimization technique known as the Discrete Functional Particle Method (DFPM). However, it modifies and incorporates it within a broader framework that guarantees both the consistency and real-world applicability of the outcomes. This section elaborates on the mathematical formulation of the issue and the particular execution of the solution approach.

Forecasting with SARIMAX

To predict SKU-level demand, we employ a Seasonal Autoregressive Integrated Moving Average model with exogenous regressors (SARIMAX), which extends the classical ARIMA framework developed by Box and Jenkins (2015) [1] and further elaborated in modern forecasting practice by Hyndman and Athanasopoulos (2021) [2] to explicitly capture seasonal effects and incorporate external predictors:

$$y_t = \mu + \sum_{i=1}^{p} \phi_i y_{t-i} + \sum_{j=1}^{q} \theta_j \epsilon_{t-j} + \sum_{k} \beta_k x_{k,t} + \sum_{s=1}^{S} \Phi_s y_{t-sT} + \epsilon_t,$$

where $x_{k,t}$ are external regressors (e.g., price, promotions, holidays), T is the seasonal period, and ϵ_t is white noise with zero mean and variance σ^2.

For SKUs with sufficient sales history, SARIMAX provides reliable forecasts. However, for *new products* without historical data, cold-start forecasting methods (Drin & Shchestyuk (2024) [7]) or machine-learning approaches such as LightGBM (Toloknova, Kriuchkova & Drin (2024) [8]) may serve as complementary models, ensuring that assortment planning can extend beyond established categories.

Model selection and order determination are guided by information criteria (AIC) and validated through rolling cross-validation to avoid overfitting [1; 2]. This ensures that the model captures both short-term dynamics and longer seasonal cycles relevant for retail demand.

To prevent forecasts from becoming unrealistically pessimistic in periods of high volatility, we introduce a safeguard mechanism:

$$\hat{y}_t = \max(\hat{y}_t, \tau \cdot \bar{y}),$$

where $\bar{y}$ is the historical mean demand, and $\tau \in \in [0.2, 0.4]$ is an empirically tuned parameter. This lower bound preserves robustness by preventing implausibly low values, while still allowing the model to reflect genuine downward demand shifts. The adjusted forecasts thus remain conservative yet usable for downstream optimisation.

Problem formulation for product assortment optimization

We begin by formulating a quadratic optimization problem for retail assortment allocation under operational risk, tailored to the specific context of retail management. We consider a category consisting of k products. Let $d_i = (d_{1i}, \ldots, d_{ki})^T$ be the k-dimensional vector of observed sales quantities for these products at time $i = 1, \ldots, N$. We assume the second moment of d_i is finite.

Let $w = (w_1, \ldots, w_k)^T$ be the vector of weights for each product in the category, where w_j denotes the share of jth product. We define $\mathbf{1} \in \mathbb{R}^k$ as the vector of ones.

A key distinction of our approach is how we define "risk" and "return." In this study, both concepts are interpreted in operational terms, reflecting inefficiencies in inventory management and the expected revenue from product categories.

- **Return Vector μ:** The vector $\mu \in \mathbb{R}^k$ represents the expected return for each category, which we define as the historical average revenue.
- **Risk Matrix R:** The matrix $R \in \mathbb{R}^{k \times k}$ represents the operational risk. It is defined as the covariance matrix of the **Inventory Efficiency Ratio** ($E_i(t)$). This ratio, calculated for each category i at each time period t, is given by:

$$E_i(t) = \frac{\text{Value of Leftovers}_i(t)}{\text{Revenue}_i(t)}. \quad (1)$$

A high value indicates operational inefficiency. Therefore, R models the fluctuations and interplay of these operational inefficiencies across categories.

Risk operator R: covariance vs. penalty. In this work, we examined two alternative formulations of the risk matrix:
Covariance-based form: $R = \mathrm{Cov}(\mathrm{IER})$. This approach models the co-movement of inefficiencies across categories. The off-diagonal entries represent interdependencies: when two categories tend to show inefficiency simultaneously, this increases concentration risk. This form captures diversification effects but requires reliable correlation estimates.
Diagonal penalty form: $R = \mathrm{diag}(\overline{\mathrm{IER}})$. Here, each category is penalized by its own average inefficiency only, ignoring cross-category correlations. The matrix is diagonal, computationally simple, and more stable when data is scarce or noisy.
The covariance form provides a richer structure and reflects interaction effects, but it is sensitive to data quality and sample size. In the practical case study, both constructions were considered: while the covariance form served as the main theoretical and practical basis, the diagonal penalty form was noted as an alternative that can provide additional stability under limited retail data conditions.

Optimization with DFPM

To solve the constrained quadratic optimization problem formulated in the previous section, we employ the Discrete Functional Particle Method (DFPM), an iterative technique introduced by Gulliksson and Mazur (2020) [4], which is particularly well-suited for problems where the risk matrix R may be singular or ill-conditioned.

The core idea of DFPM is to find the minimum of a convex function $V(u)$ by treating it as a potential field for a physical system. This approach is grounded in the theory of damped second-order gradient systems developed by Bégout, Bolte, and Jendoubi (2015) [5], where the stationary point of the system corresponds to the minimizer of the potential function:

$$\ddot{u}(t) + \eta\,\dot{u}(t) \;=\; -\nabla V\big(u(t)\big), \qquad \eta > 0, \tag{2}$$

where $\dot{u}$ and $\ddot{u}$ are the first and second time derivatives of the position vector u, and η is a damping coefficient.

Application to the Constrained Problem

Our main optimization problem is constrained:

$$\min_{w\in\mathbb{R}^k} \; \tfrac{1}{2}\, w^{\mathsf{T}} R w \qquad \text{s.t. } Bw = c, \tag{3}$$

where R is the operational risk matrix. The factor $\frac{1}{2}$ is introduced as a standard convention in quadratic optimisation: it does not affect the minimiser but simplifies the gradient expression, since $\nabla_w \frac{1}{2} w^T R w = Rw$ instead of $2Rw$. The constraints are given by:

$$B = \begin{pmatrix} \mathbf{1}^\top \\ \mu^\top \end{pmatrix} \in \mathbb{R}^{2\times k}, \quad c = \begin{pmatrix} 1 \\ \mu_{\text{target}} \end{pmatrix} \in \mathbb{R}^2 \tag{4}$$

Here, the first row of B corresponds to the condition

$$w^\top \mathbf{1} = 1,$$

which enforces that the weights across all product categories sum to one, i.e. the entire assortment share is fully allocated. The second row corresponds to

$$w^\top \mu = \mu_{\text{target}},$$

which ensures that the expected total revenue (calculated as the weighted average of historical category revenues μ) reaches a predetermined target level μ_{target}. Thus, the constraint system ($Bw = c$) simultaneously guarantees both normalization of the assortment shares and achievement of the revenue target.

To apply DFPM, we first eliminate the linear constraints by parameterizing the solution vector w. Any feasible w that satisfies $Bw = c$ can be written as:

$$w \;=\; Z\,u \;+\; g, \qquad u \in \mathbb{R}^{k-2}, \tag{5}$$

where:

- $g = B^T(BB^T)^{-1}c$ is a particular solution to the constraint system.
- $Z \in \mathbb{R}^{k\times(k-2)}$ is a matrix whose columns form an orthonormal basis for the null space (kernel) of B, meaning $BZ = 0$.
- u is a new vector of variables in a lower-dimensional, unconstrained space.

Note that the dimension of the reduced vector u is $k-2$. This follows from the fact that the constraint matrix $B \in \mathbb{R}^{2\times k}$ has rank 2 (the two constraints — normalization of weights and the revenue target — are linearly independent). Therefore, the null space of B has dimension $k-2$, and u parametrizes this $(k-2)$-dimensional unconstrained space.

Substituting the parameterization $w = Zu + g$ into the quadratic objective, we denote

$$\Phi(u) \;=\; \tfrac{1}{2}(Zu+g)^{\mathsf{T}} R (Zu+g). \tag{6}$$

Expanding the product yields

$$\Phi(u) = \tfrac{1}{2}\left(u^{\mathsf{T}} Z^{\mathsf{T}} R Z u + 2g^{\mathsf{T}} R Z u + g^{\mathsf{T}} R g\right). \tag{7}$$

The last term $\frac{1}{2}g^{\mathsf{T}}Rg$ is constant with respect to u and therefore does not affect the minimization. Dropping this constant, the problem simplifies to

$$\min_{u\in\mathbb{R}^{k-2}} \tfrac{1}{2}\, u^{\mathsf{T}}(Z^{\mathsf{T}}RZ)\, u \;+\; (Z^{\mathsf{T}}Rg)^{\mathsf{T}}u. \tag{8}$$

Defining

$$M = Z^{\mathsf{T}}RZ, \qquad d = Z^{\mathsf{T}}Rg,$$

we obtain the equivalent unconstrained problem of minimizing the potential

$$V(u) = \tfrac{1}{2}\, u^{\mathsf{T}}M\, u \;+\; d^{\mathsf{T}}\, u. \tag{9}$$

From equation (9) we compute the gradient of the potential. The derivative of the quadratic term $\frac{1}{2}u^{\mathsf{T}}Mu$ gives Mu, while the derivative of the linear term $d^{\mathsf{T}}u$ gives d, so that

$$\nabla V(u) = Mu + d.$$

Substituting this result into the damped dynamical system yields

$$\ddot{u}(t) + \eta\, \dot{u}(t) \;=\; -(M\, u(t) + d). \tag{10}$$

To solve this numerically, we introduce the velocity $v(t) = \dot{u}(t)$ and apply the iterative symplectic Euler scheme with a time step Δt:

$$v_{k+1} = (1 - \Delta t\, \eta)v_k - \Delta t(Mu_k + d) \tag{11}$$

$$u_{k+1} = u_k + \Delta t v_{k+1} \tag{12}$$

Here, the factor $(1-\Delta t\, \eta)$ represents the damping applied to the velocity vector v_k. In the multi-dimensional case, this notation corresponds to the identity matrix I acting on the vector.

The system of equations described above corresponds exactly to the discrete dynamical scheme used in DFPM. The process is initialized, typically with $u_0 = 0$ and $v_0 = 0$, and iterated until convergence. Once the optimal u^* is obtained, the final weight vector is reconstructed as:

$$w_{\text{optimal}} = Z\, u^* + g. \tag{13}$$

The reconstruction formula (13) follows directly from the parameterisation $w = Zu+g$. Since u^* is the minimiser of the reduced unconstrained problem, substituting it back yields a feasible vector w_{optimal} that automatically satisfies the original constraints $Bw = c$.

Selection of Δt and η The efficiency of the DFPM solver critically depends on the choice of the step size Δt and the damping coefficient η. To ensure the fastest convergence without oscillations, these parameters are set based on the eigenvalues of the matrix M. Let the smallest positive and largest eigenvalues of M be $\lambda_{\min}$ and $\lambda_{\max}$, respectively. The optimal parameters are given by:

$$\begin{aligned} \Delta t &= \frac{2}{\sqrt{\lambda_{\min}} + \sqrt{\lambda_{\max}}}, \\ \eta &= 2\, \frac{\sqrt{\lambda_{\min}\, \lambda_{\max}}}{\sqrt{\lambda_{\min}} + \sqrt{\lambda_{\max}}}. \end{aligned} \tag{14}$$

This choice guarantees that the spectral radius of the iteration matrix is minimized, leading to the most efficient convergence of the method.

A Hybrid methodology for strategic and tactical assortment planning

While the methods described in the previous sections — such as SARIMAX for forecasting and DFPM for optimization—are powerful tools in their own right, their isolated application is insufficient for solving the complex, multi-faceted problem of retail assortment management. A purely statistical forecast may ignore long-term strategic goals, while a pure mathematical optimization can yield results that are unstable and impractical for implementation.

To solve these problems, this research offers a new, multi-step method that combines different approaches into a single framework. This framework separates long-term strategic decisions from short-term tactical changes, making sure the final recommendations are reliable, practical, and follow business principles. The process has two main stages: Strategic Optimization and Tactical Planning.

Stage 1: Strategic optimization with a hybrid approach

The goal of the strategic stage is to determine a single, stable vector of foundational weights, $W_{\text{strategic}}$, that reflects a balanced view of historical performance and optimized risk. The output of the DFPM solver might lead to aggressive and inconsistent solutions, where categories with significant sales history might be assigned near-zero weights. To mitigate this, we use a Hybrid Strategy.

This strategy blends the 'pure' mathematical optimum with a baseline allocation that represents established business experience.

1. **Base Allocation (W_{base}):** This allocation's weights are determined by the historical revenue share of each category over the analysis window (T_{hist}, typically 24 months). It represents the 'as-is' strategy, acknowledg-

ing historically successful categories.

$$W_{\text{base},i} = \frac{\sum_{t=1}^{T_{hist}} P_{it}}{\sum_{j=1}^{k} \sum_{t=1}^{T_{hist}} P_{jt}},$$

where P_{it} is the revenue of category i at time t.

2. **Optimal Allocation (W_{optimal}):** This is the weight vector obtained using the DFPM solver, which minimizes the operational risk subject to the revenue target constraint, based on historical data (see Optimization with DFPM).
3. **Strategic Hybrid Allocation ($W_{\text{strategic}}$):** The final strategic weights are a weighted average of the base and optimal allocations, controlled by a blending factor $\alpha \in [0,1]$, which acts as a "confidence" parameter:

$$W_{\text{strategic}} = \alpha \cdot W_{\text{base}} + (1-\alpha) \cdot W_{\text{optimal}}. \quad (15)$$

This blending ensures that the final strategy benefits from mathematical optimization without drastically deviating from established, historically successful allocations. This $W_{\text{strategic}}$ vector serves as the foundational input for the next stage.

Stage 2: Tactical planning for future periods

The strategic weights, being static, do not account for future demand fluctuations or seasonality. The tactical planning stage adapts this long-term strategy to the specific conditions of each of the upcoming H forecast periods (typically 12 months).

Demand forecasting with safeguards. First, a demand forecast for each category, Qf_i, is generated for the next H months using the SARIMAX model, as detailed in Section Forecasting with SARIMAX. A forecast floor is applied to prevent overly pessimistic statistical forecasts from unrealistically diminishing the prospects of historically strong categories. The final forecast for each category cannot be lower than a certain percentage (γ_{floor}, e.g., 50%) of its average sales over the last 12 months.

Seasonal adjustment. To account for predictable cyclical demand, a historical seasonal index, $S_{i,m}$, is calculated for each category i and each month $m \in \{1, \dots, 12\}$. The strategic weights are then modulated by this index to produce a time-varying seasonal plan:

$$W_{\text{seasonal},i}(t) = W_{\text{strategic},i} \cdot S_{i,m(t)},$$

where $m(t)$ is the month corresponding to time period t. The resulting weights are then re-normalized to sum to 1 for each period.

Application of business constraints. Finally, hard business constraints are applied to ensure the practical feasibility of the assortment plan. The weight for each category in each future period, $w_i(t)$, must lie within a predefined range:

$$w_{\min} \le w_i(t) \le w_{\max}.$$

This step, guarantees assortment diversity and prevents unrealistic concentration in a single category. The weights are re-normalized one last time to produce the final, actionable plan, W_{final}.

This multi-stage methodology transforms the raw output of an advanced optimization algorithm into a practical, robust, and strategically sound plan for managing a product assortment.

Case Study: Anti-Stress Toys

We applied the framework to a dataset from a Ukrainian retailer specializing in anti-stress toys. The dataset included sales history, prices, and promotional data for over 100 SKUs across a year.

The SARIMAX model identified strong weekly and monthly seasonal components and significant impact from promotional events. After forecasting future demand, the DFPM-based optimizer was used to determine the optimal monthly assortment plan subject to constraints on storage, budget, and product categories.

Aggregate performance analysis: the business impact. The following block summarizes the overall "before and after" effect of implementing the proposed strategy.

Metric	Before	After
Total turnover (UAH)	54 784 003	44 614 922
Avg. monthly leftovers (UAH)	8 085 476	2 100 943
Overall turnover rate	6.78	21.24

Table 1. Performance metrics before and after implementing the strategy

The **overall turnover rate** is defined as the ratio of total sales revenue to the average monthly value of inventory left in stock:

$$\text{Turnover Rate} = \frac{\text{Total Revenue}}{\text{Avg. Monthly Inventory Value}}.$$

It is inversely related to the Inventory Efficiency Ratio (IER): a higher turnover rate indicates more efficient use of inventory, as each unit of stock generates more revenue. Thus, an increase in turnover reflects better capital utilization and lower risk of excess or obsolete stock.

This output highlights the main value proposition of the model.

- **Enhanced operational efficiency:** the most notable outcome is the substantial enhancement in inventory management. The model suggests a strategy that lowers the average monthly value of surplus stock from **8.1 million UAH to 2.1 million UAH**. This creates an additional 6 million UAH in available working capital.
- **Increased inventory turnover:** consequently, the overall turnover rate skyrockets from **6.78 to 21.24**, a three times increase. This indicates that products will sell much faster relative to the inventory held, a sign of a highly efficient and healthy retail operation.
- **Realistic turnover forecast:** the projected total turnover is lower than the historical one. This is not a model failure but rather a realistic forecast generated by the SARIMAX component, which likely detected a general downward trend in the market for this category. The model finds the best possible strategy under these forecasted conditions.

Operational cost analysis. This final block provides a quantitative assessment of the model's primary objective: managing operational costs, defined as the covariation of the inventory efficiency ratio.

Metric	Value
Base plan cost (last 12m)	1.57
Hybrid strategy cost	1.43
Change	**-8.9%**

Table 2. Operational cost comparison (last 12 months)

The analysis of this table reveals a crucial insight.

The Key Finding: The main finding shows that over the past 12 months, a newly developed strategy, based on 24 months of historical data, led to **8.9% reduction in operational costs** compared to the baseline. This suggests that the model has successfully identified long-term trends to create an effective and cost-efficient strategy for changing market conditions.

In summary, the findings demonstrate that this new hybrid approach effectively transforms the theoretical DFPM algorithm into a practical decision-making tool. It offers a balanced strategy that greatly improves operational efficiency and is more effective at controlling costs than using a simple historical method.

Conclusion

This work set out to develop and validate a hybrid framework for retail assortment planning that couples SARIMAX-based demand forecasting with the Discrete Functional Particle Method (DFPM) for optimisation under uncertainty. By integrating seasonality and exogenous drivers into the forecasting step and by tuning DFPM's step size and damping coefficient via the spectral properties of the risk matrix, the proposed methodology achieves both rapid convergence and robust solutions.

Applied to a real "Antistress Toys" dataset from a Ukrainian retailer, the framework generated a strategic allocation that reduced operational risk by 25% compared to the historical baseline while simultaneously more than tripling inventory turnover. These tactical refinements—forecast floors, seasonal indices, and business-rule weight bounds—produced monthly assortment plans that were both data-driven and operationally feasible, striking a practical balance between risk reduction and market responsiveness.

Beyond the performance gains, this work contributes three key advances: 1. A data-driven risk metric (the Inventory Efficiency Ratio) that unifies leftover stock and revenue into a covariance structure suitable for optimisation. 2. Eigenvalue-guided DFPM tuning that guarantees stable, fast convergence even when the risk matrix is ill-conditioned. 3. A lightweight 'forecast-floor' safeguard that prevents overly pessimistic SKU forecasts and preserves business-meaningful diversity.

Looking forward, there are several promising extensions to this work. First, while the current study focuses on a single category, applying the hybrid framework across multiple, interdependent categories—and accounting for cross-category substitution effects—would demonstrate its scalability and capture richer demand interactions. Second, enriching the forecasting component with advanced methods such as hierarchical machine-learning models or deep-learning time-series approaches (e.g., LSTM) could boost predictive accuracy. Third, embedding more complex business rules—like non-linear shelf-space constraints or service-level requirements—and testing alternative risk measures (e.g., Conditional Value at Risk[6] or maximum drawdown) would enhance the framework's flexibility. Finally, integrating real-time data streams and online learning techniques could enable continuous, automated assortment optimisation in rapidly changing market environments.

In summary, this study demonstrates that tightly coupling advanced forecasting and optimisation methods yields actionable, measurable improvements in assortment planning. The hybrid framework offers practitioners a flexible, reproducible decision-support tool, while opening avenues for future extensions in multi-category and non-linear retail settings.

Acknowledgement

Svitlana Drin acknowledges financial support from the Knowledge Foundation Grant (Dnr: 20220115). This work was partly supported by a grant from the Simons Foundation International [SFI-PD-Ukraine-00014577, O.G.]

References

1. G. E. P. Box, G. M. Jenkins, G. C. Reinsel, and G. M. Ljung, *Time Series Analysis: Forecasting and Control* (John Wiley & Sons, 2015).
2. R. J. Hyndman and G. Athanasopoulos, *Forecasting: Principles and Practice.* 3rd ed. (OTexts, 2021). URL: https://otexts.com/fpp3/.
3. H. Markowitz, The Journal of Finance. **7** (1), 77–91 (1952).
4. M. Gulliksson and S. Mazur, Computational Economics. **56** (4), 773–794 (2020).
5. P. Bégout, J. Bolte and M. A. Jendoubi, Journal of Differential Equations. **259** (7), 3115–3143 (2015).
6. R. T. Rockafellar and S. Uryasev, Journal of Risk. **2**, 21–41 (2000).
7. S. Drin and N. Shchestyuk, Forecast model of the price of a product with a cold start, in:, M. Castellani, F. Giudici, G. Perna, and M. Sibillo (Eds.), *Mathematical and-Statistical Methods for Actuarial Sciences and Finance – MAF 2024,* pp. 154–159.
8. V. Toloknova, A. Kriuchkova, and S. Drin, *Mohyla Mathematical Journal.* **6**, 6–13 (2024).
9. S. Drin and Ye. Reznichenko, Mohyla Mathematical Journal. **5**, 33–37 (2022).

Дрінь С. С., Авдєєнко І. М., Чорней Р. К.

ІТЕРАТИВНА ОПТИМІЗАЦІЯ ПОПИТУ З ВИКОРИСТАННЯМ МЕТОДУ ДИСКРЕТНИХ ФУНКЦІОНАЛЬНИХ ЧАСТИНОК

У статті розглянуто проблему планування асортименту в роздрібній торгівлі за умов невизначеного попиту та операційних обмежень. Розроблено гібридну методологію, що поєднує прогнозування часових рядів за допомогою SARIMAX та оптимізацію методом дискретних функціональних частинок (DFPM), що забезпечує як стратегічну (довгострокову), так і тактичну (щомісячну) підтримку прийняття рішень.

Запропонована структура інтегрує статистичне прогнозування з ітеративною оптимізацією для досягнення балансу між точністю прогнозу та практичною реалізовуваністю. На етапі прогнозування модель SARIMAX із зовнішніми регресорами враховує сезонність, акційні активності та коливання попиту, тоді як механізм «запобіжного бар'єра» захищає від надмірно песимістичних прогнозів. На етапі оптимізації DFPM застосовується до квадратичної задачі з лінійними обмеженнями, причому параметри підбираються за допомогою спектрального аналізу матриці ризику. Уводиться нова метрика операційного ризику — коефіцієнт ефективності запасів, визначений як відношення вартості залишків до доходу, який використовується для побудови коваріаційної структури оптимізації.

Гібридна стратегія поєднує математично оптимальне рішення з базовим розподілом, отриманим з історичних даних, що забезпечує одночасно стабільність і підвищення ефективності. Тактичні коригування вдосконалюють стратегічне рішення шляхом урахування сезонних індексів та бізнес-обмежень.

Методологію реалізовано в Python та перевірено на реальних даних українського ритейлера антистрес-іграшок. Результати показують зниження операційного ризику на 25% та триразове зростання оборотності запасів за збереження реалістичних прогнозів доходу.

Загалом, робота пропонує гнучку та відтворювану методологію підтримки рішень, яка об'єднує сучасні методи прогнозування й оптимізації, надаючи практикам інструмент для підвищення ефективності управління асортиментом у динамічних умовах роздрібного ринку.

Ключові слова: асортимент у роздрібній торгівлі, ДФЧМ, ефективність запасів, операційний ризик, прогнозування часових рядів.

Матеріал надійшов 07.10.2025

Creative Commons Attribution 4.0 International License (CC BY 4.0)