

Drin, Svitlana; Avdieienko, Ivan; Chornei, Ruslan

Working Paper

Iterative demand optimization using the discrete functional particle method

Working Paper, No. 17/2025

Provided in Cooperation with:

Örebro University School of Business

Suggested Citation: Drin, Svitlana; Avdieienko, Ivan; Chornei, Ruslan (2025) : Iterative demand optimization using the discrete functional particle method, Working Paper, No. 17/2025, Örebro University School of Business, Örebro

This Version is available at:

<https://hdl.handle.net/10419/339200>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



WORKING PAPER 17/2025 (STATISTICS)

Iterative Demand Optimization Using the Discrete Functional Particle Method

Svitlana Drin, Ivan Avdieienko and Ruslan Chornei

ISSN 1403-0586

Örebro University School of Business
SE-701 82 Örebro, Sweden

Iterative Demand Optimization Using the Discrete Functional Particle Method

Svitlana Drin^{1,2}, Ivan Avdieienko², and Ruslan Chornei²

¹School of Business, Örebro University, 70182 Örebro, Sweden

²Department of Mathematics, National University of Kyiv-Mohyla Academy, 04070 Kyiv, Ukraine

December 5, 2025

Abstract

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.

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.

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.

Keywords: Retail assortment planning, SARIMAX forecasting, DFPM, inventory efficiency ratio, operational risk optimisation.

1 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.

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.

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.

1.1 Literature review

Research on retail assortment planning has developed along several related directions. A first important stream considers assortment planning as a central problem in revenue management and retail operations. Early contributions and later reviews emphasise the role of assortment decisions for retailer profitability and link assortment models to consumer choice theory, space limitations, and category management Kök et al. (2015). Within this stream, several authors integrate assortment planning with other operational decisions. For example, joint models of assortment and shelf-space optimisation show how product selection interacts with physical display constraints in stores Hübner et al. (2016), while combined assortment and inventory models for vertically differentiated products demonstrate that higher demand uncertainty may lead to deeper assortments and different inventory allocations across quality levels Transchel et al. (2016). Dynamic assortment planning with demand learning further extends these ideas and shows how retailers can adapt offered assortments when customer preferences are only partially known Sauré and Zeevi (2013).

A second, closely related literature studies demand forecasting for retail and supply chain management. Classical time series models such as ARIMA and SARIMA remain standard tools for business forecasting because of their transparency and well under-

stood statistical properties Box et al. (2015); Hyndman and Athanasopoulos (2021). In many applications, these models are extended to SARIMAX in order to incorporate exogenous variables such as prices, promotions, holidays, and special events. Empirical studies show that including such explanatory variables improves forecast accuracy compared to pure seasonal models, especially for products with strong promotion effects or pronounced seasonality Arunraj and Ahrens (2016); Falatouri et al. (2021); Saputra and Kumar (2020). Despite these advances, forecasting and optimisation are often treated as separate modules: a forecasting model produces point forecasts which then serve as inputs to an optimisation problem, with only limited joint design of the two components.

Risk modelling provides a third methodological pillar. In finance, the classical mean-variance framework of Markowitz interprets risk as the variance of portfolio returns and leads to quadratic optimisation problems under linear constraints Markowitz (1952). Later, coherent risk measures such as Conditional Value-at-Risk (CVaR) were introduced and formulated in an optimisation-friendly way, for example through linear programming techniques Rockafellar and Uryasev (2000a). These ideas have inspired a large body of work on portfolio selection with singular or ill-conditioned covariance matrices, including regularisation techniques and iterative numerical schemes Gulliksson and Mazur (2020). In the retail and operations management context, risk is more often captured indirectly through service-level constraints, stockout penalties, or safety-stock formulas. Practical key performance indicators such as inventory turnover, days of inventory on hand, and inventory efficiency ratios are widely used by practitioners to monitor how effectively inventory is converted into sales Silver et al. (2017). However, these metrics usually play a descriptive role and are rarely embedded as primary risk measures in formal optimisation models.

The Discrete Functional Particle Method (DFPM) and related dynamical-systems-based methods form a fourth relevant strand. DFPM was proposed as a way to solve equations and optimisation problems by recasting them as damped second-order dynamical systems with a suitable potential function Gulliksson and Mazur (2020). The underlying theory of such damped gradient systems, including conditions for convergence and rates of decay, has been analysed in the mathematical literature on second-order gradient systems Bégout et al. (2015). Subsequent work applied DFPM-type approaches to ill-conditioned optimisation problems and rank-deficient covariance matrices, where they offer robust convergence to numerically stable solutions Gulliksson and Mazur (2020). Despite this potential, DFPM has been used mainly in numerical analysis and financial optimisation, with very limited applications to business analytics or retail assortment problems.

The present study is positioned at the intersection of these four strands of literature. It combines a SARIMAX-based demand forecasting module with a DFPM-based optimisation module in a single framework tailored to assortment planning. Forecasts explicitly account for seasonality and promotional effects, while the optimisation problem is formulated in terms of an inventory efficiency metric that links the value of leftovers to achieved revenue. In this way, the proposed approach exploits well-established time series methods Box et al. (2015); Hyndman and Athanasopoulos (2021), connects to risk-based optimisation concepts from portfolio theory Markowitz (1952); Rockafellar and Uryasev (2000a); Gulliksson and Mazur (2020), and responds to managerial practice where inventory turnover and related indicators are central measures of performance Silver et al. (2017). To the best of our knowledge, there are only few, if any, integrated frameworks

that combine SARIMAX forecasting, DFPM-based optimisation, and inventory-efficiency-oriented risk metrics for retail assortment planning, which highlights the novelty and contribution of the present work.

This study contributes to the literature on retail analytics and optimisation by proposing an integrated framework that tightly couples SARIMAX-based demand forecasting with DFPM-based assortment optimisation. In contrast to most existing approaches where forecasting and optimisation are treated as separate stages, our model directly links statistically grounded demand projections with the optimisation process. We introduce an inventory-efficiency-oriented risk representation based on the ratio of leftover inventory to realised revenue, which provides a managerial interpretation of risk in terms of capital locked in unsold stock. From a numerical perspective, the study extends the application of the Discrete Functional Particle Method to retail assortment planning and demonstrates its stability and computational efficiency under realistic constraints. The framework is further enhanced by a hybrid allocation strategy that combines the DFPM solution with historical revenue shares to improve robustness and practical implementability. An empirical case study using data from a Ukrainian retailer confirms significant improvements in inventory efficiency and turnover, highlighting the practical relevance of the proposed approach.

The remainder of this paper is organised as follows. In Section 2, we present the overall methodological framework of the proposed approach, including the demand forecasting and optimisation components. Section 2.1 describes the SARIMAX model used for demand forecasting, together with the treatment of seasonality, promotional effects, and the forecast floor mechanism. In Section 2.3, we introduce the optimisation model based on the Discrete Functional Particle Method and formulate the inventory-efficiency-oriented risk operator and the associated constraints.

Section 3 is devoted to the empirical case study based on data from a Ukrainian retailer. It details the dataset, model calibration, and implementation procedure. The main numerical results and business performance indicators are reported and discussed in Section 3, where we analyse the impact of the proposed strategy on inventory levels, turnover, and capital efficiency.

Finally, Section 4 summarises the main findings, discusses managerial implications, and outlines limitations of the current study as well as directions for future research.

2 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.

2.1 Forecasting with SARIMAX

To predict SKU-level demand, we employ a Seasonal Autoregressive Integrated Moving Average model with exogenous regressors (SARIMAX). This model extends ARIMA by

explicitly capturing seasonal effects and incorporating 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 .

Model selection and order determination are guided by information criteria (AIC) and validated through rolling cross-validation to avoid overfitting Box et al. (2015); Hyndman and Athanasopoulos (2021). 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 $\tau \in [0.2, 0.4]$ 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. Forecast accuracy was additionally evaluated using standard error measures such as MAE and RMSE.

2.2 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 $d_i = (d_{1i}, \dots, d_{ki})^T$ be the k -dimensional vector of observed sales quantities for these products at time $i = 1, \dots, N$. We assume the second moment of d_i is finite.

Let $w = (w_1, \dots, w_k)^T$ be the vector of weights for each product in the category, where w_j denotes the share of j th 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:** $R = \text{Cov}(\text{IER})$. 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: $R = \text{diag}(\overline{\text{IER}})$. 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.

2.3 Optimization with DFPM

To solve the constrained quadratic optimization problem formulated in the previous section, we employ the Discrete Functional Particle Method (DFPM), described by Guliksson and Mazur (2020). This iterative method 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. The minimum of the function corresponds to the stationary point of a damped dynamical system, described by the second-order differential equation Bégout et al. (2015):

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

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

2.4 Application to the Constrained Problem

Our main optimization problem is constrained:

$$\min_{w \in \mathbb{R}^k} \frac{1}{2} w^\top R w \quad \text{s.t. } B w = c, \quad (3)$$

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 $\nabla_w \frac{1}{2} w^\top R w = R w$ instead of $2R w$. The constraints are given by:

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

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:

$$w = Zu + g, \quad u \in \mathbb{R}^{k-2}, \quad (5)$$

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

$$\Phi(u) = \frac{1}{2}(Zu + g)^\top R(Zu + g). \quad (6)$$

Expanding the product yields

$$\Phi(u) = \frac{1}{2} (u^\top Z^\top RZ u + 2g^\top RZ u + g^\top Rg). \quad (7)$$

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

$$\min_{u \in \mathbb{R}^{k-2}} \frac{1}{2} u^\top (Z^\top RZ) u + (Z^\top Rg)^\top u. \quad (8)$$

Defining

$$M = Z^\top RZ, \quad d = Z^\top Rg,$$

we obtain the equivalent unconstrained problem of minimizing the potential

$$V(u) = \frac{1}{2} u^\top M u + d^\top u. \quad (9)$$

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

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

Substituting this result into the damped dynamical system yields

$$\ddot{u}(t) + \eta \dot{u}(t) = -(Mu(t) + d). \quad (10)$$

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 :

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

$$u_{k+1} = u_k + \Delta t v_{k+1} \quad (12)$$

Here, the factor $(1 - \Delta t \eta)$ represents the damping applied to the velocity vector v_k . In the multidimensional 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:

$$w_{\text{optimal}} = Zu^* + g. \quad (13)$$

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 λ_{\min} and λ_{\max} , respectively. The optimal parameters are given by:

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

In practice, the extreme eigenvalues of M were computed using a standard symmetric eigensolver with numerical tolerance 10^{-6} . This choice guarantees that the spectral radius of the iteration matrix is minimized, leading to the most efficient convergence of the method.

2.5 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.

- **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, acknowledging historically successful categories.

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

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

- **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.
- **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:

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

In the empirical study, α was selected through grid search in the interval $[0.2, 0.8]$ by minimising the out-of-sample operational cost over a validation window.

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, $Q_{f,i}$, is generated for the next H months using the SARIMAX model, as detailed in

Section 2.1. 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:

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

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.

3 Empirical study

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.

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

Table 1: Aggregate performance before and after implementation of the proposed strategy

The relative reduction in average monthly inventory leftovers amounts to 74.0%, while the turnover rate increases by 213.3% compared to the historical baseline.

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 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.

4 Conclusions

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 RiskRockafellar and Uryasev (2000b) or maximum drawdown) would enhance the framework’s flexibility. Finally, integrating real-time data streams and on-line 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.

References

- Arunraj, N. S. and Ahrens, D. (2016). A hybrid seasonal autoregressive integrated moving average and quantile regression approach for daily food demand forecasting. *International Journal of Production Economics*, 170:321–335.
- Bégout, P., Bolte, J., and Jendoubi, M. A. (2015). On damped second-order gradient systems. *Journal of Differential Equations*, 259(7):3115–3143.

- Box, G. E. P., Jenkins, G. M., Reinsel, G. C., and Ljung, G. M. (2015). *Time Series Analysis: Forecasting and Control*. John Wiley & Sons, 5 edition.
- Falatouri, T., Tirkolaee, E. B., and Ivanov, D. (2021). Demand forecasting in retail supply chains using sarimax and lstm models. *Annals of Operations Research*.
- Gulliksson, M. and Mazur, S. (2020). An iterative approach to ill-conditioned optimal portfolio selection. *Computational Economics*, 56(4):773–794.
- Hübner, A., Kuhn, H., and Sternbeck, M. G. (2016). Demand and supply chain planning in grocery retail: An operations planning framework. *International Journal of Retail & Distribution Management*, 44(6):512–530.
- Hyndman, R. J. and Athanasopoulos, G. (2021). *Forecasting: Principles and Practice*. OTexts, 3 edition.
- Kök, A. G., Fisher, M. L., and Vaidyanathan, R. (2015). *Assortment Planning: Review of Literature and Industry Practice*. Springer.
- Markowitz, H. (1952). Portfolio selection. *The Journal of Finance*, 7(1):77–91.
- Rockafellar, R. T. and Uryasev, S. (2000a). Optimization of conditional value-at-risk. *Journal of Risk*, 2:21–41.
- Rockafellar, R. T. and Uryasev, S. (2000b). Optimization of conditional value-at-risk. *Journal of Risk*, 2:21–41.
- Saputra, R. and Kumar, S. (2020). Retail demand forecasting using sarimax with intervention analysis. *Journal of Retailing and Consumer Services*, 52:101924.
- Sauré, D. and Zeevi, A. (2013). Optimal dynamic assortment planning with demand learning. *Manufacturing & Service Operations Management*, 15(3):387–404.
- Silver, E. A., Pyke, D. F., and Thomas, R. P. (2017). Inventory and production management in supply chains. *Operations Research*. Inventory turnover metrics and efficiency indicators.
- Transchel, S., Minner, S., Kallrath, J., and Becker, C. (2016). Joint optimization of assortment and inventory with vertically differentiated demand. *European Journal of Operational Research*, 254(1):172–187.