Article

Not peer-reviewed version

Attributing Inventory Performance via Shapley-Based Counterfactual Decomposition

Lu Xu *

Posted Date: 31 March 2026

doi: 10.20944/preprints202603.2499.v1

Keywords: inventory management; Shapley value; counterfactual decomposition; performance attribution

 doi:10.20944/preprints202603.2499.v1

Article

Attributing Inventory Performance via Shapley-Based Counterfactual Decomposition

Lu Xu

Independent Researcher, USA; shufesifexulu@gmail.com

Abstract

Inventory systems are typically evaluated using aggregate performance metrics such as out-of-stock and average inventory. In supply chain management, it is important to understand the underlying reasons for a period's performance— specifically, how previous inventory management decisions, such as order placement, lead to the result and what their contributions are. Traditional methods are often restrictive and cannot be applied to broader cases. This paper proposes a Shapley-based decomposition framework that attributes the realized performance gap between an observed inventory policy and an optimized benchmark to individual order decisions. To numerically illustrate this method, we consider a finite-horizon periodic-review inventory system with fixed lead time and integer order quantities. Performance is evaluated using a composite metric that balances stockout frequency and normalized average inventory. Given an observed order vector and an optimized reference vector obtained via a genetic algorithm, we compute the average marginal contribution of each decision using a Monte Carlo Shapley estimator. The resulting decomposition allocates the total performance gap between the observed policy and optimal policy while accounting for nonlinear interactions and eliminating ordering bias. Compared to traditional methods, the proposed approach directly explains a realized benchmark-relative performance difference and is applicable to integer-constrained, non-differentiable, and simulation-based inventory systems. Managerially, it enables transparent inventory management performance evaluation and effective root-cause analysis.

Keywords: inventory management; Shapley value; counterfactual decomposition; performance attribution

---

1. Introduction

Inventory management systems are typically evaluated using aggregate performance metrics such as service level, stockout frequency, fill rate, average inventory, or composite cost functions [1]. In practice, firms assess the quality of an ordering policy over a finite horizon by examining summary indicators that balance service performance and inventory efficiency. While such evaluation metrics are informative at the system level, they provide limited insight into which specific order decisions contributed most to performance deviations from a desired benchmark. To understand the underlying drivers of inventory performance during a given period, it is necessary to conduct retrospective analysis and quantify the contribution of past decisions to realized outcomes.

Such a task, although practically important, is analytically challenging. For example, in periodic-review systems with fixed lead time, order decisions are made at discrete epochs and propagate dynamically through future inventory trajectories. Because inventory evolves recursively, and current stock depends on prior stock levels, arrivals, and realized demand, the impact of a single decision is inherently intertemporal. A suboptimal decision at one review epoch may influence multiple downstream periods through altered stock levels, stockout occurrences, and holding patterns. Consequently, it is difficult to attribute overall performance to individual decisions in a principled manner.

Several quantitative approaches exist for analyzing the impact of inventory management decisions. Gradient-based sensitivity analysis evaluates local marginal effects [2], while variance-based global

 doi:10.20944/preprints202603.2499.v1

sensitivity methods decompose output variability across input factors [3,4]. Although valuable, these methods have limitations for retrospective reason decomposition. Sensitivity analysis is inherently local and may not apply in integer-constrained or non-differentiable systems, whereas variance-based methods rely on strong assumptions. As discussed in detail later, these approaches are not ideally suited for attributing realized deviations from an optimized policy.

To address these limitations, this paper proposes a novel and broadly applicable decomposition framework that attributes the gap between an observed inventory performance metric and an ideal benchmark to individual decisions. The proposed method is inspired by the Shapley value from cooperative game theory [5], which measures the average marginal contribution of a player across all possible coalitions. In our setting, the "players" are decision variables. For each decision, we compute its average marginal effect on the evaluation metric by comparing counterfactual policies in which subsets of other decisions are replaced by their ideal counterparts. This procedure yields an attribution measure that captures, in expectation, how much the evaluation metric would improve if a particular order were aligned with its ideal value, while accounting for interaction effects among decisions.

To illustrate the applicability of the decomposition method, we consider a finite-horizon periodic-review inventory system with fixed lead time and lost sales. The system is evaluated using a composite metric that combines the proportion of periods experiencing stockouts and normalized average inventory. The decision variables are the order quantity at each ordering epoch. Applying the proposed method, we quantitatively calculate the contribution of each order decision that leads to the gap of observed inventory evaluation results and the optimal results.

The contribution of this work is threefold. First, we introduce a decision-epoch attribution framework for inventory systems that bridges dynamic inventory simulation and cooperative-game-based decomposition. Second, we formalize the decomposition under both uniform-subset averaging and Shapley-weighted schemes, and implement scalable Monte Carlo estimators suitable for realistic planning horizons. Third, we demonstrate how the framework can be integrated with numerical optimization methods, such as genetic algorithms, to generate a reference policies and quantify the marginal impact of deviations from optimal ordering. The proposed decomposition provides managerial insight beyond aggregate performance evaluation. Instead of merely reporting that a policy underperforms relative to a benchmark, the method identifies which specific ordering decisions contribute most to service degradation or inventory.

The remainder of the paper is organized as follows. Section 2 does a literature review. Section 3 describes the existing traditional approaches. Section 4 introduces the proposed decomposition methodology and Monte Carlo estimation procedure, and discusses its intuition. Section 5 presents computational experiments. Section 6 discusses managerial implications, and Section 7 concludes.

2. Literature Review

This paper builds on three streams of literature: (i) cooperative game theory and Shapley-value-based allocation in inventory systems, (ii) decomposition and sensitivity analysis for complex simulation models, and (iii) Shapley-based attribution methods used for explaining the output of black-box functions.

2.1. Cooperative Game Theory in Inventory Management

A substantial body of work applies cooperative game theory to inventory systems, primarily to allocate the gains from cooperation (e.g., risk pooling, inventory centralization, and joint replenishment) among participating agents. For example, ref. [6] uses the Shapley value to allocate savings generated by inventory pooling in multi-retailer supply chains, showing that Shapley-based allocations can coordinate the supply chain and yield individually rational outcomes [7]. further develops this idea in the context of inventory centralization, using Shapley-based allocations to distribute profits in centralized systems.

Beyond these specific models, surveys and reviews synthesize the broader class of *inventory games* and highlight the role of solution concepts such as the core and the Shapley value in cost/profit allocation problems [8]. Related research on deterministic multi-period replenishment also formulates cooperative games for lot-sizing and joint ordering [9].

2.2. Sensitivity Analysis and Variance-Based Decomposition

A second stream explores the decomposition of model outputs via sensitivity analysis. Local (derivative or finite-difference) sensitivity methods quantify marginal responsiveness near a baseline configuration but are often limited in non-smooth or integer-constrained systems. Global sensitivity analysis addresses this by apportioning output uncertainty to input factors, typically via variance-based indices. Ref. [3] provides a foundational treatment of global sensitivity indices and Monte Carlo estimators for nonlinear models. Refs. [4,10] develop practical estimators and best practices for computing total effect indices, emphasizing interaction effects and computational efficiency, and provide a comprehensive synthesis in a widely used monograph.

These approaches are valuable for understanding how *uncertainty* in inputs drives *variance* in outputs, but they generally do not answer the counterfactual question at the core of our study: given a realized underperformance relative to a benchmark, *which particular order decisions contributed most to the observed gap*? Our decomposition targets the realized performance difference $y(\mathbf{x}^*) - y(\mathbf{x})$ rather than variance, and it does so without requiring distributional assumptions over the decision variables.

2.3. Shapley-Based Attribution for Explaining Complex Functions

A third, adjacent stream uses Shapley values for attributing the output of a complex function to its inputs, most prominently in interpretable machine learning. Ref. [11] propose SHAP, a unified framework that uses Shapley values to assign feature attributions for individual predictions and discusses desirable axioms such as consistency. The theoretical origin of this attribution principle traces back to Shapley's seminal axiomatization of a value for n-person games [5]. Our method is conceptually aligned with Shapley-based attribution in that it computes average marginal contributions under coalitions of "corrected" inputs.

3. Existing Approaches

Various methodologies have been proposed in the inventory and operations literature to analyze the drivers of system performance. However, these approaches differ fundamentally in what they decompose and how attribution is assigned. This section situates our proposed decomposition framework relative to existing methods.

3.1. Domain Knowledge-Based Decomposition

In practice, a widely used approach to diagnosing inventory performance relies on domain knowledge and heuristic reasoning. Managers often analyze operational metrics such as order placement quality (e.g., whether orders are systematically under- or over-sized), demand forecast accuracy, and realized lead-time delays to infer the underlying causes of inventory outcomes. By examining these factors, it is possible to form a qualitative understanding of the drivers behind observed performance.

While such domain knowledge-based analysis is intuitive and operationally convenient, it has several limitations. First, it does not provide a principled quantitative framework for evaluating the relative importance of different factors. Multiple drivers may simultaneously influence inventory outcomes, and their effects are often interdependent, making it difficult to disentangle their individual contributions. Second, this approach typically relies on manually selected metrics and human interpretation. As a result, the conclusions may depend on subjective judgment and may vary across analysts. In complex systems with nonlinear dynamics and intertemporal interactions, such heuristic reasoning may overlook subtle but important effects. Finally, domain knowledge-based methods do not naturally support counterfactual analysis. While they can indicate potential causes, they do not quantify how much performance would improve if a particular factor were corrected.

In contrast, the proposed decomposition framework provides a systematic and quantitative attribution mechanism. By evaluating counterfactual scenarios and averaging marginal contributions

 doi:10.20944/preprints202603.2499.v1

across different interaction contexts, it enables a consistent and a reproducible assessment of the relative importance of individual decisions or underlying drivers.

3.2. Gradient-Based Analysis

A common approach in inventory control and optimization is local sensitivity analysis, where the impact of decision variables on system performance is evaluated through marginal effects [2,12]. Let

$$\mathbf{x} = (x_1, x_2, \ldots, x_p)$$

denote the vector of decisions, where x_j represents the decision quantity(assuming all decision variables are numerical) at decision epoch j, and let $y(\mathbf{x})$ denote the corresponding inventory performance metric.

In this setting, the sensitivity of performance with respect to decision x_j is typically approximated using derivatives or finite differences, such as

$$\frac{\partial y}{\partial x_j} \quad \text{or} \quad \frac{y(x_1, \ldots, x_j + \epsilon, \ldots, x_p) - y(\mathbf{x})}{\epsilon}, \tag{1}$$

where ϵ is a small perturbation. These quantities measure the marginal change in system performance resulting from an infinitesimal or small change in the decision variable x_j.

Such methods provide useful directional information and are widely used in continuous optimization to study the stability and local behavior of optimal solutions [2]. In particular, they can indicate whether increasing or decreasing a specific order quantity is likely to improve performance in the neighborhood of The current policy.

However, several limitations arise when applying gradient-based analysis of inventory systems. First, these methods are inherently local: they describe the responsiveness of the system near a fixed operating point $\mathbf{x}$ and do not capture how decisions influence performance globally across the decision space. As a result, they are not suitable for explaining large deviations between two policies. Second, inventory systems often exhibit nonlinear and non-smooth dynamics. For example, lost-sales models involve max operators, inventory truncation at zero, and discrete event transitions . In such settings, the mapping from decisions to performance may not be differentiable, and derivatives may be ill-defined or unstable. Third, decision variables are frequently non-numeric or constrained in practice, which further limits the applicability of differential methods. Finally, gradient-based methods do not capture interaction effects among decisions. In dynamic inventory systems, the impact of one decision may depend on other decisions through state propagation over time. While higher-order derivatives could, in principle, account for such interactions, they are rarely tractable in simulation-based or high-dimensional settings.

3.3. Variance-Based Decomposition and Surrogate-Model-Based Attribution

Another class of approaches for explaining inventory performance is global sensitivity analysis. Unlike gradient-based methods, which characterize local responsiveness near a fixed operating point, global sensitivity methods study how variation in the input decisions contributes to variation in the output over a broader decision space [3,4]. In variance-based sensitivity analysis, the objective is not to decompose a realized gap between two policies, but rather to decompose the output variance $\mathrm{Var}(y(\mathbf{x}))$ into components associated with individual decision variables and their interactions.

A standard formulation is the functional ANOVA decomposition, under which the model output is represented as

$$y(\mathbf{x}) = f_0 + \sum_{j=1}^{p} f_j(x_j) + \sum_{i<j} f_{ij}(x_i, x_j) + \cdots + f_{1,\ldots,p}(x_1, \ldots, x_p), \tag{2}$$

where f_0 is a constant term, $f_j(x_j)$ captures the main effect of decision x_j, and higher-order terms represent interaction effects among decisions [4]. This decomposition induces a variance decomposition of the form

$$\mathrm{Var}(y) = \sum_{j=1}^{p} V_j + \sum_{i<j} V_{ij} + \cdots , \tag{3}$$

where $V_j = \mathrm{Var}\big(\mathbb{E}[y \mid x_j]\big)$ is the main-effect contribution of decision x_j, and V_{ij} denotes the pairwise interaction contribution between x_i and x_j [3,4]. The corresponding first-order Sobol index is $S_j = \frac{V_j}{\mathrm{Var}(y)}$, which measures the share of output variance explained by the decision x_j alone [3]. In principle, such quantities can be estimated by generating many decision combinations, simulating the inventory system under those combinations, and applying Monte Carlo estimators [3,4].

When the inventory simulator is computationally expensive or the decision space is high-dimensional, one may instead fit a surrogate model $\hat{y}(\mathbf{x}) \approx y(\mathbf{x})$. for example using a neural network, and then studying variable importance through the surrogate. In that case, one may use variance-based importance measures or feature-attribution tools such as SHAP values to quantify how the surrogate prediction depends on each input variable [11]. This strategy is attractive because it can reduce repeated simulation costs and leverage modern predictive models.

Despite their usefulness, these methods have several limitations for the retrospective decomposition problem considered in this paper. First, variance-based methods explain *performance variability* over a hypothetical distribution of decisions, rather than the realized performance difference between an observed policy and a benchmark policy. Second, these methods require the specification of a sampling distribution over decision variables. In practical inventory problems, this is not straightforward, because the decision space may be constrained and the decisions may be dependent. For example, later order quantities may be limited by earlier replenishment choices, budget constraints, or policy rules. Classical functional ANOVA and Sobol decompositions are most naturally formulated under independent inputs, and dependent inputs complicate both the interpretation and estimation of variance-based effects [4]. If simulated decision combinations do not respect the actual dependence structure, the resulting importance measures may be misleading. Third, the surrogate-model approach introduces an additional layer of approximation error. Reliable attribution then depends not only on the decomposition method itself, but also on the predictive accuracy of the fitted model. In high-dimensional settings, this may require a large number of simulated decision combinations for training and validation, which can be computationally demanding. Moreover, if the sampled combinations are infeasible or operationally implausible, the surrogate may learn relationships that do not reflect the true decision environment.

4. Decomposition Methodology

4.1. Mathematical Formulation

Let $\mathbf{x}$ denote the observed decision vector and let $\mathbf{x}^*$ denotes an optimized (ideal) decision vector obtained via numerical optimization. Define

$$y = y(\mathbf{x}), \qquad y^* = y(\mathbf{x}^*), \tag{4}$$

and the performance gap

$$\Delta = y^* - y. \tag{5}$$

Our objective is to decompose Δ into contributions associated with each decision variable x_j.

Let $N = \{1, 2, \ldots, p\}$ denote the index set of decisions. For any subset $A \subseteq N$, define a hybrid decision vector

$$\mathbf{x}^{(A)} = \begin{cases} x_j^*, & \text{if } j \in A, \\ x_j, & \text{if } j \notin A. \end{cases} \tag{6}$$

Thus, decisions in A are replaced by their ideal counterparts.

doi:10.20944/preprints202603.2499.v1

Define the value function

$$v(A) = y(\mathbf{x}^{(A)}). \tag{7}$$

Observe that

$$v(\varnothing) = y(\mathbf{x}), \qquad v(N) = y(\mathbf{x}^*). \tag{8}$$

The decomposition problem becomes attributing

$$v(N) - v(\varnothing) \tag{9}$$

across the p decisions.

For decision $j \in N$, define its marginal contribution relative to subset $A \subseteq N \setminus \{j\}$ as

$$\Delta_j(A) = v(A \cup \{j\}) - v(A). \tag{10}$$

The canonical Shapley value for decision j is

$$\phi_j = \sum_{A \subseteq N \setminus \{j\}} \frac{|A|!(p - |A| - 1)!}{p!} [v(A \cup \{j\}) - v(A)]. \tag{11}$$

The Shapley value satisfies the efficiency property

$$\sum_{j=1}^{p} \phi_j = v(N) - v(\varnothing), \tag{12}$$

ensuring that the total performance gap is fully allocated.

Exact computation requires evaluating 2^p coalitions, which becomes computationally infeasible for large p. We therefore employ a Monte Carlo permutation-based estimator.

Let π be a random permutation of N. For each permutation:

1. Start from baseline $\mathbf{x}$.
2. Sequentially replace decisions by their ideal values following π.
3. Record the marginal effect when decision j is switched:

$$\Delta_j^{(\pi)} = v(S_\pi(j) \cup \{j\}) - v(S_\pi(j)), \tag{13}$$

where $S_\pi(j)$ is the set of decisions preceding j in permutation π.

The Monte Carlo Shapley estimator is

$$\hat{\phi}_j = \frac{1}{M} \sum_{m=1}^{M} \Delta_j^{(\pi_m)}, \tag{14}$$

where M is the number of sampled permutations. As $M \to \infty$, $\hat{\phi}_j$ converges to ϕ_j.

As an alternative decomposition, we define

$$\tilde{\phi}_j = \mathbb{E}_{A \subseteq N \setminus \{j\}} [v(A \cup \{j\}) - v(A)], \tag{15}$$

where each subset is assigned equal probability. In practice, we approximate this expectation via Monte Carlo sampling of subsets.

Unlike the Shapley value, this scheme does not guarantee that

$$\sum_j \tilde{\phi}_j = v(N) - v(\varnothing), \tag{16}$$

but it provides an interpretable average marginal effect.

doi:10.20944/preprints202603.2499.v1

4.2. Properties of the Decomposition

The proposed decomposition inherits key properties from the Shapley value in cooperative game theory. For completeness, we state these properties in the context of our formulation.

Proposition 1. *Let $v(A) = y(\mathbf{x}^{(A)})$ be the value function defined over subsets of decision indices. The Shapley-based contribution ϕ_j satisfies:*

- ***Efficiency: Contributions sum to the total gap***

$$\sum_{j=1}^{p} \phi_j = v(N) - v(\varnothing).$$

- ***Symmetry: If two decisions have identical effects in all coalitions, they receive identical contributions.*** *If for two decisions i and j,*

$$v(A \cup \{i\}) = v(A \cup \{j\}) \quad \forall A \subseteq N \setminus \{i, j\},$$

  then $\phi_i = \phi_j$.
- ***Null player: If correcting a decision never changes performance, its contribution is zero.*** *If*

$$v(A \cup \{j\}) = v(A) \quad \forall A \subseteq N \setminus \{j\},$$

  then $\phi_j = 0$.

These properties follow directly from the linearity and symmetry of the Shapley value construction. The efficiency property arises from telescoping marginal contributions across permutations. Symmetry follows from identical marginal effects across all coalitions, and the null player property holds when marginal contributions are zero in all contexts. See [5] for the original derivation.

4.3. Intuition Behind the Shapley-Based Contribution

The Shapley-style contribution used in our framework can be interpreted as the expected marginal improvement obtained when a particular decision is aligned with its ideal value, averaged across all possible interaction contexts.

In a dynamic decision system like inventory management, the impact of a single decision depends on other decisions made before and after it. For example, correcting an early under-order may have little effect if subsequent orders compensate for the shortfall. Conversely, the same correction may substantially reduce stockouts if later decisions are also aligned with the benchmark policy. Thus, the marginal effect of correcting decision j is inherently context-dependent.

Formally, the marginal contribution of decision j depends on the subset A of other decisions already corrected:

$$\Delta_j(A) = v(A \cup \{j\}) - v(A). \tag{17}$$

Different subsets A may produce different marginal effects due to interaction among decisions.

The quantity ϕ_j can be interpreted as:The expected reduction in the performance gap if order decision j were aligned with the ideal policy, accounting for interaction with other decisions.

If ϕ_j is large in magnitude, decision j plays a substantial role in explaining the gap between the observed and ideal policies. If ϕ_j is close to zero, correcting that decision alone has a limited impact once interactions are averaged.

One may view the ideal policy as a fully corrected system. Starting from the observed policy, imagine gradually correcting decisions in random order. Each time a decision is corrected, the performance improves by some amount. The Shapley contribution assigns to each decision its average improvement across all possible correction sequences.

 doi:10.20944/preprints202603.2499.v1

4.4. Comparison with Traditional Decomposition Approaches

The proposed decomposition framework differs from traditional approaches primarily in the object being explained. Rather than analyzing local responsiveness, output variability, or heuristic operational indicators, our method directly decomposes the realized benchmark-relative performance gap between observed decision vector and the optimized benchmark policy. This distinction is important because, in practice, managers are often less concerned with abstract average behavior over hypothetical decision spaces than with understanding why the current policy underperformed relative to an attainable target.

The detailed comparison with each existing tradition method is as follows:

- **vs.domain knowledge-based analysis:** the proposed method provides a more rigorous and reproducible attribution rule. Domain knowledge can often identify plausible drivers of poor performance, such as forecast error, lead-time delay, or under-ordering, but it typically does not provide a principled mechanism for quantifying their relative importance when several factors interact. Our method addresses this limitation by assigning each decision or driver a formal contribution value based on counterfactual comparisons with the benchmark policy. As a result, the analysis becomes less reliant on subjective judgment and more suitable for consistent managerial evaluation and retrospective diagnosis.
- **vs. gradient-based analysis:** the proposed approach is global rather than local. Gradient-based methods measure marginal responsiveness near a fixed operating point, and are therefore informative for small perturbations around the current policy. However, they do not explain large deviations between the observed policy and an optimized benchmark. In addition, many inventory systems involve integer decisions and non-smooth state transitions, which limit the interpretability or even the existence of classical derivatives. In contrast, our decomposition does not rely on differentiability and remains directly applicable in simulation-based, integer-constrained environments.
- **vs. the variance-based decomposition:** the proposed method is more closely aligned with managerial diagnosis. Variance-based approaches explain how uncertainty in decision variables contributes to the variability of performance over a hypothetical distribution of input combinations. While this is useful for understanding uncertainty propagation, it does not answer the decision-relevant question of which past decisions explain the realized performance gap in a specific planning horizon. Moreover, variance-based methods suffer from possible model misspecification. Our method avoids these issues by focusing only on the observed policy, the benchmark policy, and structured hybrid policies formed by transitioning decisions from observed to benchmark values.

Overall, the proposed method is better suited to the retrospective reason-decomposition problem studied in this paper. It is benchmark- relative, interaction-aware, globally interpretable, and applicable to nonlinear, non-differentiable, and integer-constrained inventory systems. These properties make it more informative than domain knowledge-based diagnosis, more globally relevant than gradient-based analysis, and more managerially meaningful than variance-based decomposition for explaining realized inventory performance.

4.5. Comparison with Expectation-Based Attribution

To further explain the advantages of the proposed decomposition method, we comparee the method with an alternative, more traditional perspective for defining the contribution of each decision variable. Such a traditional approach is based on comparing the realized performance with an expected baseline. Specifically, one may define the contribution of decision x_j as

$$\phi'_j = y(\mathbf{x}) - \mathbb{E}_{x_{-j}}\left[y(x_j, x_{-j})\right], \tag{18}$$

 doi:10.20944/preprints202603.2499.v1

where x_{-j} denotes all decision variables except x_j, and the expectation is taken over a distribution of possible values for the remaining decisions.

While this formulation is conceptually appealing in some contexts, it is less suitable for the inventory performance decomposition problem considered in this paper.

First, the two approaches differ fundamentally in their reference points. The expectation-based formulation attributes the difference between the realized performance $y(\mathbf{x})$ and an average performance across hypothetical decision combinations. In contrast, the proposed method attributes the difference between the realized performance and an optimized benchmark, i.e.,

$$y(\mathbf{x}^*) - y(\mathbf{x}).$$

From a managerial perspective, the benchmark-relative comparison is more meaningful, as decision makers are typically concerned with how far the current policy deviates from an attainable optimal or target policy, rather than from an abstract average over all possible decisions.

Second, the expectation-based approach requires specifying a distribution over x_{-j} and evaluating the performance over all possible combinations of decision variables. In high-dimensional settings, this leads to significant computational challenges. Moreover, many such combinations may be operationally infeasible (e.g., violating capacity or policy constraints), which can distort the resulting attribution. In contrast, the proposed decomposition considers a structured set of counterfactual scenarios in which decisions are transitioned from their observed values to their benchmark values. This reduces the computational burden and ensures that all evaluated combinations remain interpretable and relevant to the underlying decision problem.

5. Numerical Experiment

This section illustrates the decomposition method via a numerical example.

5.1. Problem Setup

Consider a finite-horizon periodic-review inventory system over periods $t = 0, 1, \ldots, T-1$. Orders are placed every x periods, resulting in p order decision epochs. Let the order decision vector be

$$\mathbf{x} = (x_1, x_2, \ldots, x_p), \tag{19}$$

where each $x_j \in \mathbb{Z}_{\geq 0}$ denotes the nonnegative integer order quantity placed at decision epoch j. Lead time is fixed at L periods.

Let demand $D(t)$ be exogenous and realized over the horizon. The end-of-period inventory evolves according to

$$S(t) = \max\{0,\ S(t-1) + \text{Arrival}(t) - D(t)\}, \tag{20}$$

where arrivals satisfy

$$\text{Arrival}(t) = O(t-L). \tag{21}$$

We evaluate performance using a composite metric

$$y(\mathbf{x}) = w \cdot \text{OOS}(\mathbf{x}) + (1-w) \cdot \frac{\overline{S}(\mathbf{x})}{\overline{D}}, \tag{22}$$

where:

- $\text{OOS}(\mathbf{x})$ is the proportion of periods with stockout,
- $\overline{S}(\mathbf{x})$ is the average inventory level,
- $\overline{D}$ is the average demand,
- $w \in [0, 1]$ is a weighting parameter.

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 31 March 2026 doi:10.20944/preprints202603.2499.v1

5.2. Parameters

Table 1 summarizes the key parameters used in the inventory simulation, genetic algorithm optimization, and Monte Carlo Shapley estimation.

Table 1. Model and Algorithm Parameters.

Inventory System Parameters		
Parameter	Description	Value
T	Planning horizon length (periods)	52
x	Review interval (periods)	4
p	Number of order decisions	T/x
L	Fixed lead time (periods)	2
S_0	Initial inventory level	10
$D(t)$	Demand process	Deterministic example sequence
w	Weight on stockout proportion	0.6
$q_{\min}$	Minimum order quantity	0
$q_{\max}$	Maximum order quantity	80
Genetic Algorithm Hyperparameters		
Parameter	Description	Value
Population size	Number of candidate solutions	140
Generations	Number of evolution iterations	220
Elite fraction	Proportion retained each generation	0.10
Tournament size	Tournament selection parameter	3
Crossover rate	Probability of crossover	0.9
Mutation rate	Probability of mutation per gene	0.25
Mutation σ	Std. dev. of Gaussian mutation step	6.0
Monte Carlo Shapley Estimation		
Parameter	Description	Value
Permutations (M)	Number of random permutations	3000
Subset samples	Samples per decision (uniform scheme)	4000
Estimator	Shapley estimator type	Permutation-based MC

6. Illustrative Computational Results

To illustrate the proposed decomposition framework, we implemented the integer-constrained genetic algorithm and Monte Carlo attribution procedure described in Section 3 on a stylized periodic-review inventory problem.

6.1. Experimental Setup

The planning horizon was set to $T = 52$ periods, with order decisions made every $x = 4$ periods, resulting in $p = 13$ decision epochs. The lead time was fixed at $L = 2$ periods, and the initial inventory level was $S_0 = 10$. The evaluation metric was defined as

$$y(\mathbf{x}) = w \cdot \mathrm{OOS}(\mathbf{x}) + (1 - w) \cdot \frac{\overline{S}(\mathbf{x})}{\overline{D}},$$

with weight parameter $w = 0.6$.

For illustration, demand was generated deterministically as

$$D(t) = 20 + (t \bmod 6), \qquad t = 0, 1, \ldots, 51.$$

The observed order vector was specified as

$$\mathbf{x} = (35, 25, 30, 40, 22, 28, 33, 29, 37, 31, 26, 34, 30).$$

Preprints.org (www.preprints.org) | NOT PEER-REVIEWED | Posted: 31 March 2026 doi:10.20944/preprints202603.2499.v1

The benchmark order vector $\mathbf{x}^*$ was obtained using the integer genetic algorithm described earlier, with population size 140, 220 generations, mutation rate 0.25, crossover rate 0.9. See Table 1 for parameter details. Decision quantities were constrained to be nonnegative integers in the range $[0, 80]$.

6.2. Benchmark Optimization Results

The genetic algorithm returned the following ideal benchmark vector:

$$\mathbf{x}^* = (45, 41, 49, 45, 41, 49, 45, 41, 49, 45, 41, 49, 45).$$

The resulting evaluation values were:

$$y(\mathbf{x}) = 0.489108, \qquad y(\mathbf{x}^*) = 0.402573.$$

Thus, the optimized benchmark reduced the evaluation metric by

$$y(\mathbf{x}^*) - y(\mathbf{x}) = -0.086535.$$

Table 2 summarizes the main performance statistics.

Table 2. Summary of illustrative computational results.

Metric	Observed policy	Ideal benchmark
Evaluation metric y	0.489108	0.402573
Stockout proportion	0.7500	0.5000
Average inventory	2.1923	5.7500
Average demand	22.4231	22.4231

The results indicate that the benchmark policy improves overall performance primarily through a reduction in stockout frequency, while increasing average inventory. This is consistent with the structure of the objective, which balances service performance against inventory holding.

6.3. Monte Carlo Decomposition Results

We then applied the Monte Carlo Shapley decomposition using 3000 random permutations. The estimated contributions summed exactly (up to numerical rounding) to the total performance gap:

$$\sum_{j=1}^{13} \hat{\phi}_j = y(\mathbf{x}^*) - y(\mathbf{x}) = -0.086535.$$

Table 3 reports the estimated contribution of each order decision epoch.

Table 3. Monte Carlo Shapley contributions by decision epoch.

Epoch j	Actual order	Ideal order	Order change	Contribution
1	35	45	10	-0.008108
2	25	41	16	-0.006050
3	30	49	19	-0.005020
4	40	45	5	-0.009823
5	22	41	19	-0.005020
6	28	49	21	-0.004334
7	33	45	12	-0.007422
8	29	41	12	-0.007422
9	37	49	12	-0.007422
10	31	45	14	-0.006736
11	26	41	15	-0.006393
12	34	49	15	-0.006393
13	30	45	15	-0.006393
Total				-0.086535

A more negative contribution indicates that aligning that decision with the benchmark policy yields a larger expected reduction in the evaluation metric. In this example, the fourth decision epoch had the largest individual contribution, followed by the first epoch and the seventh through ninth epochs. This suggests that early and mid-horizon order choices were especially important in explaining the gap between the observed and benchmark policies.

6.4. Interpretation

Several observations emerge from this illustrative experiment. First, the decomposition produces a complete allocation of the benchmark-relative performance gap across decision epochs. Second, the contribution values are heterogeneous, implying that not all order decisions are equally important for explaining underperformance. Third, the method provides diagnostic information beyond aggregate policy comparison: although the benchmark policy increases inventory on average, the decomposition reveals which specific order decisions account for the improvement in the overall evaluation metric.

Because this example is intentionally stylized, the numerical values are illustrative rather than empirical. Nevertheless, the experiment demonstrates that the proposed method can be implemented end-to-end: an optimized reference policy can be constructed, the performance gap can be quantified, and that gap can then be decomposed into interpretable decision-level contributions.

7. Managerial Implications

The proposed decomposition framework has several practical implications for inventory governance, performance evaluation, and operational diagnosis. By attributing the realized performance gap between an observed ordering policy and an ideal benchmark to individual decision epochs, the method translates aggregate system performance into actionable managerial insights.

The framework enables quantitative and transparent evaluation of inventory management performance at the decision level. Traditional performance metrics (e.g., service level, stockout rate, or total cost) are aggregated over time and do not distinguish which specific order decisions contributed to underperformance. As a result, accountability is often diffuse.

By computing the contribution of each order decision to the overall performance gap, the proposed method provides a principled mechanism for attributing performance outcomes across decision epochs. Because the decomposition is based on Shapley-style marginal contributions, it accounts for interaction effects among decisions and avoids order bias in attribution. This property supports a fair evaluation framework, in which inventory managers or staff can be assessed based on the quantified impact of their decisions relative to a benchmark policy.

Such transparency can enhance incentive design, facilitate post-period performance reviews, and support more objective managerial assessments.

Also, the decomposition framework can be extended beyond direct order quantities to analyze the underlying drivers of ordering decisions. In practice, order quantities are often determined by rules that depend on demand forecasts, safety stock calculations, and anticipated lead times. Moreover, realized arrivals may deviate from planned lead times due to logistical delays.

Let the order decision at epoch j be expressed as a function

$$x_j = f(\hat{D}_j, \hat{L}_j, \theta),$$

where $\hat{D}_j$ denotes demand forecasts, $\hat{L}_j$ denotes expected lead time, and θ represents policy parameters.

By redefining the "players" in the decomposition as components such as forecast inputs or realized lead-time deviations, the same counterfactual framework can be applied to attribute performance gaps to forecasting accuracy or delivery delays. For example, it is possible to replace realized demand forecasts with ideal or unbiased forecasts while holding other inputs fixed. It is also possible to replace realized lead times with planned lead times to measure the contribution of delay.

 doi:10.20944/preprints202603.2499.v1

In this way, the method extends naturally from decision-level attribution to root-cause analysis, enabling managers to distinguish whether underperformance is primarily driven by forecasting error, execution delays, or discretionary ordering behavior.

8. Summary

8.1. Conclusions

This paper proposes a Shapley-based decomposition framework for attributing inventory performance gaps to individual order decisions. We introduce a counterfactual contribution mechanism that allocates the realized performance difference between an observed policy and an optimized benchmark across decision epochs.

The key innovation of the framework lies in treating decisions as contributors within a cooperative game defined over dynamic inventory evaluation. By averaging marginal effects across all possible correction contexts, the Shapley-based decomposition provides a straightforward insight into each decision's contribution to the gap between the desired and the observed evaluation results.

Compared to traditional knowledge-based reason analysis, sensitivity analysis, and variance-based decomposition, the proposed approach directly addresses a distinct diagnostic question: *which specific decisions explain the realized gap between two policies?* This focus on counterfactual gap attribution provides a decision-level perspective that complements existing inventory evaluation tools. This approach is also distinct from the traditional idea of Shapley value, with the former being more relevant to business applications and less computational burden.

From a managerial standpoint, the framework enhances transparency and accountability in inventory governance. It enables quantitative assessment of decision contributions, supports root-cause analysis of forecasting error or lead-time delay, and facilitates continuous policy improvement relative to optimized benchmarks. Because the method does not rely on differentiability or convexity, it can be integrated into modern analytics platforms built on simulation or data-driven optimization.

Several avenues for future research remain open. Extensions to multi-echelon systems, stochastic lead times, and adaptive learning-based policies would broaden applicability. Further investigation into benchmark selection under multiple optimal solutions and constraint-aware attribution in interdependent decision spaces may strengthen theoretical foundations. Scalable variance-reduction techniques for Monte Carlo Shapley estimation also represent a promising direction.

Overall, this work contributes a novel bridge between dynamic inventory modeling and cooperative-game-based attribution. By transforming aggregate performance evaluation into granular, interaction-aware diagnostics, the proposed method advances both the analytical understanding and practical management of inventory systems.

8.2. Limitations and Future Research

While the proposed decomposition framework provides a principled mechanism for attributing inventory performance gaps to individual decision epochs, several methodological and practical limitations deserve discussion.

8.2.1. Non-uniqueness of the Ideal Benchmark

A central component of the framework is the ideal decision vector $\mathbf{x}^*$, obtained by minimizing the evaluation metric. However, in nonlinear and integer-constrained systems, the optimization problem may admit multiple solutions achieving the same optimal performance level y^*.

When multiple optimal solutions exist, the Shapley-based attribution is benchmark-relative and may depend on which optimal solution is selected as $\mathbf{x}^*$. Different optimal vectors can induce different marginal contribution structures, even if they yield identical performance levels.

This issue does not invalidate the decomposition but changes its interpretation. The method should be viewed as attributing the performance gap relative to a *chosen reference policy*. Several approaches can mitigate ambiguity:

 doi:10.20944/preprints202603.2499.v1

- **Regularized selection:** Select among optimal solutions the one closest to the observed policy (e.g., minimizing $\|\mathbf{x}^* - \mathbf{x}\|$).
- **Averaged benchmark:** Compute attributions relative to multiple optimal solutions and average the resulting contributions.
- **Managerially constrained benchmark:** Restrict the optimization to a policy class reflecting operational practice, thereby reducing multiplicity.

Future research may explore theoretical conditions under which the decomposition is invariant to the choice of optimal benchmark.

8.2.2. Interdependence and Feasibility Constraints Among Decisions

The framework treats order decisions as components of a vector $\mathbf{x} \in \mathbb{Z}^p_{\geq 0}$. In some settings, however, decisions may be interdependent because of budget limits, capacity constraints, service requirements, or state-dependent action spaces, all of which are common in inventory and replenishment problems [13,18]. When feasible decisions lie in a constrained set $\mathcal{X}$ that is not a Cartesian product of independent intervals, counterfactual replacement of subsets of decisions may produce infeasible hybrid vectors $\mathbf{x}^{(A)}$.

Two approaches can address this issue.

1. **Feasibility-preserving projection.** After constructing $\mathbf{x}^{(A)}$, one may project it onto the feasible region $\mathcal{X}$ using a deterministic repair or projection rule. The value function then becomes

$$v(A) = y\big(\Pi_{\mathcal{X}}(\mathbf{x}^{(A)})\big),$$

where $\Pi_{\mathcal{X}}$ denotes a projection operator. This idea is consistent with the broader optimization literature, where projection and repair procedures are widely used to recover feasible solutions under constraint violations [16,17].

2. **Policy-parameter decomposition.** Instead of decomposing raw order quantities directly, one may define the “players” as independent policy parameters, such as base-stock levels, reorder points, or safety stock multipliers. This is often natural in inventory control because many operational policies are parameterized in terms of a small number of interpretable control variables, such as base-stock or (s, S) policy parameters [13–15]. Because these parameters may be independent even when realized orders are not, this approach avoids many infeasible counterfactual combinations.

Thus, while dependence among decisions complicates interpretation, the decomposition framework remains applicable provided that the feasibility is handled explicitly. Importantly, the Shapley mechanism itself still accounts for interaction effects among decisions; the main issue is not the existence of interaction, but ensuring that the counterfactual coalitions being evaluated remain operationally meaningful.

Funding: This research is not funded.

Institutional Review Board Statement: This article does not contain any studies with animals/human participants performed by any of the authors.

References

1. Silver, Edward Allen, David F. Pyke, and Rein Peterson. Inventory management and production planning and scheduling. Vol. 3. New York: Wiley, 1998.
2. Fiacco, Anthony V. "Introduction to sensitivity and stability analysis in non linear programming." (1983).
3. Sobol, Ilya M. "Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates." Mathematics and computers in simulation 55.1-3 (2001): 271-280.
4. Saltelli, Andrea, et al. Global sensitivity analysis: the primer. John Wiley & Sons, 2008.
5. Shapley, Lloyd S. "A value for n-person games." (1953): 307-317.

6. Bartholdi III, John J., and Eda Kemahlioğlu-Ziya. "Using Shapley value to allocate savings in a supply chain." Supply chain optimization. Boston, MA: Springer US, 2005. 169-208.
7. Kemahlıoğlu-Ziya, Eda, and John J. Bartholdi III. "Centralizing inventory in supply chains by using Shapley value to allocate the profits." Manufacturing & Service Operations Management 13.2 (2011): 146-162.
8. Fiestras-Janeiro, M. Gloria, et al. "Cooperative game theory and inventory management." European Journal of Operational Research 210.3 (2011): 459-466.
9. Van den Heuvel, Wilco, Peter Borm, and Herbert Hamers. "Economic lot-sizing games." European Journal of Operational Research 176.2 (2007): 1117-1130.
10. Saltelli, Andrea, et al. "Variance based sensitivity analysis of model output. Design and estimator for the total sensitivity index." Computer physics communications 181.2 (2010): 259-270.
11. Lundberg, Scott M., and Su-In Lee. "A unified approach to interpreting model predictions." Advances in neural information processing systems 30 (2017).
12. Bertsimas, Dimitris, and John N. Tsitsiklis. Introduction to linear optimization. Vol. 6. Belmont, MA: Athena scientific, 1997.
13. Axsäter, Sven. Inventory control. Boston, MA: Springer US, 2006.
14. Graves, Stephen C. "A single-item inventory model for a nonstationary demand process." Manufacturing & Service Operations Management 1.1 (1999): 50-61.
15. Geevers, Kevin, Lotte Van Hezewijk, and Martijn RK Mes. "Multi-echelon inventory optimization using deep reinforcement learning." Central European Journal of Operations Research 32.3 (2024): 653-683.
16. Censor, Yair, et al. "On the effectiveness of projection methods for convex feasibility problems with linear inequality constraints." Computational Optimization and Applications 51.3 (2012): 1065-1088.
17. Barratt, Shane, Guillermo Angeris, and Stephen Boyd. "Automatic repair of convex optimization problems." Optimization and Engineering 22.1 (2021): 247-259.
18. Harsha, Pavithra, et al. "Deep policy iteration with integer programming for inventory management." Manufacturing & Service Operations Management 27.2 (2025): 369-388.