

Article

Closed-Form Equations for the Reorder Point and Order-Up-To Level in a Lost-Sales Periodic-Review (R, s, S) Inventory Policy

Samir Žic ^{1,*} and Jasmina Žic ^{2,*}¹ Faculty of Engineering, University of Rijeka, Vukovarska 58, 51000 Rijeka, Croatia² Faculty of Mechanical Engineering and Naval Architecture, University of Zagreb, Ivana Lucića 5, 10002 Zagreb, Croatia

* Correspondence: samir.zic@uniri.hr (S.Ž.); jasmina.zic@gmail.com (J.Ž.)

Abstract

This paper develops explicit equations for setting the reorder point s and the order-up-to level S in a periodic-review (R, s, S) inventory policy in a lost-sales environment. The objective is to support direct policy parameterization from demand level, demand variability, review period, lead time, and type-II unit fill-rate target. A long-horizon discrete-event simulation was combined with exhaustive enumeration of integer policy pairs to construct a policy-consistent reference dataset of five million observations. Symbolic regression was then used to convert this simulation-derived reference map into compact closed-form equations for both policy parameters. Over the full tested domain, the equations achieved $R^2 = 0.941$ for the reorder point s and $R^2 = 0.989$ for the order-up-to level S . On the common domain where analytical comparison is possible, the proposed equations reduced mean absolute error by approximately 65% for the reorder point and 85% for the order-up-to level. The equations also remain directly evaluable at the finite-horizon zero-lost-sales boundary corresponding to a 100% fill rate, where standard normal-loss logic has no finite safety-factor solution. The study provides an interpretable, auditable equation system for initial estimation of policy parameters for periodic-review lost-sales inventory policies within the tested normal-demand domain.

Keywords: periodic-review (R, s, S) inventory policy; lost sales; reorder point s ; order-up-to level S ; type-II unit fill rate; closed-form equations; exhaustive enumeration; simulation-derived reference map; symbolic regression; inventory policy parameterization

MSC: 90B05; 90B30; 90C59; 68T20

Academic Editor: Andrea Scozzari

Received: 8 June 2026

Revised: 29 June 2026

Accepted: 2 July 2026

Published: 6 July 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Inventory control links stochastic customer demand with finite replenishment and storage capacity. In periodic-review single-echelon systems, replenishment policies include base-stock (R, S) policies and reorder-point/order-up-to policies such as (R, s, S). The (R, s, S) inventory policy is particularly relevant in practice because it combines periodic decision-making with an explicit reorder level s and an order-up-to level S , making it compatible with many enterprise resource planning (ERP) and replenishment-planning environments. However, analytical design of (R, s, S) policies remains substantially more

difficult than the design of one-parameter periodic-review base-stock policies, especially when service-level constraints are imposed under lost-sales dynamics [1,2].

In retail, fast-moving consumer goods, e-commerce, and related environments, service performance is frequently measured by the type-II unit fill rate, defined as the fraction of demand satisfied immediately from on-hand stock. This measure is operationally meaningful because it reflects product availability at the moment of customer demand, which is often the relevant service criterion in retail and distribution systems [2,3]. The same immediate-fill concept can be used in both backorder and lost-sales settings, but the treatment of unmet demand differs: in backorder systems, unmet units remain as future obligations, whereas in lost-sales systems, they are permanently removed. Lost sales are especially relevant in consumer-facing sectors, where stockouts often lead to brand switching, store switching, purchase cancellation, or substitution rather than delayed fulfillment [4,5]. In the present study, the service measure is therefore interpreted strictly within a lost-sales state transition: unmet product units are permanently removed from the system and are not counted as later fulfilled demand. This convention keeps the simulation model, the exhaustive search, and the symbolic regression equations aligned with the lost-sales (R, s, S) setting.

The analytical literature contains important results on fill rates for periodic-review inventory systems, but most of these results concern policies or design problems that are simpler or structurally different from the (R, s, S) lost-sales setting considered here. Exact and approximate fill-rate expressions and related formal refinements have been developed for periodic-review base-stock (R, S) systems, including normally distributed demand and discrete-demand formulations [6–9]. These studies show that even for base-stock systems with a single policy parameter S , fill-rate evaluation can require careful treatment of demand discreteness, review timing, and service definition. However, base-stock formulas do not determine a reorder point s , and therefore do not solve the full (R, s, S) design problem.

The closest formal analytical predecessor to the present study is the Tijms–Groenevelt approximation for (s, S) -type service-level systems. Their paper explicitly focuses on the service measure requiring that a specified fraction of demand is met directly from stock on hand and presents approximations for the reorder point s [1]. However, their formulation requires the order-up-to gap $S - s$ as an external input, for example, from an economic order quantity (EOQ)-type calculation, and therefore solves the conditional problem $(\mu, \sigma, R, L, FR, S - s) \rightarrow s$ rather than the direct design problem $(\mu, \sigma, R, L, FR^*) \rightarrow (s, S)$. Although the method includes a lost-sales modification, it is not a full-spectrum lost-sales design equation for $FR = 1\%$ to 100% .

Other periodic-review and lost-sales studies provide important but non-equivalent reference points. Some methods address periodic-review lost-sales systems with lot sizing, case-pack restrictions, safety-stock approximations, or aggregate service constraints rather than a direct (R, s, S) equation pair [10,11]. Other lost-sales studies analyze structural properties, asymptotic optimality, robustness, or broader lost-sales theory rather than deriving explicit service-level design equations for both s and S in the periodic-review (R, s, S) setting [12–15].

Consequently, the gap addressed in this paper is specific. To the authors' knowledge, no existing method provides a direct, explicit equation pair that maps demand descriptors, review period, lead time, and type-II unit fill rate to both s and S for a periodic-review (R, s, S) lost-sales system over a broad operational domain. Existing methods either evaluate fill rate for a given policy, solve one-parameter base-stock special cases, require cost or $S - s$ inputs, impose lot-sizing or aggregate service constraints, or apply only in narrower service-level regimes. This gap is practically important because planners often need to set s and S for many items using only average demand, demand variability, review

period, lead time, and desired service level. Without a direct equation pair, practitioners must rely on repeated simulation, numerical search, adapted heuristics, or manual tuning.

Exhaustive simulation and symbolic regression are used here because existing analytical tools do not directly supply both policy thresholds for the present lost-sales periodic-review (R, s, S) setup. Periodic-review base-stock formulations address the one-threshold (R, S) policy class; they are analytically important, but they do not solve the two-threshold (R, s, S) setup problem because no separate reorder point s is determined. Conditional (s, S) -type approximations require an externally specified order-up-to gap $S - s$. EOQ-type or other cost-based rules introduce holding and ordering cost assumptions, whereas the present study is formulated as a service-level parameterization problem. Lot-sizing, case-pack, and aggregate-service formulations address related but structurally different settings. Consequently, the role of simulation in this study is to construct a policy-consistent reference map under the stated lost-sales dynamics, integer demand, review-period, and lead-time grid, and type-II unit fill-rate criterion, while symbolic regression is used to compress that reference map into explicit equations.

A further limitation of the normal-loss formulation is the $FR = 100\%$ boundary. Under an unbounded continuous normal approximation, exact zero expected shortage requires an infinite safety factor, so the corresponding normal-loss equation has no finite solution. In contrast, finite-horizon integer-demand simulations can identify finite (s, S) pairs that produce zero lost demand on the evaluated simulated demand paths. The present study, therefore, explicitly includes the $FR = 100\%$ zero-lost-sales boundary as part of the reference policy map and evaluates whether symbolic regression (SR) can approximate this boundary together with the remaining fill-rate spectrum.

The present study addresses this gap by constructing a large simulation-derived reference dataset for the periodic-review (R, s, S) lost-sales system and then deriving explicit equations for both s and S using symbolic regression. The system is evaluated under a stationary, normally distributed demand, implemented as a non-negative integer demand per period, with a deterministic review period R , a deterministic lead time L , lost sales, and at most one outstanding replenishment order.

The experimental domain covers average demand from 10 to 1000 units per period, coefficient-of-variation construction classes from 0.1 to 0.3, review periods $R = 1, \dots, 15$, lead times $L = 0, \dots, 15$, and fill-rate levels from 1% to 100%. Thus, the complete $R \times L$ grid is evaluated, rather than a single timing regime such as $R < L$, $R = L$, or $R > L$, because the review-period/lead-time relationship can materially affect periodic-review lost-sales dynamics, and the previous literature has noted that restrictive assumptions about the lead-time/review-period relationship are often not satisfied in practical settings [16]. Simulation-based analysis is appropriate for such inventory-control settings because it allows the actual behavior of complex replenishment rules to be evaluated when analytical assumptions or closed-form performance expressions are unavailable or incomplete [17]. The grid also includes $R = 1$ and $L = 0$, allowing the highest-frequency review case and the shortest replenishment-delay boundary of the tested periodic-review domain to be included.

For each demand scenario, demand replica, review period, lead time, and fill-rate class, an exhaustive enumeration of integer (s, S) policies is performed. Candidate policies are evaluated by discrete-event simulation, and the first feasible policy satisfying the service class is retained as the lowest feasible policy under the enumeration order. The exact achieved fill rate of each accepted policy is recorded and used as the service-level input for SR. The final dataset contains 5,002,469 accepted observations, including 50,400 observations at $FR = 100\%$. This creates a dense empirical characterization of the lowest feasible (s, S) policy mapping over the tested domain.

Symbolic regression is used because it searches for interpretable mathematical expressions rather than black-box predictive models. This is important in inventory-control

applications, where equations must often be implemented in spreadsheets, ERP systems, or planning tools and must remain auditable by practitioners. In this study, symbolic regression is implemented using PySR(version 1.5.10), a Python interface to Symbolic-Regression.jl that performs multi-objective search over mathematical expressions while balancing predictive error and expression complexity [18,19]. PySR is used to convert the exhaustive simulation benchmark into compact, auditable equations for s and S .

The resulting equations approximate the simulation-derived policy mapping with high aggregate accuracy over the tested domain. Across approximately five million accepted observations, the equation for S achieves a coefficient of determination (R^2) of 0.989 with a full-domain mean absolute error (MAE) of 187.2 units, while the equation for s achieves $R^2 = 0.941$ with a full-domain MAE of 266.03 units. These errors should be interpreted relative to the scale of the policy space: in the simulation benchmark, s ranges from 0 to 37,797 units, and S ranges from 1 to 51,948 units. Thus, the equations provide a compact, computationally efficient approximation to a policy mapping that would otherwise require extensive simulation search.

The contributions of this paper are sixfold. First, it develops a large-scale simulation and exhaustive-search framework for evaluating periodic-review (R, s, S) inventory policies in a lost-sales environment under a type-II unit fill-rate criterion, within the tested normal-demand domain. Second, it constructs, to the authors' knowledge, the first large-scale policy-consistent simulation-generated reference map for this setting, containing 5,002,469 policy-consistent accepted observations and simultaneously linking μ, σ, R, L , and FR to the corresponding policy parameters s and S . Third, it explicitly includes the FR = 100% zero-lost-sales boundary, for which the unbounded normal-loss fill-rate formulation has no finite safety-factor solution, and provides finite simulation-derived (s, S) policies for every operating configuration in the experimental grid. Fourth, it provides a detailed numerical evaluation of a hybrid normal-loss analytical comparator against this reference map on the FR < 100% domain, thereby quantifying the limits of directly transferring standard safety-stock logic to the lost-sales periodic-review (R, s, S) setting. Fifth, it derives explicit closed-form equations for both s and S , interpreted jointly as a coordinated (s, S) policy-pair parameterization of the simulation-derived map $(\mu, \sigma, R, L, \text{FR}) \rightarrow (s, S)$. This allows the algebraic structure of the policy-design map to be discovered from simulation reference data rather than imposed from classical safety-stock theory, while also covering the finite-horizon FR = 100% zero-lost-sales boundary. Sixth, it demonstrates under the tested conditions that SR can convert an exhaustive inventory simulation benchmark into interpretable inventory policy rules, reducing reliance on repeated simulation-based search during first-pass parameterization.

Together, these contributions position the study as both a theoretical inventory-control contribution and a practical policy-parameterization contribution: the paper develops a direct inverse mapping for a lost-sales periodic-review (R, s, S) design problem and translates demand descriptors and service targets into auditable values for the reorder point s and order-up-to level S within the tested domain.

2. Materials and Methods

2.1. Research Framework, Notation, and Assumptions

This study investigates the design of periodic-review (R, s, S) inventory policies for a single-item, single-echelon stocking system operating under lost sales. The objective is to derive explicit equations that map demand descriptors, review period, lead time, and type-II unit fill rate to the two policy parameters s and S . The modeling framework follows the lost-sales inventory literature, where unmet demand is removed from the system rather than carried forward as backlog [14], and is motivated by periodic-review lost-sales applications with target fill-rate requirements [10]. Throughout the Materials and Meth-

ods section, the main quantities are μ and σ for demand, R and L for deterministic timing, FR^* and FR for target and achieved fill rate, s and S for the policy thresholds, T for the simulation horizon, and r for the demand-replica index.

The system consists of one non-perishable item stored at one location. Capacity constraints, lateral transshipments, substitution, rationing, minimum order quantities, and explicit lot-size restrictions are excluded. The replenishment side is assumed to be uncappeditated and reliable: any order generated by the policy is delivered in full after the deterministic lead time. Customer demand that cannot be satisfied immediately from on-hand inventory is recorded as lost sales. The reference policy is defined as the lowest feasible (s, S) pair identified by the exhaustive enumeration procedure under the stated service criterion. The study is therefore service-level-oriented rather than cost-optimized; no explicit holding, ordering, or shortage-cost parameters are specified.

Inventory is controlled by a periodic-review (R, s, S) rule. Time is discretized into generic periods. The review period R and replenishment lead time L are deterministic non-negative integers measured in the same time unit as the demand observations. At review epochs, a replenishment order may be placed only if no previous order for the same product at the same stocking location is outstanding. If the system is eligible to order and the inventory level is at or below s , an order is issued to raise inventory to S ; otherwise, no order is placed. The restriction to at most one outstanding replenishment order is imposed to keep the lost-sales state representation tractable and is consistent with prior lost-sales inventory models that use a single outstanding order assumption, including periodic-review settings with lead-time information [20,21]. This restriction does not limit the number of replenishment orders that a warehouse or firm may place for other products; it only excludes overlapping open replenishment orders for the same product at the same stocking location. The restriction is especially relevant when the lead time exceeds the review period because several review epochs may occur before the previously issued order arrives. Allowing each of these reviews to generate an additional order for the same product at the same stocking location would create a different pipeline-ordering model, in which repeated orders can be issued before the effect of the first order is observed. The grid includes $R = 1$, the highest-frequency review case in the tested periodic-review domain, and $L = 0$, the shortest replenishment-delay boundary. Their exact timing interpretation is specified in Section 2.3.

Demand per period is assumed to be stationary and independent. At the latent level, demand follows a normal distribution, while the simulation uses non-negative integer-valued demand series generated and statistically validated as described in Section 2.4. This controlled stationary normal-demand setting is a standard framework in inventory analysis because it allows demand level and variability to be represented by μ and σ . Intermittent, seasonal, trend-driven, promotion-driven, or heavily skewed demand processes would require additional demand descriptors and a separately constructed simulation-derived reference map. The experimental demand domain is organized by mean demand μ and nominal coefficient-of-variation construction class CV^* . In the retained demand replicas, μ denotes the achieved mean demand per period. Because the generator imposes an exact mean constraint, this value is identical to the target mean level for all replicas in the same demand class. By contrast, σ denotes the achieved sample standard deviation of the specific demand replica; it may differ across the ten replicas within the same nominal CV^* class, although all accepted values remain within the prescribed tolerance. Thus, the symbolic regression input uses the realized demand tuple (μ, σ) , not only nominal demand-class labels.

Service performance is measured exclusively by the type-II unit fill rate. For a simulation run, let D_{total} denote total demand over the simulation horizon and let D_{lost} denote total demand not satisfied immediately from on-hand stock. The realized fill rate is:

$$FR = 1 - \frac{D_{\text{lost}}}{D_{\text{total}}}. \quad (1)$$

Thus, FR measures the fraction of demand units satisfied immediately from on-hand stock. This immediate-fill definition is standard in service-level inventory analysis [2]. In this study, unmet units are permanently lost and do not affect later replenishment decisions.

The exhaustive search is organized over target fill-rate classes FR^* , defined as one-percentage-point intervals $[1\%, 2\%), [2\%, 3\%), \dots, [98\%, 99\%), [99\%, 100\%)$, with $FR^* = 100\%$ treated as a separate zero-lost-sales boundary class. For each accepted policy, however, the exact achieved fill rate FR is recorded and retained in decimal form as the service-level input for SR. Thus, FR^* denotes the search class, whereas FR denotes the exact service level used as an explanatory variable in symbolic regression. Each recorded observation has the form $(\mu, \sigma, R, L, FR, s, S)$, where s and S are the lowest feasible policy parameters found by enumeration. The resulting dataset is used to fit symbolic regression models of the following form:

$$s = f_s(\mu, \sigma, R, L, FR), \quad S = f_S(\mu, \sigma, R, L, FR). \quad (2)$$

Simulation-based analysis is appropriate here because the behavior of lost-sales (R, s, S) systems is difficult to capture with general closed-form expressions, especially under integer demand, finite review periods, non-negative lead times, and service-level constraints [17]. Symbolic regression is used to convert the simulation-derived policy mapping into compact equations because it searches for explicit mathematical expressions that balance predictive accuracy and structural simplicity, rather than producing only black-box predictions [22,23]. This makes SR suitable for inventory policy parameterization, where the equations should be interpretable, auditable, and implementable in spreadsheets, ERP systems, or replenishment-planning software. Recent work has also applied SR to engineering and control problems [24], while earlier inventory-related work used simulation and SR to derive replenishment-planning equations for systems operating under an (R, s, S) policy [25]. In the present study, PySR is used to derive equations for the direct policy-parameter mapping from (μ, σ, R, L, FR) to (s, S) [18,19].

2.2. Analytical Reference Methods and Comparator Definition

The purpose of this subsection is to define how existing analytical inventory formulas are used in the present study. The reviewed literature does not provide a directly equivalent explicit equation pair of the form $(\mu, \sigma, R, L, FR) \rightarrow (s, S)$ for a periodic-review (R, s, S) lost-sales system under a type-II unit fill-rate criterion. Therefore, existing analytical methods are treated as reference formulations or illustrative comparators rather than as exact benchmarks for the proposed symbolic regression equations.

2.2.1. The Tijms–Groenevelt Approximation

The closest formal analytical predecessor is the approximation of Tijms and Groenevelt for (s, S) -type systems with service-level constraints [1]. Their service measure is the fraction of demand satisfied directly from stock on hand, which corresponds to the type-II unit fill rate measure denoted by FR in this paper. This method was also adopted in later textbook treatment as the decision rule for (R, s, S) systems with a specified fraction P_2 of demand satisfied directly from shelf [26]. However, the Tijms–Groenevelt approximation does not provide a direct equation pair for both (s, S) . Its formulation assumes that the order-up-to gap $Q = S - s$ is predetermined, for example by an EOQ-type calculation, and then derives an implicit approximation for the reorder point s [1]. In the notation

of the present study, the method therefore solves a conditional problem of the following form:

$$(\mu, \sigma, R, L, \text{FR}, Q) \rightarrow s, \quad Q = S - s$$

rather than the direct design problem considered here:

$$(\mu, \sigma, R, L, \text{FR}) \rightarrow (s, S)$$

This distinction is central because s and S jointly determine (R, s, S) behavior: s triggers replenishment, while S determines the post-replenishment inventory position and affects future stockout exposure, AIL, replenishment frequency, and shipment size. Consequently, changing either parameter changes the service behavior of the policy pair. The proposed symbolic regression equations, therefore, estimate both policy parameters directly from demand descriptors, operating conditions, and the exact achieved fill rate. By contrast, the Tijms–Groenevelt approximation requires $Q = S - s$ to be specified before s is computed. A numerical comparison with this method would therefore be conditional on an externally supplied Q ; if Q were taken from the simulation benchmark itself, the comparison would use information that the proposed equations are intended to predict [1].

Tijms and Groenevelt also discuss a lost-sales modification, in which the shortage fraction term in the backorder approximation is adjusted to account for lost sales. This makes the method highly relevant as an analytical reference. However, the distinction between the Tijms–Groenevelt approximation and the present study is structural rather than only numerical. The reviewed literature does not provide a directly equivalent closed-form equation pair for the inverse design problem $(\mu, \sigma, R, L, \text{FR}) \rightarrow (s, S)$ in the lost-sales periodic-review setting considered here. The Tijms–Groenevelt approximation remains a reorder-point approximation conditional on a previously specified Q . A numerical benchmark based on this method is therefore not unique unless an additional rule for Q is imposed. Using Q from the simulation-derived reference policy would be circular because it would supply the comparator with information that the proposed equations are intended to predict. Using an EOQ-type value of Q , or another cost-based order-quantity rule, would introduce holding-cost and ordering-cost assumptions that are outside the service-level-oriented formulation of the present study and would change the problem into a mixed service-cost optimization problem. In the present framework, costs are treated as a subsequent decision layer that can be evaluated after service-feasible policy parameters, inventory levels, order frequencies, and shipment quantities have been determined. For this reason, the Tijms–Groenevelt approximation is retained as the closest formal analytical predecessor, but not as a primary numerical benchmark for the proposed equations.

Other analytical methods are also non-equivalent to the present design problem. Periodic-review base-stock (R, S) fill-rate formulas evaluate or set a single order-up-to parameter and do not determine a reorder point s [6–9]. Periodic-review lost-sales studies with lot-sizing or aggregate-service constraints address related but different replenishment structures [10].

2.2.2. Hybrid Normal-Loss Heuristic Used as an Illustrative Comparator

Because no directly equivalent published equation pair of the form $(\mu, \sigma, R, L, \text{FR}) \rightarrow (s, S)$ was identified for the lost-sales periodic-review (R, s, S) policy, we define a simple hybrid normal-loss heuristic as an illustrative comparator. This heuristic is not claimed to be a formal analytical solution of the (R, s, S) problem. Rather, it combines two standard safety-stock components: lead-time reorder-point protection for s and periodic-review protection-period logic for S . The normal-loss and type-II unit fill rate relationship used

in this construction is standard in service-level inventory analysis [2], while related periodic-review fill-rate expressions have been studied for base-stock (R, S) systems [6–9].

Let $G(k)$ denote the following standard normal loss function:

$$G(k) = \int_k^{\infty} (z - k) \phi(z) dz = \phi(k) - k[1 - \Phi(k)], \quad (3)$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ are the standard normal density and distribution function. Unlike a cycle-service calculation, the type-II unit fill rate calculation cannot be obtained directly from the usual standard normal quantile table. A quantile $k = \Phi^{-1}(\alpha)$ controls the probability that demand does not exceed a threshold, whereas fill rate depends on expected shortage volume. Therefore, the safety factor k_H must be obtained by numerical inversion of the loss function. This distinction is important because FR in this study measures the fraction of demanded product units delivered immediately from stock, not the fraction of periods or cycles without a stockout [2,6–8].

For $FR < 100\%$, the heuristic safety factor k_H is defined as the solution of:

$$G(k_H) = \frac{(1 - FR)\mu R}{\sigma\sqrt{R + L}}. \quad (4)$$

Since $G(k)$ is strictly decreasing in k , the equation has a unique finite solution for every positive right-hand side. In this study, k_H is computed numerically for each observation with $FR < 100\%$ using a one-dimensional Newton–Raphson solver. The boundary case $FR = 100\%$ corresponds to a zero-shortage target and would require $k_H \rightarrow \infty$ under the unbounded normal approximation; therefore, the heuristic is undefined at $FR = 100\%$. This structural limitation is important for the present study because the simulation-derived reference dataset contains finite zero-lost-sales policies at $FR = 100\%$, whereas the hybrid normal-loss comparator can only be evaluated on the common domain $FR < 100\%$.

The heuristic reorder point and order-up-to level are then defined as:

$$\begin{aligned} s_H &= \lceil \mu L + k_H \sigma \sqrt{L} \rceil, \\ S_H &= \lceil \mu(R + L) + k_H \sigma \sqrt{R + L} \rceil. \end{aligned} \quad (5)$$

The first expression applies lead-time protection logic to the reorder point s , while the second applies periodic-review protection-period logic to the order-up-to level S . Their direct combination is introduced here only as a transparent practical comparator. It should not be interpreted as a published textbook rule for the full periodic-review (R, s, S) lost-sales design problem.

2.2.3. Comparison Strategy

The primary reference in this study is the exhaustive simulation-derived dataset constructed in Sections 2.3–2.5. The proposed symbolic regression equations are therefore evaluated primarily by comparing their predicted values, s_{SR} and S_{SR} , with the reference values s and S retained in the final policy-consistent dataset.

The hybrid normal-loss heuristic defined in Section 2.2.2 is used solely as a secondary, illustrative analytical comparator in the common domain $FR < 100\%$; primary validation of the proposed equations is performed against the simulation-derived reference values. Its purpose is to quantify how a transparent adaptation of standard normal-loss safety-stock logic performs when applied to the present lost-sales periodic-review (R, s, S) simulation-derived policy map. It is not treated as a formal published competitor because it is not a direct analytical solution to the inverse design problem $(\mu, \sigma, R, L, FR) \rightarrow (s, S)$ and is undefined at $FR = 100\%$.

Accordingly, Section 3.3 evaluates the hybrid normal-loss comparator only on the common $FR < 100\%$ domain, while Section 3.4 evaluates the symbolic regression equations

on both $FR < 100\%$ and the full $FR \leq 100\%$ domain, including the $FR = 100\%$ zero-lost-sales boundary. The Results section reports the common-domain comparison, while the Discussion interprets it as an illustrative analytical reference rather than a direct comparison of equivalent methods. Accordingly, the Tijms–Groenevelt approximation is retained as the closest formal analytical predecessor, while numerical validation is based on the simulation-derived reference dataset. This avoids treating conditional, one-parameter, or structurally non-equivalent formulas as direct competitors and keeps the evaluation focused on the study’s objective.

2.3. Discrete-Event Simulation of the Lost-Sales (R, s, S) System

The simulation model evaluates candidate (s, S) policies under the statistically validated demand replicas described in Section 2.4. The system is represented as a single-item, single-echelon, discrete-time lost-sales inventory model with a finite horizon of T periods. One period is the common unit for demand observation, inventory control, R , and L ; in the computational experiments, it can be interpreted as one day. Discrete-event simulation is used because lost-sales inventory systems with periodic reviews, lead times, integer-valued demand, and service-level constraints are difficult to represent by general closed-form performance expressions [17].

Let D_t denote customer demand in period t , where $t = 1, \dots, T$. Demand is non-negative, integer-valued, exogenous, and independent of inventory-control decisions. Let I_t^{start} denote on-hand inventory available at the start of the period t , after any replenishment scheduled for that period was received. Let I_t^{end} denote on-hand inventory after demand in the period t . Let Y_t denote the number of demand units satisfied immediately from on-hand stock in the period t , let $D_{\text{lost},t}$ denote lost demand in the period t and let Q_t denote the replenishment order issued at the end of period t , if any. At the start of each period, any scheduled replenishment order for that period becomes available and is added to the on-hand inventory. Demand is then realized. Sales, lost demand, and end-of-period inventory are calculated as:

$$\begin{aligned} Y_t &= \min(D_t, I_t^{\text{start}}), \\ D_{\text{lost},t} &= \max(D_t - I_t^{\text{start}}, 0), \\ I_t^{\text{end}} &= \max(I_t^{\text{start}} - D_t, 0). \end{aligned} \quad (6)$$

Thus, on-hand inventory is never negative. Any demand exceeding available stock is recorded as lost demand and not carried forward to later periods. This lost-sales treatment is standard in inventory models where unmet demand is removed from the system rather than accumulated as backlog [14].

The review period R and replenishment lead time L are deterministic non-negative integers measured in periods. Review epochs are defined as $t = \{R, 2R, 3R, \dots\}$. At the end of each review epoch, the (R, s, S) rule is applied. If no replenishment order is outstanding, the order quantity is defined as:

$$Q_t = \begin{cases} S - I_t^{\text{end}}, & \text{if } I_t^{\text{end}} \leq s, \\ 0, & \text{if } I_t^{\text{end}} > s. \end{cases} \quad (7)$$

If a review epoch occurs while an earlier order is still in transit, no additional order is placed. Outside review epochs, $Q_t = 0$. An order issued at the end of the period t becomes available at the start of the period $t + L + 1$. Thus, $L = 0$ denotes zero full intervening periods between order placement and replenishment availability: the order is not available during the already completed period t , but it is available at the start of the next period. For $L > 0$, the order becomes available after L complete intervening periods. This timing convention is consistent with the simulation’s discrete-time structure, in which re-

view and ordering occur after demand within each review period. Continuous-review proximity is represented by $R = 1$, the highest-frequency review case in the periodic-review grid.

The restriction that at most one replenishment order may be outstanding is imposed throughout the simulation. If a review epoch occurs while an earlier order is still in transit, no additional order is placed. This single-open-order restriction is common in lost-sales inventory models because it simplifies the state representation while preserving relevant replenishment dynamics [20,21].

Each candidate policy (s, S) is simulated separately on a fixed demand replica. The initial state is policy-consistent $I_1^{\text{start}} = S$, with no outstanding replenishment order. Thus, every candidate policy starts from a just-replenished state corresponding to its own order-up-to level. Since S differs across candidate policies, the initial inventory is policy-specific rather than fixed externally. For each simulation run, total demand and total lost demand are recorded as:

$$\begin{aligned} D_{\text{total}} &= \sum_{t=1}^T D_t, \\ D_{\text{lost}} &= \sum_{t=1}^T D_{\text{lost},t}. \end{aligned} \quad (8)$$

The realized fill rate is then computed using Equation (1). For every evaluated candidate policy, the simulator returns the tuple $(\mu, \sigma, R, L, s, S, FR)$. This realized finite-horizon FR is used in Section 2.5 to determine whether the candidate policy satisfies the target fill-rate interval FR^* . For accepted policies, the exact realized FR is retained as the service-level input for SR.

2.4. Demand Input Construction and Statistical Validation

The simulation experiments (SEs) use fixed, finite-demand replicas so that all candidate (s, S) policies within the same demand scenario are evaluated on the same realized demand path. This design removes resampling noise from policy comparisons and supports reproducible construction of the $(\mu, \sigma, R, L, FR, s, S)$ dataset. Simulation-generated demand inputs are commonly used in inventory-control experiments [17].

Demand is represented as a sequence of non-negative integer observations, $D_1, D_2, \dots, D_T$, with $T = 3650$ periods. The experimental demand grid contains seven target mean levels, $\mu \in \{10, 25, 50, 100, 250, 500, 1000\}$. For each mean level, three nominal coefficient of variation (CV) classes were constructed $CV^* \in \{0.1, 0.2, 0.3\}$ with a tolerance of ± 0.01 . These classes define the relative demand-variability levels used during demand generation; subsequent inventory simulation and SR use the achieved demand descriptors μ and σ .

Candidate replicas were generated from a latent normal model. For each scenario, a latent coefficient of variation was sampled uniformly from the scenario-specific interval $[CV_{\min}, CV_{\max}]$. The corresponding latent standard deviation was set to μCV , and T continuous observations were sampled independently from $N(\mu, (\mu CV)^2)$. Continuous values were rounded to integers, and candidate sequences containing negative rounded values were rejected rather than truncated. A candidate sequence was retained only if its total demand exactly matched the target mean condition:

$$\sum_{t=1}^T D_t = \mu T,$$

Equivalently:

$$\frac{1}{T} \sum_{t=1}^T D_t = \mu$$

and if its achieved demand-variability satisfied the prescribed CV-interval condition. The achieved standard deviation used in this study is the population standard deviation:

$$\sigma = \sqrt{\frac{1}{T} \sum_{t=1}^T (D_t - \mu)^2} \quad (9)$$

This definition is consistent with the demand-generation procedure used to accept or reject candidate replicas.

Candidate sequences satisfying the moment constraints were then screened using statistical diagnostics for normality and temporal randomness. Normality was evaluated using the D'Agostino–Pearson, Shapiro–Wilk, and Anderson–Darling tests. Temporal randomness and serial dependence were evaluated using the runs test, the Ljung–Box test at lag 20, and the maximum absolute sample autocorrelation over lags 1–30. These tests and diagnostics were used to assess whether the retained integer-valued demand replicas were sufficiently consistent with the intended normal-like and independent demand structure [27–32]. The purpose of this screening was not to claim that all real market demand is normally distributed, but to construct controlled integer-valued demand inputs that are as consistent as feasible with the normal-demand scope of the experiment.

For each (μ, CV^*) construction class, ten independent demand replicas were retained, producing $7 \times 3 \times 10 = 210$ demand inputs. All retained replicas in the same mean class have identical achieved μ , while their achieved σ values differ slightly within the accepted demand-variability interval. The resulting demand replicas are used as fixed inputs in the exhaustive policy enumeration described in Section 2.5.

2.5. Exhaustive Policy Enumeration and Dataset Construction

The exhaustive enumeration procedure constructs the reference dataset used for symbolic regression. For each validated demand replica from Section 2.4, the simulation model from Section 2.3 is evaluated over the review-period grid $R \in \{1, 2, \dots, 15\}$, the lead-time grid $L \in \{0, 1, \dots, 15\}$, and 100 target fill-rate classes FR^* . The target classes consist of one-percentage-point intervals $[1\%, 2\%), [2\%, 3\%), \dots, [98\%, 99\%), [99\%, 100\%)$ with $FR^* = 100\%$ treated as a separate zero-lost-sales boundary class. The resulting 240 timing combinations include 120 cases with $R > L$, 15 cases with $R = L$, and 105 cases with $R < L$, including the 15 zero-lead-time cases with $L = 0$. Therefore, the reference dataset does not condition SR on a single review period/lead-time regime. The complete experimental design contains $7 \times 3 \times 10 \times 15 \times 16 \times 100 = 5,040,000$ potential design points, corresponding to mean-demand levels, demand-variability classes, demand replicas, review periods, lead times, and fill-rate classes. The use of exhaustive simulation is appropriate because no directly equivalent closed-form equation pair is available for determining both s and S in the lost-sales (R, s, S) setting considered here, while simulation-based inventory analysis is well-suited for systems whose performance cannot be represented adequately by simple analytical expressions [17,33].

The simulation and exhaustive enumeration were implemented in OptimInventory version 5.1, an inventory-simulation software environment previously applied in simulation-based studies of periodic-review (R, s, S) inventory systems, including analyses of inventory levels, costs, emissions, and replenishment-planning outcomes [25,34–36].

For each fixed $(\mu, \sigma, R, L, \text{replica})$ configuration, integer policy pairs are enumerated by increasing S ; for each fixed S , admissible values of s are then evaluated in increasing

order. The search begins with the smallest admissible order-up-to level and expands the feasible policy space one inventory unit at a time. For each value of S , all admissible reorder points below S are evaluated before S is increased. The admissible policy space is:

$$S \in \mathbb{Z}_{\geq 1}, s \in \{0, 1, \dots, S - 1\}. \quad (10)$$

Each candidate pair (s, S) is evaluated using a lost-sales simulation over $T = 3650$ periods, and its exact realized fill rate FR is calculated. The target service classes are one-percentage-point intervals over $1\% \leq FR < 100\%$, with a separate boundary class for $FR = 100\%$. A simulated policy is assigned to the class containing its realized fill rate. The 100% class is assigned only when no lost demand occurs over the full simulation horizon, $D_{\text{lost}} = 0$.

For each $(\mu, \sigma, R, L, \text{replica})$ configuration and each fill-rate class FR^* , only the first policy encountered in the enumeration order is retained. The retained pair is therefore the lowest- S , then the lowest- s , feasible policy found for that service class. This construction deliberately follows a lowest-threshold feasible-policy principle: it identifies the earliest policy that reaches the required service class while keeping the structural inventory thresholds as low as possible. This construction is economically meaningful because the enumeration rule prioritizes lower structural inventory thresholds and, as the auxiliary average inventory level (AIL) check below indicates, generally identifies policies with low or near-low average inventory levels among service-equivalent alternatives. A full cost-based comparison, including item values, holding costs, ordering costs, and shortage penalties, is left for future research.

Whenever a policy is retained, the recorded observation is $(\mu, \sigma, R, L, FR, s, S)$, where FR is the exact achieved fill rate of the accepted simulation run. The class label FR^* is used only to organize the enumeration; it is not used as the service-level input for symbolic regression. Consequently, the final dataset contains the exact achieved fill-rate values rather than only 100 categorical service labels.

The $FR = 100\%$ class is available for every operating configuration. By contrast, some lower classes may be absent because they are not attainable by any integer (s, S) policy on the finite demand path. This occurs when the discrete mapping $(s, S) \rightarrow FR$ jumps over a narrow target interval. The effect is most visible at low mean demand levels and shorter review-period/lead-time combinations, where total demand volume is smaller, and each lost or satisfied unit has a larger effect on realized fill rate. It becomes negligible at higher demand levels or longer review periods/lead times. Here, $FR = 100\%$ denotes zero lost demand over the finite simulated demand path, not zero shortage probability under an unbounded continuous demand distribution. Therefore, finite (s, S) policies can exist in the simulation even though the normal-loss comparator has no finite safety-factor solution at $FR = 100\%$.

Across the full experimental design, the exhaustive search produced 5,009,739 accepted simulation records out of 5,040,000 possible design points, corresponding to approximately 99.4% coverage, and required evaluating approximately 2.82×10^{12} candidate (s, S) policies. After applying the (R, s, S) policy-consistency filter described below, 7270 accepted observations were excluded, leaving 5,002,469 observations for comparator evaluation and symbolic regression. The $FR = 100\%$ class remained complete after filtering, with 50,400 zero-lost-sales observations, one for each operating configuration.

The resulting dataset is a large-scale simulation-derived reference map of the lowest threshold feasible (s, S) policies under lost-sales dynamics and type-II unit-fill-rate classification. This distinction is important because the reviewed literature does not provide an equivalent equation pair for recommending both s and S under this setting. The dataset, therefore, serves as the empirical foundation for evaluating the illustrative comparator in

Section 2.2 and for training and validating the symbolic regression equations in Section 2.6.

After the exhaustive search, the raw first-feasible simulation experiments were screened using an (R, s, S) inventory policy-consistency filter. The purpose of this step was not to correct the simulation search but to construct a coherent reference policy map for symbolic regression. The exhaustive search identifies the first-feasible (s, S) pair separately for each target fill-rate class FR^* ; it does not impose a global ordering of policy parameters across target service classes. Therefore, under finite-horizon integer demand, local irregularities may occur in which a lower FR^* class is associated with a higher order-up-to level S , or with the same S but a higher reorder point s , than a higher FR^* class for the same generated market-demand replica and the same (R, L) setting. The filter was applied independently within each sequence defined by the same demand replica and the same (R, L) . Starting from the highest retained service class, each lower- FR^* row was compared with the last retained higher- FR^* row. The lower- FR^* row was retained only if its S value was lower, or if S was equal and its s value was lower or equal. Thus, for two consecutive retained target fill-rate classes with $FR_{k+1}^* < FR_k^*$, the cleaned reference dataset satisfies:

$$S_{k+1} < S_k, \text{ or } S_{k+1} = S_k \text{ and } s_{k+1} \leq s_k. \quad (11)$$

Rows excluded by this rule were not treated as simulation errors. They were valid finite-simulation results, but were not used in the final reference dataset because they violated the intended (R, s, S) policy envelope when the target service was relaxed. The resulting policy-consistent dataset serves as the final reference dataset for SR in Section 2.6.

Because the enumeration proceeds from the lowest admissible inventory thresholds upward, the first service-feasible (s, S) pair is used as the lowest-inventory feasible policy under the adopted search construction. As an auxiliary check, the AIL behavior of this first-retained policy was evaluated on a high-service subset defined by $\mu = 50$, $CV^* = 0.1$, 0.2 , 0.3 , $R = 1, 2, \dots, 10$, $L = 0, 1, \dots, 10$, and $FR^* = 90\%, 91\%, \dots, 100\%$. For each evaluated scenario group, the first retained policy was compared with 99 later service-feasible alternatives, giving 2209 groups and 220,900 evaluated candidate-policy records. In 217,215 records (98.33%), later retained policies did not produce a lower AIL than the corresponding first retained policy; 124,110 records (56.18%) produced a higher AIL, with an average relative increase of 143.22%, while only 3685 records (1.67%) produced a lower AIL, with an average relative reduction of 2.11%. These auxiliary results support using the first service-feasible (s, S) pair as the low-inventory reference policy in the present equation-development dataset, while a full-domain analysis of AIL behavior across later feasible alternatives remains a topic for future research.

2.6. Symbolic Regression Protocol

Symbolic regression was used to convert the simulation-derived policy mapping from Section 2.5 into explicit equations for the two (R, s, S) policy parameters. Each accepted observation has the form $(\mu, \sigma, R, L, FR, s, S)$, where FR is the exact fill rate achieved in simulation. The variables μ and σ are numerical demand-rate descriptors, R and L are integer time parameters, and FR is stored as a decimal fraction. In the constructed dataset, the retained service range covers approximately $FR \in [0.01, 1]$, corresponding to fill rates from 1% to 100%.

Two independent symbolic regression searches were conducted, one for the reorder point s and one for the order-up-to level S , consistent with Equation (2). The resulting closed-form equations are interpreted jointly as a coordinated (s, S) policy-pair parameterization of the map $(\mu, \sigma, R, L, FR) \rightarrow (s, S)$.

Symbolic regression was implemented in Python using PySR version 1.5.10. PySR is a symbolic regression framework based on SymbolicRegression.jl that performs evolu-

tionary search over mathematical expressions and returns a Pareto front of candidate equations balancing predictive accuracy and expression complexity [18,19]. Generally, symbolic regression is used for interpretable scientific model discovery [22,23]. Recent work has also used SR as supervised machine learning for control synthesis, emphasizing its ability to discover both expression structure and parameters in engineering decision problems [24]. The simulation-and-symbolic regression methodology also follows the authors' earlier work on (R, s, S) -based replenishment planning, while the present study addresses a different inverse design problem [25]. Related work has also used genetic programming for inventory-control policy derivation, supporting the broader relevance of evolutionary symbolic methods in inventory-policy design [37–39].

The symbolic search space was restricted to addition, subtraction, multiplication, division, and the square-root operator. This restriction was imposed to keep the resulting equations compact, auditable, and implementable, but not to force the discovered formulas to reproduce the algebraic structure of classical inventory approximations. The purpose of symbolic regression in this study is equation discovery rather than fitting a pre-specified safety-stock form. Candidate expressions were therefore allowed to contain non-classical combinations of demand, review period, lead time, and fill-rate variables, provided they satisfied dimensional constraints and achieved acceptable validation accuracy. Numerical constants were permitted, but final equations were selected with preference for small integer constants when this did not significantly reduce predictive accuracy. This reflects the practical objective of producing equations that can be implemented easily in spreadsheets, ERP systems, or replenishment-planning tools.

Dimensional consistency was considered a necessary screening criterion. The variables μ and σ are empirical demand descriptors measured in the same demand unit per period; R and L are time quantities measured in periods; FR is dimensionless; and the outputs s and S are inventory quantities. Candidate expressions were therefore required to preserve the correct output dimension. Expressions that numerically fit the data but violated dimensional homogeneity were not considered admissible policy equations. Once dimensional validity, operational feasibility, validation accuracy, and acceptable complexity were satisfied, non-classical algebraic structure was treated as a potentially informative outcome of the discovery process. This use of dimensional information is consistent with symbolic regression methodology in which physical units and dimensional constraints are used to restrict the search space and improve equation discovery [40].

Before symbolic regression, the dataset was randomly shuffled and divided into training and test subsets. After shuffling, 25% of the dataset, approximately 1.25 million observations, was assigned to the training subset, while the remaining 75%, approximately 3.77 million observations, was reserved for out-of-sample testing. The same split logic was applied separately to the s - and S -equation searches.

The primary search objective was the following mean absolute error:

$$\text{MAE} = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i| \quad (12)$$

where y_i is the simulation-derived value of s or S , and $\hat{y}_i$ is the corresponding equation prediction. MAE was used because it is expressed in inventory units and is directly interpretable for practitioners [41]. Additional metrics, including R^2 , root mean square error (RMSE), and absolute-error diagnostics such as median AE, 95th percentile AE, and maximum AE, are reported in Section 3 to assess both average accuracy and extreme deviations [42,43]. Because R^2 can be sensitive to the variance structure of the evaluated dataset and may be misleading when groups with different response scales are pooled, it is interpreted together with absolute-error metrics rather than used as a sole accuracy criterion [44,45].

Final symbolic regression searches used the MAE objective defined above. The main searches were executed as full 64-core workstation runs using 192 populations, population size of 60, ncycles_per_iteration = 3000, maxsize = 60, parsimony = 0.0005, and PySR model_selection = score. The iteration limit was set sufficiently high so that the search duration was controlled by practical convergence rather than by reaching the nominal iteration cap. MAE-based searches ran for approximately 30 days per target variable. Searches were continued until no new Pareto-front expression improving the relevant accuracy–complexity trade-off was obtained for 48 h.

Because the hybrid normal-loss comparator is undefined at $FR = 100\%$, symbolic regression accuracy is reported separately on three domains: the common comparator domain $FR < 100\%$, the zero-lost-sales boundary $FR = 100\%$, and the full reference domain $FR \leq 100\%$. This separation allows the proposed equations to be evaluated both against the comparator-compatible subset and against the full domain for which the equations are intended.

The final equations were selected from the PySR Pareto front. Candidate equations were jointly assessed based on test-set MAE, R^2 , maximum absolute error, algebraic simplicity, and ease of implementation. This selection rule reflects the aim of deriving inventory equations that are accurate enough for policy design while remaining compact enough for practical use. The exhaustive enumeration search budget is reported by the mean demand level in Section 3.2.

3. Results

3.1. Statistical Quality of Simulated Market-Demand Series

The reference dataset was constructed from 210 validated market-demand series: seven mean-demand levels $\mu \in \{10, 25, 50, 100, 250, 500, 1000\}$, three nominal coefficient-of-variation classes $CV^* \in \{0.1, 0.2, 0.3\}$ each with tolerance of ± 0.01 , and ten independent replicas per (μ, CV^*) class. Each series contained 3650 non-negative integer demand observations. The achieved mean μ was equal to the target by construction, while the achieved standard deviation was replica-specific.

Classical inventory formulas depend not only on mean demand but also on demand-variability σ or $CV = \sigma/\mu$. The demand input was therefore constructed using nominal demand-variability classes, each with a tight tolerance around its target. For a fixed μ , this produces an approximately threefold increase in σ between the lowest and highest variability classes, allowing demand variability to be represented in the resulting (R, s, S) parameters. Table 1 reports the achieved CV values and the median minimum and maximum demand observed across the ten replicas in each demand class. The latter values are relevant because extreme realized demand within the 3650-period horizon is expected to affect the minimum feasible reorder point s and order-up-to level S , especially near high fill-rate levels. Although symbolic regression uses σ directly, Table 1 reports CV because the demand design was constructed from relative-demand-variability classes.

Table 1. Achieved demand variability and daily demand range of validated market-demand series.

Demand Scenarios (μ, CV^*)	Achieved CV (Median [Min–Max])	Daily Demand Range (Median of Min/Max)
(10, 0.1)	0.098 [0.094–0.109]	6.0–13.5
(10, 0.2)	0.197 [0.193–0.21]	3.0–17.0
(10, 0.3)	0.301 [0.291–0.308]	0.0–21.0
(25, 0.1)	0.103 [0.093–0.108]	16.0–34.0
(25, 0.2)	0.201 [0.193–0.209]	7.5–42.0
(25, 0.3)	0.299 [0.297–0.308]	0.5–51.0

(50, 0.1)	0.101 [0.094–0.107]	31.5–68.0
(50, 0.2)	0.201 [0.197–0.208]	14.5–85.0
(50, 0.3)	0.299 [0.296–0.306]	1.0–105.0
(100, 0.1)	0.108 [0.103–0.11]	59.0–138.5
(100, 0.2)	0.203 [0.19–0.209]	27.5–171.0
(100, 0.3)	0.298 [0.291–0.304]	3.0–209.0
(250, 0.1)	0.1 [0.092–0.109]	157.0–341.0
(250, 0.2)	0.202 [0.192–0.21]	69.0–427.0
(250, 0.3)	0.297 [0.291–0.307]	3.0–511.5
(500, 0.1)	0.104 [0.09–0.11]	321.5–683.5
(500, 0.2)	0.2 [0.193–0.207]	130.5–846.5
(500, 0.3)	0.292 [0.29–0.3]	9.5–1048.0
(1000, 0.1)	0.103 [0.092–0.11]	627.0–1364.5
(1000, 0.2)	0.205 [0.194–0.208]	309.0–1728.0
(1000, 0.3)	0.298 [0.29–0.306]	10.5–2075.0

Table 1 also shows the expected widening of the realized demand support as CV^* increases. For a fixed μ , moving from $CV^* = 0.1$ to $CV^* = 0.3$ shifts the lower end of the observed demand range toward zero and raises the upper end to approximately twice the mean demand in many scenarios. This is expected to be relevant for the (R, s, S) analysis because the lowest feasible s and S values are influenced not only by μ and σ but also by the realized finite-horizon demand path on which FR is evaluated.

The generated market-demand series were also evaluated for normality and temporal randomness, with results summarized in Table 2. This verification is important because the (R, s, S) inventory policy characteristic levels s and S depend not only on the distributional properties of demand, but also on its temporal ordering.

For example, a reordered demand sequence may preserve the same μ , σ , CV , and min–max values, while producing different inventory-depletion patterns, stockout timing, achieved FR, and, consequently, different first-feasible s and S . The retained replicas were therefore required to satisfy both distributional and temporal-randomness diagnostics, supporting their use as controlled normal-like and independent demand inputs.

Table 2. Summary of normality and temporal randomness diagnostics for the market-demand series grouped into 21 (μ, CV^*) classes.

Diagnostic	Reported Statistic	Result	Interpretation for the (R, s, S) Simulation Experiments
D’Agostino–Pearson	Range of class-level median p -values	0.926–0.988	All class-level medians are high. Replicas preserve the intended skewness–kurtosis structure of the latent normal model, supporting their use as normal-like demand inputs.
Shapiro–Wilk	Number of classes with median $p < 0.05$	8/21	Rejections occur in the lowest-demand classes, mainly $\mu = 10$, $\mu = 25$, and part of $\mu = 50$, where integer rounding makes discreteness most visible.
Anderson–Darling	Number of classes with median $p < 0.05$	10/21	Tail-sensitive results are stricter, again mainly in low- μ classes. This is consistent with a rounded non-negative integer demand.
Runs test	Range of class-level median p -values	0.420–0.960	No class-level median indicates systematic non-random ordering. This supports the use of the sequences for inventory simulations where demand order affects depletion and stockout timing.
Ljung–Box, lag 20	Range of class-level median p -values	0.230–0.769	Class-level medians do not indicate material autocorrelation up to lag 20, supporting the independence assumption over review and lead-time horizons.
Maximum absolute autocorrelation	Range of class-level median max. ACF , lags 1–30	0.032–0.048	Serial dependence is small across lags 1–30; this supports the use of the series where demand order affects inventory policy.

Table 2 shows that the market-demand series satisfies the intended stochastic structure sufficiently for the simulation experiments. The D’Agostino–Pearson medians are high in all demand classes, while the stricter Shapiro–Wilk and Anderson–Darling tests mainly detect the expected discreteness of low- μ , rounded non-negative demand. The temporal diagnostics are particularly relevant for the (R, s, S) policy because the order of demand observations affects inventory depletion, stockout timing, achieved FR, and, therefore, the first-feasible s and S . The runs and Ljung–Box results, together with median maximum absolute autocorrelations below 0.05, do not indicate material sequence effects in the generated demand inputs. These validated demand inputs form the basis for both the hybrid-comparator evaluation in Section 3.3 and the symbolic regression analysis in Section 3.4. The same realized μ , σ , and exact achieved FR values are used in both analyses.

3.2. Reference Dataset Coverage and Feasibility of the Exhaustive Search

The experimental design contained 5,040,000 theoretical simulation experiments, obtained from seven mean-demand levels μ , three nominal coefficient-of-variation classes CV^* , ten market-demand replicas, fifteen review periods R , sixteen lead times L , and one hundred target fill-rate classes FR^* . The exhaustive search identified 5,009,739 feasible simulation experiments. The remaining 30,261 theoretical simulation experiments were non-feasible, meaning that no (s, S) policy satisfied the simulation-experiment rule for the corresponding $(\mu, CV^*, \text{replica}, R, L, FR^*)$ combination.

After the exhaustive search, the (R, s, S) inventory policy-consistency filter described in Section 2.5 excluded 7270 feasible simulation experiments. The final reference dataset used for symbolic regression, therefore, contained 5,002,469 simulation experiments, corresponding to 99.26% of the full experimental design. The filter was highly selective, excluding only 0.145% of the feasible simulation experiments identified by the exhaustive search. The $FR^* = 100\%$ class remained complete, with one retained zero-lost-sales policy for every $(\mu, CV^*, \text{replica}, R, L)$ configuration.

Coverage by mean demand is summarized in Table 3. Each μ class contains 720,000 theoretical simulation experiments. Non-feasible simulation experiments were concentrated at the lowest mean-demand levels, while exhaustive-search feasibility became complete for $\mu \geq 250$. The policy-consistency filter was also localized: excluded rows occurred only for $\mu = 10, 25, 50$, with no exclusions for $\mu \geq 100$. Thus, the final reference dataset preserves essentially the entire feasible reference map while removing a small number of local policy-sequence irregularities before SR. In Table 3, “Tested policies” denotes the exhaustive-search budget, i.e., the number of candidate policy pairs evaluated during enumeration.

Table 3. Exhaustive-search coverage, exhaustive-search scale, and final reference-dataset size by mean demand (SE = simulation experiment).

μ	Tested Policies	Non-Feasible SEs	Excluded by Policy Filter	Final Reference SEs	Final Coverage (%)
10	2.127×10^8	24,423	6528	689,049	95.701
25	1.329×10^9	4848	656	714,496	99.236
50	5.320×10^9	907	86	719,007	99.862
100	2.117×10^{10}	83	0	719,917	99.988
250	1.316×10^{11}	0	0	720,000	100
500	5.281×10^{11}	0	0	720,000	100
1000	2.131×10^{12}	0	0	720,000	100

The two reductions from the theoretical design have different meanings. The 30,261 non-feasible simulation experiments correspond to target classes for which the exhaustive search did not identify a feasible (s, S) policy satisfying the simulation-experiment rule.

These cases occur mainly in low-demand, low-service, short- R , L settings, where integer policy increments and finite-horizon integer demand can make some one-percentage-point FR^* classes unattainable. The 7270 excluded simulation experiments, in contrast, were feasible results of the exhaustive search. They were excluded only because they did not satisfy the (R, s, S) policy-consistency envelope across decreasing FR^* classes. The final reference dataset is therefore the feasible and policy-consistent part of the simulation-derived policy map.

Table 4 summarizes coverage across the target fill-rate spectrum. Each 10-percentage-point FR^* range contains 504,000 possible simulation experiments and 16,800 operating cells of the form (μ, FR^*, R, L) . The table reports both retained simulation-experiment counts and empty operating cells because these two quantities describe different aspects of the learning domain: simulation-experiment coverage measures row-level density, while empty operating cells indicate where symbolic regression has no retained observation for a nominal operating-service condition.

Table 4. Final reference-dataset coverage by target fill-rate range.

FR* Range	Final Reference SEs	SE Coverage (%)	Empty Operating Cells	Operating-Cell Coverage (%)
1–10%	494,065	98.03	267	98.41
11–20%	496,319	98.48	172	98.98
21–30%	498,353	98.88	119	99.3
31–40%	499,610	99.13	80	99.52
41–50%	500,627	99.33	50	99.7
51–60%	501,384	99.48	31	99.82
61–70%	502,196	99.64	9	99.95
71–80%	502,854	99.77	0	100
81–90%	503,309	99.86	0	100
91–100%	503,752	99.95	0	100

Coverage increased across the target fill-rate spectrum from 98.03% in the $FR^* = 1$ –10% range to 99.95% in the $FR^* = 91$ –100% range. Empty operating cells were confined to lower service levels; within the 71–80% FR^* range upward, every operating cell had at least one retained simulation experiment. This confirms that the final reference dataset is particularly complete in the medium- and high-service regions, including the high-fill-rate range most relevant to practical service-level inventory planning. Importantly, the $FR^* = 100\%$ class remained complete after policy-consistency filtering. Thus, for every $(\mu, CV^*, \text{replica}, R, L)$ configuration, the final reference dataset contains a retained zero-lost-sales policy. This provides a complete upper service boundary for the symbolic regression learning domain.

Dataset completeness was also evaluated at the design-cell level. A CV^* -specific cell is one (μ, CV^*, FR^*, R, L) combination and can contain up to ten simulation experiments, one for each market-demand replica. An operating cell is one (μ, FR^*, R, L) combination aggregated across the three CV^* classes and can contain up to thirty simulation experiments. Table 5 summarizes the same coverage structure at the design-cell level, which is more directly relevant for symbolic regression than row-level counts alone.

Table 5. Cell-level coverage of the final policy-consistent reference dataset.

Cell Type	Total Cells	Full Cells	Partial Cells	Empty Cells	Cells with at Least One Retained SE
CV*-specific cells	504,000	499,524	1255	3221	500,779 = 99.36%
Operating cells	168,000	165,848	1424	728	167,272 = 99.57%

The 728 empty operating cells were highly localized. They occurred only for $\mu = 10$, 25, and 50; no empty operating cells occurred for $\mu \geq 100$. The highest FR^* with an empty operating cell was 66% for $\mu = 10$, 45% for $\mu = 25$, and 5% for $\mu = 50$. Therefore, all operating cells with $FR^* \geq 67\%$ had at least one retained simulation experiment. Empty operating cells were also absent for $R + L \geq 21$, where $R + L$ is used only as a compact descriptor of the combined review and lead-time setting. These results show that the final reference dataset is not a fragmented sample of the experimental design; it is dense over the domain used to derive the symbolic regression equations.

Table 6 reports the scale and local spread of the retained (s, S) values. The first two columns give the maximum retained values of s and S by mean demand. The next two columns report the maximum within-CV* replica spread. This spread is calculated within each fixed (μ, CV^*, FR^*, R, L) class across the ten market-demand replicas:

$$\Delta s = \max(s) - \min(s),$$

$$\Delta S = \max(S) - \min(S).$$

Because the ten replicas in each (μ, CV^*) class have the same nominal mean and coefficient of variation, and very similar achieved variability, this spread mainly reflects the effect of finite demand-path realization and temporal demand ordering. The last two columns report the broader local spread calculated within each fixed (μ, FR^*, R, L) class, pooling all CV* classes and replicas. This broader measure shows how much the retained (s, S) values can vary for the same μ , target service class, review period, and lead time when demand variability is allowed to differ.

Table 6. Final policy-parameter scale and local spread by mean demand.

μ	Max s	Max S	Max Δs Within CV*	Max ΔS Within CV*	Max Δs Across CV* and Replicas	Max ΔS Across CV* and Replicas
10	364	510	145	31	189	51
25	906	1273	384	76	472	129
50	1869	2550	779	133	990	260
100	3599	5104	1470	249	1836	461
250	8734	12,638	3727	572	4413	1111
500	17,978	25,307	7502	1200	9200	2146
1000	37,797	51,948	16,987	4138	20,533	5786

The policy-parameter scale increases strongly with mean demand. In the final reference dataset, the largest retained values were $s = 37,797$ and $S = 51,948$, both observed for $\mu = 1000$. The local spread of the retained reference map is also substantial. For $\mu = 1000$, the maximum within-CV* replica spread reached 16,987 for s and 4138 for S , while the broader spread across CV* classes and replicas reached 20,533 for s and 5786 for S . The largest spreads occurred in the high-service part of the reference map, especially near $FR^* = 100\%$ and at long R, L settings.

These ranges are important for interpreting the symbolic regression errors reported later. The reference dataset does not represent a single smooth deterministic curve; even under fixed operating and service conditions, finite demand-path realization and demand variability can generate substantial differences in the first-feasible (s, S) policy. Therefore, absolute prediction errors in s and S must be evaluated relative to both the global scale of the policy parameters and the local variability of the simulation-derived (R, s, S) reference policy map.

The final reference dataset therefore provides two complementary evaluation domains. The first is the common domain $FR < 100\%$, on which both the hybrid normal-loss comparator and the symbolic regression equations can be evaluated. The second is the full

domain $FR \leq 100\%$, including the complete $FR = 100\%$ zero-lost-sales boundary, on which only the symbolic regression equations are defined. Sections 3.3 and 3.4 use this distinction to keep the analytical comparator and the proposed equations separated.

3.3. Accuracy of the Hybrid Normal-Loss Comparator

The hybrid normal-loss comparator defined in Section 2.2.2 was evaluated against the simulation-derived reference dataset. The comparator was evaluated in its raw analytical form and only on the domain for which the normal-loss inversion has a finite solution. Because the continuous normal-loss formulation does not yield a finite safety factor for $FR = 100\%$, the 50,400 $FR = 100\%$ observations were excluded from the hybrid-comparator calculation, leaving 4,952,069 observations with $FR < 100\%$.

For each valid observation, the normal-loss safety factor k_H was obtained from Equation (4). The corresponding hybrid reorder point and order-up-to level were then calculated using Equation (5). The realized FR recorded for each individual reference-dataset entry was used in the calculation. No additional feasibility correction was imposed on the hybrid results. In particular, raw values were not truncated to enforce the definitional admissibility conditions $s \geq 0$, $S \geq 1$, and $s < S$; when the hybrid expression produced negative reorder-point values, they were retained and reported as part of the comparator's measured behavior.

Table 7 summarizes the hybrid-comparator domain and admissibility diagnostics. The hybrid comparator produced finite s_H and S_H values for 4,952,069 observations. Negative s_H values occurred in 451,559 rows, corresponding to 9.12% of the valid hybrid-comparison domain. No negative S_H values and no cases with $S_H \leq s_H$ were recorded. The negative s_H values occurred across all tested μ levels and all three CV levels. In the observed dataset, they were associated with $R \in [2, 15]$, $L \in [1, 9]$, and $FR \in [0.01, 0.73]$. The resulting negative reorder-point values ranged from -3093 to -1 .

Table 7. Raw-domain and admissibility diagnostics for the hybrid normal-loss comparator.

Diagnostic	Result
Total rows loaded	5,002,469
Valid hybrid rows with finite (s_H , S_H)	4,952,069
$FR = 100\%$ rows excluded from hybrid evaluation	50,400
Negative s_H rows	451,559
Share of valid hybrid domain with $s_H < 0$	9.12%
Negative S_H rows	0
Rows with $S_H \leq s_H$	0
Observed μ range for $s_H < 0$	10–1000
Observed CV range for $s_H < 0$	0.1–0.3
Observed R range for $s_H < 0$	2–15
Observed L range for $s_H < 0$	1–9
Observed FR range for $s_H < 0$	0.01–0.73
Observed s_H range for $s_H < 0$	−3093 to −1
Observed S_H range in rows with $s_H < 0$	11–11,951

The negative s_H values were not isolated row-level exceptions. They formed systematic regions in the tested parameter grid. The largest concentration occurred for short positive lead times and longer review periods. For $L = 1$, the hybrid comparator produced 160,344 negative s_H values, corresponding to 52.93% of valid hybrid rows with $L = 1$. The highest cell-level share occurred at $R = 15$, $L = 1$, where 15,180 rows, or 73.02% of valid rows in that cell, had $s_H < 0$. By service level, negative s_H values were concentrated at low

and moderate FR values: more than 90% of all negative s_H rows occurred at FR values below 50%.

Table 8 reports the aggregate accuracy of the hybrid comparator on the valid $FR < 100\%$ comparison domain. The comparator captures part of the broad policy scale but does not accurately reproduce the simulation-derived lowest feasible (s , S) reference map. The aggregate bias is positive for both policy parameters, indicating average overestimation on the full $FR < 100\%$ domain. The overestimation is stronger for S_H , with a bias of 757.31 units, compared with 252.26 units for s_H . The median absolute errors are considerably smaller than the 95th percentile absolute errors, indicating a right-skewed error distribution with substantial upper-tail deviations.

Table 8. Aggregate accuracy of the hybrid normal-loss comparator against the simulation-derived reference policies on the $FR < 100\%$ domain ($n = 4,952,069$).

Output	MAE	RMSE	Median AE	95th pct AE	Max AE	Bias	R ²
s_H vs. s	753.65	1713.55	155	3645	16,628	252.26	0.6
S_H vs. S	1207.14	2436.03	293	5841	15,039	757.31	0.705

The error pattern differs between the two policy parameters. The hybrid comparator produces larger aggregate errors for S_H than for s_H . This is relevant because S_H is calculated from a protection-period expression over $R + L$, whereas the simulation-derived S values arise from the full lost-sales (R , s , S) mechanism with integer policies, discrete demand, delayed replenishment, and at most one outstanding order. The positive aggregate bias for S_H indicates that the hybrid order-up-to expression tends to exceed the corresponding simulation-derived lowest feasible S values on the full $FR < 100\%$ domain.

Table 9 reports the same comparison by μ level. Absolute errors increase strongly with μ , as expected, because the magnitudes of s and S scale with demand volume. For example, at $\mu = 10$, MAE equals 27.51 for s_H and 44.83 for S_H . At $\mu = 1000$, the corresponding values increase to 2710.74 and 4333.48. The 95th percentile absolute error also increases substantially with μ , reaching 8979 units for s_H and 11,204 units for S_H at $\mu = 1000$.

Table 9. Accuracy of the hybrid normal-loss comparator by μ on the $FR < 100\%$ reference dataset.

μ	n	Output	MAE	RMSE	Median AE	95th pct AE	Max AE	Bias	R ²
10	681,849	s_H vs. s	27.51	40.04	17	92	160	13.63	0.326
		S_H vs. S	44.83	56.97	38	113	151	27.98	0.451
25	707,296	s_H vs. s	67.35	98.15	42	226	401	28.67	0.383
		S_H vs. S	109.29	140.11	90	281	377	68.62	0.469
50	711,807	s_H vs. s	134.55	196.23	83	452	853	51.95	0.41
		S_H vs. S	217.28	279.24	178	561	752	136.47	0.473
100	712,717	s_H vs. s	269.13	391.41	168	899	1587	96.53	0.426
		S_H vs. S	433.70	557.90	354	1121	1504	272.25	0.474
250	712,800	s_H vs. s	675.23	978.51	426	2241	3806	231.51	0.432
		S_H vs. S	1083.67	1394.32	885	2801	3760	679.95	0.474
500	712,800	s_H vs. s	1353.28	1961.05	853	4488	8373	448.75	0.439
		S_H vs. S	2167.34	2788.74	1771	5602	7515	1359.28	0.474
1000	712,800	s_H vs. s	2710.74	3923.41	1718	8979	16,628	882.37	0.44
		S_H vs. S	4333.48	5576.24	3540	11,204	15,039	2718.74	0.475

The R² values in Table 9 indicate that the hybrid comparator provides only limited-to-moderate explanatory accuracy within individual μ levels. For s_H , the within- μ R²

ranges from 0.326 to 0.44, while for S_H , it ranges from 0.451 to 0.475. These values are substantially lower than the aggregate R^2 values in Table 8, confirming that pooled R^2 partly reflects between- μ scale differences. Therefore, MAE, RMSE, median absolute error, and the 95th percentile absolute error are more informative because they measure deviations directly in inventory units.

The diagnostic domains used later for the symbolic regression equations provide additional insight into the hybrid comparator. In the high-service subset $0.90 \leq FR < 1$, the comparator has $MAE = 783.42$ for s_H and $MAE = 1075.99$ for S_H . Unlike the aggregate $FR < 100\%$ domain, the bias in this high-service subset is negative for both parameters, -780.70 for s_H and -938.59 for S_H . Thus, the hybrid comparator does not merely overestimate the reference map; its bias varies across service-level regimes.

The timing-related diagnostic subsets show additional structural differences. For $R = 1$, the comparator has $MAE = 1147.28$ and $R^2 = -0.041$ for s_H , while S_H has $MAE = 1143.48$ and $R^2 = 0.652$. For $L = 0$, the formula gives:

$$s_H = \lceil \mu \cdot 0 + k_H \sigma \sqrt{0} \rceil = 0$$

by construction. Consequently, the $L = 0$ reorder-point error reflects the extent to which the simulation-derived reorder point remains positive even when replenishment delay is zero. In this subset, s_H has $MAE = 63.71$ and a negative bias of the same magnitude.

By contrast, S_H is very close to the simulation-derived S values when $L = 0$, with $MAE = 1.4$ and $R^2 \approx 1.0$. This boundary-case result is informative because, when replenishment delay is zero, the hybrid order-up-to expression reduces to a periodic-review protection-period form over R :

$$S_H = \lceil \mu R + k_H \sigma \sqrt{R} \rceil.$$

Thus, the normal-loss protection-period logic remains highly informative for S in the tested zero-lead-time regime. In the combined $R = 1$, $L = 0$ subset, S_H is also very close to the simulation-derived S , with $MAE = 0.27$, while s_H remains structurally limited because it is zero for all rows.

Overall, the hybrid normal-loss comparator provides a transparent analytical benchmark, but its raw output differs substantially from the simulation-derived (R, s, S) reference map over the tested $FR < 100\%$ domain. Its limitations are evident in three ways: it is not defined for $FR = 100\%$, it produces negative s_H values in 9.12% of the valid hybrid-comparison rows, and it yields large aggregate and upper-tail errors for both s_H and S_H . These results provide the analytical context for the symbolic regression analysis in Section 3.4, where explicit equations are evaluated in both the common $FR < 100\%$ domain and the full $FR \leq 100\%$ reference domain.

3.4. Symbolic Regression Equations and Predictive Accuracy

The symbolic regression objective was full-domain MAE because it reports prediction error directly in inventory units. Final equation selection, however, also considered diagnostic subsets representing practically important or analytically difficult regions of the reference map: $FR < 100\%$, $FR = 100\%$, $R = 1$, $L = 0$, and the combined $R = 1$, $L = 0$ case. The $R = 1$, $L = 0$ and combined $R = 1$, $L = 0$ subsets are therefore boundary diagnostics within the tested domain, not extrapolation cases or separately optimized regimes. Across these subsets, MAE, RMSE, median AE, 95th percentile AE, maximum AE, Bias, and R^2 were considered jointly. On the evaluated Pareto front, increasing symbolic complexity did not yield simultaneous improvements across these metrics and subsets; rather, it shifted the error trade-off among average error, squared-error scale, upper-tail error, maximum error, explanatory fit, and interpretability. The final equations were therefore selected as compact, dimensionally admissible, non-negative, and operationally interpreta-

ble full-domain approximations, consistent with the MAE training objective but not chosen solely by the smallest attainable full-domain MAE.

The selected equation for the reorder point s is:

$$s_{SR} = \frac{11\mu LFR^4}{1 + 8FR^4} + \sigma RFR^4. \quad (13)$$

The selected equation for the order-up-to level S is:

$$S_{SR} = \mu FR \sqrt{R^2 + L^2(1 + 5FR - FR^3)}. \quad (14)$$

The closed-form equations also provide a direct analytical sensitivity interpretation. The reorder-point equation is monotone nondecreasing in μ , σ , R , L , and FR , while the order-up-to equation is monotone nondecreasing in μ , R , L , and FR . Thus, within the fitted equation pair, increases in the main demand, timing, and service-level inputs do not decrease the corresponding estimated policy thresholds.

The two selected equations can be interpreted through their main variable groups. In the s -equation, the first term, $11\mu LFR^4/(1 + 8FR^4)$, is driven by average demand and lead time, and therefore represents lead-time demand exposure adjusted by the target fill rate. The second term, σRFR^4 , represents a variability and review-period buffer whose contribution increases with the target fill rate. In the S -equation, μFR scales the order-up-to level by demand scale and service requirement, while the square-root term combines review-period exposure, R^2 , and lead-time exposure, $L^2(1 + 5FR - FR^3)$. The absence of σ from the selected S -equation should be interpreted as a result of Pareto-front model selection, not as a general theoretical claim that demand variability cannot affect S . In the selected equation pair, variability is represented through the reorder-point equation s , while the order-up-to level S is represented by a simpler demand-scale, timing, and service-level structure that retained favorable MAE and R^2 performance over the tested domain. Because the equations were obtained by symbolic regression, individual fitted coefficients should not be interpreted as independent causal mechanisms; the appropriate interpretation is at the level of the main components and variable groups.

In Equations (13) and (14), FR is expressed as a decimal fraction; therefore, $FR = 1$ corresponds to a fill rate of 100%. Since s and S are integer policy parameters, the continuous equation outputs were converted to integer policy parameters using the ceiling function before accuracy metrics were calculated. The selected equations produced no negative predicted values over the tested reference domain.

Operationally, an error in s changes the reorder trigger and therefore the exposure to stockout and lost-sales risk, whereas an error in S changes the replenishment ceiling and therefore the inventory exposure after ordering. As stated earlier in the manuscript, the present model is formulated as service-level-oriented policy parameterization and contains no holding, ordering, or shortage-cost parameters; accordingly, these errors are not converted into monetary terms.

Table 10 reports the predictive accuracy of the selected reorder-point s equation across the full reference domain and across the predefined diagnostic subsets. On the full $FR \leq 100\%$ domain, the equation achieved MAE = 266.03 inventory units, RMSE = 669.2, median AE = 41, and $R^2 = 0.941$ across 5,002,469 observations. This indicates strong aggregate agreement with the simulation-derived reorder points, especially considering that the retained reference values of s span from 0 to 37,797 inventory units. On the common comparator domain $FR < 100\%$, performance was very similar, with MAE = 262.17, RMSE = 654.2, median AE = 40, and $R^2 = 0.942$. Thus, the equation retains essentially the same accuracy on the domain where comparison with the hybrid normal-loss comparator is possible. The zero-lost-sales boundary $FR = 100\%$ produced MAE = 645.56, RMSE = 1548.43, median AE = 142, and maximum AE = 15,653. The maximum AE in this subset is

also the full-domain maximum AE, showing that the largest observed reorder-point deviation occurs at the finite zero-lost-sales boundary. Although these errors exceed those in the neighboring $90\% \leq \text{FR} < 100\%$ subset, they should be interpreted relative to the larger reorder-point scale and wider local policy ranges at the zero-lost-sales boundary. The $\text{FR} = 100\%$ subset is therefore best interpreted as a distinct and practically important part of the policy map, not simply as a proportional deterioration in equation performance. The high-service subset $90\% \leq \text{FR} < 100\%$ is operationally important because high fill-rate targets are common in supply-chain planning. Errors in this subset should be interpreted relative to the increased scale of s at high service levels. The selected full-domain equation, therefore, provides a compact general approximation, while specialized high-service equations remain a relevant direction for reducing upper-tail errors.

The timing-related diagnostic subsets show different behavior. For $R = 1$, which corresponds to the highest-frequency review case in the tested grid, the equation maintained strong accuracy, with $\text{MAE} = 228.63$, $\text{RMSE} = 630.16$, median AE = 32, and $R^2 = 0.94$. This indicates that the equation remains effective in the highest-frequency review case in the tested periodic-review grid. By contrast, the $L = 0$ and $R = 1$, $L = 0$ subsets had low absolute-error values but also much lower R^2 . For $L = 0$, MAE was 66.44, median AE was 6, and $R^2 = 0.609$. In the combined $R = 1$, $L = 0$ case, MAE was 59.18, median AE was 4, and $R^2 = 0.223$. This should not be interpreted as poor operational accuracy. In these subsets, many reference reorder-point values are zero or close to zero, thereby compressing the response variable's variance and making R^2 less informative. For such subsets, MAE, median AE, 95th percentile AE, and maximum AE provide a clearer interpretation of equation performance than R^2 .

The bias was small relative to MAE on the full and common comparator domains, with values of -19.54 and -13.88 , respectively. This indicates that the selected s -equation has only limited average directional error on the main evaluation domains, despite larger absolute deviations in the upper tail.

Overall, the selected reorder-point equation provides a strong full-domain approximation while preserving acceptable behavior across the diagnostic subsets. The main limitations are concentrated in the high-service and finite zero-lost-sales regions, where policy-parameter levels and local spreads are naturally larger. These regions should therefore be treated as priority targets for future scale-normalized analysis and specialized high-service equation development.

Table 10. Accuracy of the selected symbolic regression equation for reorder-point s .

Evaluation Domain	n	MAE	RMSE	Median AE	95th pct AE	Max AE	Bias	R^2
Full domain, $\text{FR} \leq 100\%$	5,002,469	266.03	669.2	41	1348	15,653	-19.54	0.941
Common comparator domain, $\text{FR} < 100\%$	4,952,069	262.17	654.2	40	1334	12,572	-13.88	0.942
Zero-lost-sales boundary, $\text{FR} = 100\%$	50,400	645.56	1548.43	142	2788.05	15,653	-576.26	0.915
High-service domain, $90\% \leq \text{FR} < 100\%$	503,679	470.41	1047.55	103	2252.10	12,572	-36.33	0.944
Highest-frequency review case, $R = 1$	327,948	228.63	630.16	32	1186	13,434	41.52	0.94
Zero lead-time case, $L = 0$	300,448	66.44	185.62	6	361	3728	23.55	0.609
Combined shortest timing case, $(R = 1, L = 0)$	16,281	59.18	144.61	4	335	1423	-57.27	0.223

Table 11 reports the corresponding accuracy of the selected order-up-to S equation. On the full $\text{FR} \leq 100\%$ domain, the equation achieved $\text{MAE} = 187.2$ inventory units, RMSE

= 487.31, median AE = 32, and $R^2 = 0.989$ across 5,002,469 observations. This indicates very strong agreement with the simulation-derived order-up-to levels, particularly given that the retained reference values of S range from 1 to 51,948 inventory units. On the common comparator domain $FR < 100\%$, performance was slightly stronger, with MAE = 181.3, RMSE = 463.53, median AE = 32, and $R^2 = 0.989$. Thus, the selected S -equation retains high accuracy on the domain where direct comparison with the hybrid normal-loss comparator is possible. The zero-lost-sales boundary $FR = 100\%$ was again more difficult than the ordinary comparator domain, with MAE = 765.98, RMSE = 1568.20, median AE = 205, and maximum AE = 15,205. The maximum AE again occurs at this boundary, confirming that the finite zero-lost-sales subset is the most difficult part of the policy map for both s and S .

The high-service subset $90\% \leq FR < 100\%$ produced higher errors than the broader $FR < 100\%$ domain, with MAE = 393.78, RMSE = 892.15, 95th percentile AE = 1868, and maximum AE = 9771. Nevertheless, $R^2 = 0.985$ indicates that the selected equation still captures the main structure of the order-up-to policy map in this practically important region. As with the reorder point, this suggests that the selected equation is suitable as a compact full-domain approximation, while specialized high-service equations may further improve upper-tail accuracy.

The timing-related diagnostic subsets show particularly strong performance for S . For $R = 1$, the selected equation achieved MAE = 87.41, RMSE = 208.92, median AE = 17, and $R^2 = 0.997$. For $L = 0$, MAE was only 10.96, median AE was 1, and $R^2 = 0.998$. In the combined $R = 1, L = 0$ case, MAE was 9.94, median AE was 0, and $R^2 = 0.963$. These results show that the selected S -equation remains highly accurate in the shortest timing regimes, including cases where the reorder-point equation has low R^2 because many s -values are zero or close to zero. The stronger behavior of the S -equation in these subsets is consistent with the smoother empirical behavior of S in the reference dataset.

Bias was also small relative to MAE across the main domains, with values of 30.08 on the full domain and 37.46 on the common-comparator domain. This contrasts with the much larger positive aggregate bias reported for the hybrid comparator in Section 3.3.

Overall, the selected order-up-to level S equation provides a very accurate full-domain approximation and strong diagnostic-subset performance. Its main limitations, like those of the reorder-point s equation, are concentrated in the high-service and zero-lost-sales regions, especially at $FR = 100\%$. This again supports future work on specialized symbolic regression equations for high-service operational ranges.

Table 11. Accuracy of the selected symbolic regression equation for order-up-to level S .

Evaluation Domain	n	MAE	RMSE	Median AE	95th pct AE	Max AE	Bias	R^2
Full domain, $FR \leq 100\%$	5,002,469	187.2	487.31	32	907	15,205	30.08	0.989
Common comparator domain, $FR < 100\%$	4,952,069	181.3	463.53	32	882	9771	37.46	0.989
Zero-lost-sales boundary, $FR = 100\%$	50,400	765.98	1568.20	205	3431	15,205	−694.71	0.967
High-service domain, $90\% \leq FR < 100\%$	503,679	393.78	892.15	86	1868	9771	−90.23	0.985
Highest-frequency review case, $R = 1$	327,948	87.41	208.92	17	413	3032	10.43	0.997
Zero lead-time case, $L = 0$	300,448	10.96	87.97	1	23	3620	−9	0.998
Combined shortest timing case, $R = 1, L = 0$	16,281	9.94	48.18	0	46	1193	−9.92	0.963

Table 12 compares the selected symbolic regression equations with the hybrid normal-loss comparator in the common domain $FR < 100\%$, as this is the only domain in which both methods are defined. The comparison shows substantially lower prediction errors for the symbolic regression equations for both policy parameters. For s , MAE decreases from 753.65 under the hybrid comparator to 262.17 under the symbolic regression equation, while the bias changes from 252.26 to -13.88 , and R^2 increases from 0.6 to 0.942. For S , MAE decreases from 1207.14 to 181.3, bias decreases from 757.31 to 37.46, and R^2 increases from 0.705 to 0.989. RMSE, median AE, 95th percentile AE, and maximum AE are also lower for both symbolic regression equations.

These results show that the selected symbolic regression equations provide both stronger common-domain accuracy and broader domain coverage: they substantially reduce errors on $FR < 100\%$ and remain evaluable on the full $FR \leq 100\%$ reference domain, including the $FR = 100\%$ boundary where the hybrid comparator is undefined.

Table 12. Common-domain comparison between symbolic regression equations and the hybrid normal-loss comparator on $FR < 100\%$ ($n = 4,952,069$).

Policy Parameter	Method	MAE	RMSE	Median AE	95th pct AE	Max AE	Bias	R^2
s	Hybrid comparator	753.65	1713.55	155	3645	16,628	252.26	0.6
s	Symbolic regression	262.17	654.2	40	1334	12,572	-13.88	0.942
S	Hybrid comparator	1207.14	2436.03	293	5841	15,039	757.31	0.705
S	Symbolic regression	181.3	463.53	32	882	9771	37.46	0.989

4. Discussion

4.1. Interpretation of the Simulation-Derived (R , s , S) Lost-Sales Reference Map

The central contribution of this study is the derivation of explicit closed-form equations for the reorder point s and the order-up-to level S in a lost-sales periodic-review (R , s , S) inventory system under a type-II unit fill-rate criterion. These equations are not imposed from classical safety-stock approximations, but are discovered from a large simulation-derived reference map constructed under controlled operating conditions.

The reference map is novel because it directly links μ , σ , R , L , and the exact achieved unit fill rate to both policy parameters s and S . This is different from periodic-review base-stock special cases, which govern a single order-up-to level and do not determine a separate reorder point s . Existing periodic-review base-stock fill-rate studies provide important analytical foundations, but they do not solve the full two-parameter lost-sales (R , s , and S) inverse design problem considered here [2,6,8,9]. The present mapping is therefore explicitly of the form $(\mu, \sigma, R, L, FR) \rightarrow (s, S)$ where FR denotes the realized fraction of demanded product units supplied immediately from on-hand inventory.

The simulated (R , s , and S) lost-sales policy map is structured but not trivially smooth. The policy parameters generally increase with demand scale, demand variability, review period, lead time, and service requirement, as expected from inventory theory. However, substantial local variability remains even under fixed nominal operating conditions. This is expected because the reference map is generated from finite, integer-valued demand paths. Different statistically valid replicas with similar μ and σ can yield different depletion sequences, stockout timing, and first-feasible policies. Therefore, the reference map should be interpreted as a finite-horizon, path-dependent, policy-consistent empirical representation of the periodic-review lost-sales (R , s , and S) design relationship, not as a single deterministic theoretical curve.

A further structural feature of the reference map is that it combines all tested review-period/lead-time timing regimes in a single learning domain. The dataset is not separated

into equations for $R < L$, $R = L$, and $R > L$ cases, nor into separate equations by mean demand, demand variability, or fill-rate band. This design choice makes the symbolic regression task more difficult because the replenishment dynamics differ across timing regimes. When R is short relative to L , the system reviews inventory frequently while replenishment is delayed; when R is long relative to L , review frequency becomes the dominant timing restriction; and when $R = L$, review and replenishment delay operate on comparable time scales. The selected equations should therefore be interpreted as global tested-domain approximations across these regimes rather than specialized equations for a single timing configuration.

This interpretation is essential for assessing the accuracy of the equations. In the final dataset, the maximum retained values reach $s = 37,797$ and $S = 51,948$, while the local spread of policy parameters is also substantial, especially near high service levels and long R, L combinations. For $\mu = 1000$, the maximum within-CV* replica spread reaches 16,987 for s and 4138 for S , while the broader spread across CV* classes and replicas reaches 20,533 for s and 5786 for S . Consequently, symbolic regression errors and analytical-comparator errors should not be interpreted only through aggregate R^2 . Absolute-error metrics such as MAE, RMSE, median absolute error, and upper-percentile absolute error are necessary because they express deviations directly in inventory units and are therefore more relevant for operational interpretation [41–43].

4.2. Relationship to Previous Inventory Theory and the Hybrid Normal-Loss Comparator

The findings should be interpreted as complementary to previous analytical inventory theory rather than as a replacement for it. Existing periodic-review studies provide important results for fill-rate evaluation, base-stock special cases, safety-stock approximations, aggregate service constraints, and lot-sizing variants. However, these studies generally address problems that are structurally different from the inverse design problem studied here. The present objective is narrower and more operational: to determine both s and S directly from μ, σ, R, L , and FR for a periodic-review lost-sales (R, s, S) system.

This distinction explains why previous analytical formulas are not treated as direct competitors. Periodic-review base-stock models govern a single order-up-to parameter and therefore do not determine a separate reorder point s [6,8,9]. Periodic-review lost-sales studies with target service levels are closer to the present setting, but they generally address fill-rate approximation for a given safety stock, lot-sizing, case-pack restrictions, aggregate service constraints, or policy evaluation rather than a direct closed-form inverse mapping from $(\mu, \sigma, R, L, \text{FR})$ to both s and S [10,11].

The Tijms–Groenevelt approximation is the closest formal analytical predecessor because it addresses service-level constraints in (s, S) -type systems and uses a service concept aligned with the fraction of demand satisfied directly from stock on hand [1]. It assumes that $Q = S - s$ is already known externally, for example, from an EOQ-type calculation. Nevertheless, it is not equivalent to the present design problem. EOQ logic determines an order quantity based on deterministic cost trade-offs and does not account for periodic-review variables R and L , stochastic demand variability σ , or the lost-sales unit fill-rate target FR. Therefore, an EOQ-derived Q would introduce an external cost-based assumption rather than solve the direct service-constrained mapping $(\mu, \sigma, R, L, \text{FR}) \rightarrow (s, S)$. The present study instead estimates both s and S directly. Supplying Q from the simulation-derived reference dataset would make the comparison circular, because it would give the analytical method information that the proposed equations are designed to predict. For this reason, the Tijms–Groenevelt method is best interpreted as a formal reference point rather than a primary numerical benchmark.

The hybrid normal-loss comparator was included only as an illustrative analytical reference. Its purpose was to quantify how far a transparent normal-loss-style construc-

tion can approximate the simulation-derived lost-sales policy map. This framing is important because normal-loss fill-rate logic is central to classical service-level inventory analysis, while the present problem requires a direct two-parameter periodic-review (R, s, S) equation pair under lost sales and unit fill-rate measurement [2].

The comparator results should therefore be interpreted as an analytical context rather than as a comparison against an equivalent published solution. Its undefined $FR = 100\%$ boundary, negative raw reorder-point values in part of the common domain, and lower accuracy relative to the simulation-derived reference map show that direct transfer of normal-loss protection logic does not reproduce the complete lost-sales periodic-review $(R, s, \text{and } S)$ policy structure, in which both s and S must be determined. The detailed numerical evidence is reported in Section 3.3 and Table 12; in the Discussion, the main implication is that the proposed equations should be evaluated primarily against the simulation-derived reference map.

This positioning is central to the novelty claim. The present research is not a refinement of a known closed-form periodic-review lost-sales $(R, s, \text{and } S)$ equation. Rather, it constructs a simulation-derived policy map for a setting where no directly equivalent explicit equation pair appears to be available, and then converts that map into compact, dimensionally admissible equations for both policy parameters. This is consistent with the prior literature showing that lost-sales inventory systems are analytically difficult and that periodic-review lost-sales approximations can be sensitive to the exact service definition, replenishment timing, demand discreteness, and treatment of unmet demand [10,14].

4.3. Contribution of the Symbolic Regression Equation System

The symbolic regression equation system is the main outcome of this study because it converts the exhaustive simulation-derived reference map into explicit policy-parameterization rules. In the present approach, symbolic regression compresses the simulated relationship between demand descriptors, review period, lead time, fill rate, and the retained first service-feasible policy pair into two auditable closed-form expressions for s and S . This is consistent with the broader role of symbolic regression in interpretable scientific model discovery [18,22,23].

The contribution of the proposed equations is not merely that they estimate two numerical outputs. In a periodic-review $(R, s, \text{and } S)$ system, s and S jointly define the policy: the reorder point determines when replenishment is triggered, while the order-up-to level determines the post-order inventory position and future exposure to stockout. Therefore, separate formulas for s and S must be interpreted operationally as one coordinated policy-pair equation system. The proposed mapping is thus best understood as $(\mu, \sigma, R, L, FR) \rightarrow (s, S)$, not as two unrelated scalar regressions.

This joint interpretation is important because the achieved fill rate under lost sales depends on the interaction between the two thresholds. A formula for s alone would still require an externally specified S , while a formula for S alone would not determine the reorder trigger. The proposed equation system avoids this conditional structure by estimating both policy parameters directly from the same demand and operating descriptors. In this sense, it addresses the practical inverse-design problem more directly than approaches that require a predetermined $S - s$, a cost-based EOQ input, or a one-parameter base-stock approximation.

On the full $FR \leq 100\%$ domain, containing 5,002,469 policy-consistent observations, the selected s -equation achieved $MAE = 266.03$ and $R^2 = 0.941$, while the selected S -equation achieved $MAE = 187.2$ and $R^2 = 0.989$. Together with the common-domain comparison in Table 12, these results show that the equations preserve the main structure of the simulation-derived reference map, substantially improve on the illustrative hybrid comparator, and remain evaluable at the finite-horizon $FR = 100\%$ boundary. Their role is therefore

not to replace analytical inventory theory, but to provide a tested-domain closed-form parameterization for a lost-sales (R , s , and S) setting where no directly equivalent explicit equation pair is available.

The algebraic forms should be interpreted empirically. Notably, although the square-root operator was available in the symbolic search, the selected s -equation uses a variability term proportional to σRFR^4 rather than a classical square-root protection-period term, indicating that the discovered lost-sales reorder-point structure differs from standard normal-loss safety-stock logic within the tested domain. The s -equation reflects the stronger sensitivity of the reorder threshold to lead time, demand variability, review frequency, and high service levels, whereas the S -equation captures a smoother order-up-to response across the combined review and lead-time structure. These forms are not constrained to reproduce classical safety-stock equations; they are compact approximations to the finite-horizon, integer-demand, lost-sales policy map generated in this study. Their practical value lies in converting a computationally expensive exhaustive-search procedure into simple formulas that can be implemented in spreadsheets, ERP systems, or replenishment-planning software within the tested domain.

4.4. The $FR = 100\%$ Finite-Horizon Lost-Sales Boundary

The $FR = 100\%$ boundary is one of the clearest conceptual differences between the proposed approach and the normal-loss comparator. Under normal-loss service-level logic, fill rate is linked to expected shortage volume, not only to the probability of no stockout [2]. Under an unbounded normal approximation, an exact zero expected shortage requires an infinite safety factor. Therefore, the normal-loss comparator cannot produce finite policy parameters at $FR = 100\%$. This is a structural property of the approximation, not a numerical failure.

In the present study, $FR = 100\%$ indicates a finite-horizon simulation, meaning zero lost product units over the generated demand path of length $T = 3650$. Under this definition, finite (s , S) policies can exist, and the final reference dataset contains 50,400 retained zero-lost-sales observations, one for every (μ , CV^* , replica, R , and L) configuration. This gives the symbolic regression equations a complete upper service boundary that the hybrid normal-loss comparator cannot cover. The distinction is consistent with simulation-based inventory analysis, where the observed service outcome is conditional on the simulated demand path and system assumptions rather than on the full theoretical demand support [17].

The $FR = 100\%$ subset is empirically more difficult than the $FR < 100\%$ domain. For SSR , MAE increases from 262.17 on $FR < 100\%$ to 645.56 on $FR = 100\%$. For S_{SR} , MAE increases from 181.3 to 765.98. The maximum absolute errors for both equations also occur at the $FR = 100\%$ boundary. These results do not invalidate the equations. Rather, they show that the zero-lost-sales boundary is a distinct part of the policy map with a larger policy-parameter scale and larger local variability.

The reason is both operational and numerical. At very high fill-rate levels, especially at $FR = 100\%$, the retained reference policies are governed not only by average demand but also by unfavorable finite-demand-path realizations and their timing relative to review epochs, lead-time arrivals, and the one-open-order restriction. Under the lost-sales (R , s , and S) dynamics, the $FR = 100\%$ boundary requires inventory positions that prevent even a single lost unit over the simulated demand path. The enumeration must therefore continue until it finds a first-feasible (s , S) pair that eliminates that loss. This mechanism explains the larger required s and S values in the $FR = 100\%$ region and the associated increase in inventory exposure, including AIL. Because MAE is measured in inventory units, part of the increased $FR = 100\%$ error reflects the larger scale of the target policy

parameters; however, that larger scale is itself a consequence of the lost-sales (R, s, S) inventory dynamics near the zero-lost-sales boundary.

The interpretation must remain precise. $FR = 100\%$ in this paper does not mean zero stockout probability for all possible future demand realizations. It means zero lost product units over the finite simulated demand path. Therefore, the proposed equations should be interpreted as tested-domain approximations to a finite-horizon zero-lost-sales policy boundary, not as universal guarantees of complete availability under an unbounded demand distribution.

4.5. Practical Implications

The proposed equations provide a direct, explicit, and tested-domain mapping from $(\mu, \sigma, R, L, \text{ and } FR)$ to the lowest feasible (s, S) policy pair observed in the periodic-review lost-sales simulation reference map. This has direct managerial importance because s and S are operational decision parameters: s determines when replenishment is triggered, while S determines the post-order inventory position. Together, they influence product availability, average inventory level, working capital, storage capacity, replenishment frequency, shipment size, and exposure to lost sales. A closed-form, auditable mapping from demand and service descriptors to both policy thresholds therefore provides a practical decision-support contribution as well as a theoretical inventory-control contribution [2,10,26].

The equations can support master-data maintenance, rapid policy recalculation, sensitivity analysis, and first-pass parameterization of item-location combinations in ERP or replenishment-planning environments. They should be implemented together with demand monitoring because changes in μ, σ, R, L , or FR imply recalculated policy parameters. Shorter practical demand histories may contain fewer replenishment-protection cycles than the long simulation horizon used here, making estimates of μ, σ , achieved service, and required (s, S) more path-dependent [1,17,18,23].

The auxiliary AIL check reported in Section 2.5 supports the operational relevance of the first-feasible policy construction. In the evaluated high-service subset, later service-feasible alternatives did not produce a lower AIL than the first retained policy in 98.33% of evaluated candidate-policy records, while the rare lower-AIL alternatives produced only a small average relative reduction. This supports the low-inventory interpretation of the enumeration rule, while cost-optimality and full-domain AIL behavior remain separate research questions. Because AIL is directly connected to holding cost, working capital, storage requirements, logistics frequency, emissions, and inventory exposure, it should be explicitly modeled in future work as a separate response variable within the same periodic-review lost-sales ($R, s, \text{ and } S$) framework [35,46].

4.6. Limitations and Future Research

The present study establishes a controlled simulation-derived equation system for the periodic-review ($R, s, \text{ and } S$) policy under lost sales and a type-II unit fill-rate criterion. The tested domain is defined by stationary, normal-like, non-negative integer demand replicas; a single item; a single echelon; deterministic review periods and lead times; at most one outstanding replenishment order; and a finite simulation horizon of $T = 3650$ periods. Consequently, the proposed equations should not be interpreted as validated for systems that deliberately permit multiple simultaneous open replenishment orders for the same product at the same stocking location. The proposed equations should not be applied without revalidation to intermittent, seasonal, trend-driven, promotion-driven, or heavily skewed demand, because such demand regimes may require additional descriptors beyond μ and σ and would generally produce a different simulation-derived reference map. This controlled domain provides a clear basis for deriving and validating

the mapping $(\mu, \sigma, R, L, FR) \rightarrow (s, S)$ while also defining a structured research program for extending the equations to broader inventory environments. Lost-sales inventory systems are known to be analytically demanding, and periodic-review lost-sales applications with service-level requirements require careful treatment of replenishment timing, unmet demand, and service definition [10,14].

A first extension concerns service-level specialization. The full-domain equations provide compact approximations across approximately $FR \in [0.01, 1]$, including the finite zero-lost-sales boundary. The largest absolute deviations occur in high-service regions, especially $FR \geq 95\%$. These zones are operationally important in modern supply chains because critical products, contractual service items, spare parts, and high-priority retail goods often require very high immediate product availability. Future work can therefore derive specialized high-service equations for s and S , optimized for lower upper-tail errors in these regions while preserving the tested lost-sales $(R, s, \text{ and } S)$ structure. Future work should also examine AIL, search effort, and the attainability of exact fill-rate targets over finer high-service grids because the transition from very high nonzero-loss service levels to the $FR = 100\%$ boundary may be operationally large.

A second extension concerns the observation horizon. The present reference map is based on $T = 3650$ periods, while the maximum tested value of $R + L$ is 30; even the longest replenishment-protection setting is therefore evaluated across more than 120 $R + L$ -equivalent intervals. This long-horizon design gives the system many review, replenishment, depletion, and stockout-opportunity cycles, so the equations should be interpreted as approximations to a long-horizon reference policy map rather than formulas calibrated from short demand histories. Shorter observation periods pose a different, practically important research problem. If the horizon is close to the replenishment-protection period, for example, $T = 2(R + L)$, the simulated system has only a few opportunities to review inventory, place orders, receive replenishment, deplete stock, and experience lost sales. Under such conditions, the retained (s, S) values would be expected to show stronger path dependence and more local irregularity, even when the underlying demand process is stationary. This is highly relevant to real-world implementation because companies often estimate demand parameters from short historical datasets and still need to make rapid policy decisions. Future research should therefore study the stability of s , S , achieved FR , and AIL as a function of $T / (R + L)$, comparing short-horizon operational estimates with the long-horizon reference map developed in the present paper. Future work may also derive regime-specific equations for $R < L$, $R = L$, and $R > L$ subsets and compare them with the global equations reported here.

A related timing extension concerns calendar structure and workweek schedules. The present model treats time as a sequence of equal generic periods and assumes that review periods and lead times are measured directly in those periods. It does not distinguish between calendar days and working days, nor does it model weekends, holidays, supplier calendars, warehouse calendars, or transport calendars. This is relevant because a nominal lead time may correspond to different actual delivery dates depending on whether replenishment operates under a 7-day, 6-day, or 5-day workweek. Prior simulation-based studies of periodic-review $(R, s, \text{ and } S)$ systems have shown that five-day and seven-day workweek schedules can affect inventory, cost, environmental, and replenishment-planning outcomes [35,36]. More broadly, periodic-review inventory research has shown that seasonal or calendar-dependent lead times can materially affect inventory-control decisions and performance, especially for perishable and other time-sensitive products where delivery timing directly influences availability, waste, and service outcomes [47]. Future research should therefore extend the present simulation-and-symbolic regression framework to calendar-aware lost-sales $(R, s, \text{ and } S)$ systems in which nominal

R and L values are converted into effective review, order-placement, and replenishment dates under alternative workweek schedules.

A third extension concerns the demand process. The present equations are calibrated on controlled normal-like integer demand, which is appropriate for the tested research domain. Practical item demand may also exhibit intermittency, seasonality, trend, autocorrelation, promotion effects, substitution effects, or skewed and heavy-tailed behavior. Future studies can apply the same simulation-and-symbolic regression framework to these demand classes and evaluate how the $(R, s, \text{and } S)$ lost-sales mapping changes when the demand process departs from the normal-like benchmark. This would extend the practical range of the method while keeping the service criterion anchored in immediate product-unit fulfillment.

A fourth extension concerns the feasible-policy set beyond the first retained (s, S) pair. The present study retains the first feasible policy in the enumeration order for each $(\mu, \sigma, R, L, FR^*)$ configuration, thereby constructing a lowest-threshold service-feasible policy map. Later feasible policies may satisfy the same service class while producing different AIL, order frequencies, shipment sizes, replenishment quantities, costs, and emissions. Future research should therefore compare first-feasible policies with AIL-minimizing, cost-minimizing, and Pareto-efficient alternatives. This would convert the current service-constrained policy map into a broader feasible-policy frontier for periodic-review lost-sales $(R, s, \text{and } S)$ systems.

A fifth extension concerns the average inventory level. The preliminary AIL check in the present study indicates that AIL generally increases with higher service levels for the retained first-feasible policies, while finite-demand and integer-policy effects still create local irregularities. AIL should therefore be modeled explicitly as a response variable in future work. This is practically important because AIL connects service policy to holding costs, working capital, storage capacity, exposure to obsolescence, and inventory risk management. Prior periodic-review base-stock research has treated AIL as a distinct analytical object, supporting the broader relevance of modeling average inventory separately from service-level equations [46]. In the present setting, the natural next equation target is $AIL = f(\mu, \sigma, R, L, FR)$.

A sixth extension concerns the order-up-to gap $S - s$. Earlier conditional analytical approaches require $S - s$ as an external input before estimating s , whereas the present study estimates s and S directly [1]. Modeling $S - s$ as a separate response variable would provide a useful bridge between the present direct equation system and earlier conditional (s, S) -type approximations. It would also be relevant for practitioners and software systems that analyze replenishment policies through order-up-to gaps, shipment sizes, or replenishment-cycle quantities.

The final extension concerns the symbolic regression design. The present equations were derived using a compact symbolic regression operator set comprising addition, subtraction, multiplication, division, and the square root operator, to preserve interpretability, dimensional consistency, and ease of implementation. Future symbolic regression searches can test broader operator sets, including integer-oriented functions, rounding operators, modulo terms, logarithmic transformations of dimensionless variables, and additional nonlinear terms. The modulo operator may be especially relevant because earlier $(R, s, \text{and } S)$ -based simulation and symbolic regression work found that modulo terms can appear in useful replenishment-equation structures [25]. Such extensions should preserve dimensional validity, numerical stability, and transparent implementation because symbolic regression expressions can otherwise achieve high fitted accuracy while remaining difficult to implement robustly [23,48].

5. Conclusions

This study demonstrates that explicit, interpretable closed-form equations can be derived for both policy thresholds of a lost-sales periodic-review (R , s , and S) inventory system under a type-II unit-fill-rate criterion. The central conclusion is that the inverse policy-design problem can be represented directly: demand level, demand variability, review period, lead time, and fill-rate target can be mapped to both the reorder point s and the order-up-to level S without requiring an externally specified order-up-to gap or repeated simulation search for every new operating condition.

The contribution is both theoretical and practical. From a theoretical perspective, the study addresses a specific gap in inventory-control modeling by providing a direct equation pair for a two-threshold lost-sales periodic-review policy, whereas existing base-stock, conditional (s , S), cost-based, or normal-loss formulations do not provide an equivalent mapping for the tested setting. From a practical perspective, the equations provide auditable first-pass policy-parameterization rules that can be implemented in spreadsheets, ERP systems, or replenishment-planning tools when the required demand and timing descriptors are available.

The equations should be interpreted as tested-domain parameterization rules, not as universal inventory-policy formulas. Their validity is tied to the controlled simulation domain: stationary normal-like demand, a single item and echelon, deterministic review periods and lead times, lost sales, finite simulation horizon, and at most one outstanding replenishment order for the same product-location. Application outside this domain requires revalidation or extension of the simulation-derived reference map.

Future research should extend the framework to broader demand regimes, multiple open-order item-location systems, cost-based and multi-objective policy evaluation, specialized high-service equations, observation-horizon effects, average-inventory-level equations, later-feasible policies beyond the first retained pair, and explicit modeling of the order-up-to gap $S - s$. These extensions would further connect simulation-derived symbolic regression with classical inventory theory and practical replenishment-policy design.

Author Contributions: Conceptualization, S.Ž. and J.Ž.; methodology, S.Ž. and J.Ž.; software, S.Ž. and J.Ž.; validation, S.Ž. and J.Ž.; formal analysis, S.Ž. and J.Ž.; investigation, S.Ž. and J.Ž.; resources, S.Ž.; data curation, S.Ž. and J.Ž.; writing—original draft preparation, S.Ž. and J.Ž.; writing—review and editing, S.Ž. and J.Ž.; visualization, S.Ž. and J.Ž.; supervision, S.Ž. and J.Ž.; project administration, S.Ž.; funding acquisition, S.Ž. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding. The publication costs were covered by VORAX d.o.o.

Data Availability Statement: The aggregated results supporting the conclusions of this study are presented in the article. The underlying simulation records, input–output datasets, software implementation, and working files are not publicly available because they are proprietary OptimInventory development resources of VORAX d.o.o. Limited verification data may be made available from the corresponding author upon reasonable request and subject to approval by VORAX d.o.o., for verification purposes only.

Acknowledgments: The authors acknowledge VORAX d.o.o. for providing software resources related to the OptimInventory development environment.

Conflicts of Interest: The authors are associated with VORAX d.o.o., the company developing OptimInventory. VORAX d.o.o. provided software resources for this research. Apart from the authors' roles as researchers and authors, VORAX d.o.o. had no additional role in the design of the study; in

the collection, analysis, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results. The authors declare no other conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

ACF	Autocorrelation function
AE	Absolute error
AIL	Average inventory level
CV	Coefficient of variation
CV*	Nominal coefficient-of-variation construction class
EOQ	Economic order quantity
ERP	Enterprise resource planning
FR	Exact achieved type-II unit fill rate
FR*	Target fill-rate class
MAE	Mean absolute error
R ²	Coefficient of determination
RMSE	Root mean square error
SR	Symbolic regression

References

1. Tijms, H.C.; Groenevelt, H. Simple approximations for the reorder point in periodic and continuous review (s, S) inventory systems with service level constraints. *Eur. J. Oper. Res.* **1984**, *17*, 175–190. [https://doi.org/10.1016/0377-2217\(84\)90232-7](https://doi.org/10.1016/0377-2217(84)90232-7).
2. Sobel, M.J. Fill rates of single-stage and multistage supply systems. *Manuf. Serv. Oper. Manag.* **2004**, *6*, 41–52. <https://doi.org/10.1287/msom.1030.0027>.
3. Johnson, M.E.; Lee, H.L.; Davis, T.; Hall, R. Expressions for item fill rates in periodic inventory systems. *Nav. Res. Logist.* **1995**, *42*, 57–80. [https://doi.org/10.1002/1520-6750\(199502\)42:1%3C57::AID-NAV3220420107%3E3.0.CO;2-2](https://doi.org/10.1002/1520-6750(199502)42:1%3C57::AID-NAV3220420107%3E3.0.CO;2-2).
4. Campo, K.; Gijbrecchts, E.; Nisol, P. Towards understanding consumer response to stock-outs. *J. Retail.* **2000**, *76*, 219–242. [https://doi.org/10.1016/S0022-4359\(00\)00026-9](https://doi.org/10.1016/S0022-4359(00)00026-9).
5. Fitzsimons, G.J. Consumer response to stockouts. *J. Consum. Res.* **2000**, *27*, 249–266. <https://doi.org/10.1086/314323>.
6. Zhang, J.; Zhang, J. Fill rate of single-stage general periodic review inventory systems. *Oper. Res. Lett.* **2007**, *35*, 503–509. <https://doi.org/10.1016/j.orl.2006.08.003>.
7. Teunter, R.H. Note on the fill rate of single-stage general periodic review inventory systems. *Oper. Res. Lett.* **2009**, *37*, 67–68. <https://doi.org/10.1016/j.orl.2008.09.008>.
8. Silver, E.A.; Bischak, D.P. The exact fill rate in a periodic review base stock system under normally distributed demand. *Omega* **2011**, *39*, 346–349. <https://doi.org/10.1016/j.omega.2010.08.003>.
9. Guíjarro, E.; Cardós, M.; Babiloni, E. On the exact calculation of the fill rate in a periodic review inventory policy under discrete demand patterns. *Eur. J. Oper. Res.* **2012**, *218*, 442–447. <https://doi.org/10.1016/j.ejor.2011.11.025>.
10. van Donselaar, K.H.; Broekmeulen, R.A.C.M. Determination of safety stocks in a lost sales inventory system with periodic review, positive lead-time, lot-sizing and a target fill rate. *Int. J. Prod. Econ.* **2013**, *143*, 440–448. <https://doi.org/10.1016/j.ijpe.2011.05.020>.
11. van Donselaar, K.H.; Broekmeulen, R.A.C.M.; de Kok, A.G. Heuristics for setting reorder levels in periodic review inventory systems with an aggregate service constraint. *Int. J. Prod. Econ.* **2021**, *237*, 108137. <https://doi.org/10.1016/j.ijpe.2021.108137>.
12. Zipkin, P.H. Old and new methods for lost-sales inventory systems. *Oper. Res.* **2008**, *56*, 1256–1263. <https://doi.org/10.1287/opre.1070.0471>.
13. Huh, W.T.; Janakiraman, G.; Muckstadt, J.A.; Rusmevichientong, P. Asymptotic optimality of order-up-to policies in lost sales inventory systems. *Manag. Sci.* **2009**, *55*, 404–420. <https://doi.org/10.1287/mnsc.1080.0945>.
14. Bijvank, M.; Vis, I.F.A. Lost-sales inventory theory: A review. *Eur. J. Oper. Res.* **2011**, *215*, 1–13. <https://doi.org/10.1016/j.ejor.2011.02.004>.
15. Bijvank, M.; Huh, W.T.; Janakiraman, G.; Kang, W. Robustness of order-up-to policies in lost-sales inventory systems. *Oper. Res.* **2014**, *62*, 1040–1047. <https://doi.org/10.1287/opre.2014.1298>.

16. Bijvank, M.; Johansen, S.G. Periodic review lost-sales inventory models with compound Poisson demand and constant lead times of any length. *Eur. J. Oper. Res.* **2012**, *220*, 106–114. <https://doi.org/10.1016/j.ejor.2012.01.041>.
17. Gosavi, A. *Simulation-Based Optimization: Parametric Optimization Techniques and Reinforcement Learning*; Springer: New York, NY, USA, 2015. <https://doi.org/10.1007/978-1-4899-7491-4>.
18. Cranmer, M. Interpretable machine learning for science with PySR and SymbolicRegression.jl. *arXiv* **2023**, arXiv:2305.01582. <https://doi.org/10.48550/arXiv.2305.01582>.
19. Tonda, A. Review of PySR: High-performance symbolic regression in Python and Julia. *Genet. Program. Evolvable Mach.* **2025**, *26*, 7. <https://doi.org/10.1007/s10710-024-09503-4>.
20. Hill, R.M.; Johansen, S.G. Optimal and near-optimal policies for lost sales inventory models with at most one replenishment order outstanding. *Eur. J. Oper. Res.* **2006**, *169*, 111–132. <https://doi.org/10.1016/j.ejor.2004.05.022>.
21. Bendre, A.B.; Nielsen, L.R. Inventory control in a lost-sales setting with information about supply lead times. *Int. J. Prod. Econ.* **2013**, *142*, 324–331. <https://doi.org/10.1016/j.ijpe.2012.12.002>.
22. Schmidt, M.; Lipson, H. Distilling free-form natural laws from experimental data. *Science* **2009**, *324*, 81–85. <https://doi.org/10.1126/science.1165893>.
23. Makke, N.; Chawla, S. Interpretable scientific discovery with symbolic regression: A review. *Artif. Intell. Rev.* **2024**, *57*, 2. <https://doi.org/10.1007/s10462-023-10622-0>.
24. Diveev, A.; Sofronova, E.; Konyrbaev, N. Solving the Control Synthesis Problem Through Supervised Machine Learning of Symbolic Regression. *Mathematics* **2024**, *12*, 3595. <https://doi.org/10.3390/math12223595>.
25. Žic, S.; Žic, J.; Đukić, G. Efficient planning and optimization of inventory replenishments for sustainable supply chains operating under (R, s, S) policy. *Sustain. Futures* **2023**, *5*, 100110. <https://doi.org/10.1016/j.sftr.2023.100110>.
26. Silver, E.A.; Pyke, D.F.; Thomas, D.J. *Inventory and Production Management in Supply Chains*, 4th ed.; CRC Press: Boca Raton, FL, USA, 2016. <https://doi.org/10.1201/9781315374406>.
27. D’Agostino, R.B.; Pearson, E.S. Tests for departure from normality. Empirical results for the distributions of b_2 and $\sqrt{b_1}$. *Biometrika* **1973**, *60*, 613–622. <https://doi.org/10.1093/biomet/60.3.613>.
28. Shapiro, S.S.; Wilk, M.B. An analysis of variance test for normality (complete samples). *Biometrika* **1965**, *52*, 591–611. <https://doi.org/10.1093/biomet/52.3-4.591>.
29. Anderson, T.W.; Darling, D.A. Asymptotic theory of certain “goodness of fit” criteria based on stochastic processes. *Ann. Math. Stat.* **1952**, *23*, 193–212. <https://doi.org/10.1214/aoms/1177729437>.
30. Wald, A.; Wolfowitz, J. On a test whether two samples are from the same population. *Ann. Math. Stat.* **1940**, *11*, 147–162. <https://doi.org/10.1214/aoms/1177731909>.
31. Box, G.E.P.; Pierce, D.A. Distribution of residual autocorrelations in autoregressive-integrated moving average time series models. *J. Am. Stat. Assoc.* **1970**, *65*, 1509–1526. <https://doi.org/10.1080/01621459.1970.10481180>.
32. Ljung, G.M.; Box, G.E.P. On a measure of lack of fit in time series models. *Biometrika* **1978**, *65*, 297–303. <https://doi.org/10.1093/biomet/65.2.297>.
33. Schwartz, J.D.; Wang, W.; Rivera, D.E. Simulation-based optimization of process control policies for inventory management in supply chains. *Automatica* **2006**, *42*, 1311–1320. <https://doi.org/10.1016/j.automatica.2006.03.019>.
34. OptimInventory. Available online: <http://www.optiminventory.com> (accessed on 1 May 2026).
35. Žic, J.; Žic, S. Multi-criteria decision making in supply chain management based on inventory levels, environmental impact and costs. *Adv. Prod. Eng. Manag.* **2020**, *15*, 151–163. <https://doi.org/10.14743/apem2020.2.355>.
36. Žic, J.; Žic, S.; Đukić, G. Quantitative assessment of green inventory management in supply chains: Simulation-based study of economic and environmental outcomes aligned with ISO 14083 standard. *Appl. Sci.* **2024**, *14*, 9507. <https://doi.org/10.3390/app14209507>.
37. Kleinau, P.; Thonemann, U.W. Deriving inventory-control policies with genetic programming. *OR Spectr.* **2004**, *26*, 521–546. <https://doi.org/10.1007/s00291-004-0159-5>.
38. Lopes, R.L.; Figueira, G.; Amorim, P.; Almada-Lobo, B. Cooperative coevolution of expressions for (r, Q) inventory management policies using genetic programming. *Int. J. Prod. Res.* **2020**, *58*, 509–525. <https://doi.org/10.1080/00207543.2019.1597293>.
39. Ghaddar, B.; Sakr, N.; Asiedu, Y. Spare parts stocking analysis using genetic programming. *Eur. J. Oper. Res.* **2016**, *252*, 136–144. <https://doi.org/10.1016/j.ejor.2015.12.041>.
40. Udrescu, S.M.; Tegmark, M. AI Feynman: A physics-inspired method for symbolic regression. *Sci. Adv.* **2020**, *6*, eaay2631. <https://doi.org/10.1126/sciadv.aay2631>.

41. Hyndman, R.J.; Koehler, A.B. Another look at measures of forecast accuracy. *Int. J. Forecast.* **2006**, *22*, 679–688. <https://doi.org/10.1016/j.ijforecast.2006.03.001>.
42. Willmott, C.J.; Matsuura, K. Advantages of the mean absolute error (MAE) over the root mean square error (RMSE) in assessing average model performance. *Clim. Res.* **2005**, *30*, 79–82. <https://doi.org/10.3354/cr030079>.
43. Chai, T.; Draxler, R.R. Root mean square error (RMSE) or mean absolute error (MAE)? Arguments against avoiding RMSE in the literature. *Geosci. Model Dev.* **2014**, *7*, 1247–1250. <https://doi.org/10.5194/gmd-7-1247-2014>.
44. Kvålseth, T.O. Cautionary note about R^2 . *Am. Stat.* **1985**, *39*, 279–285. <https://doi.org/10.1080/00031305.1985.10479448>.
45. Chicco, D.; Warrens, M.J.; Jurman, G. The coefficient of determination R-squared is more informative than SMAPE, MAE, MAPE, MSE and RMSE in regression analysis evaluation. *PeerJ Comput. Sci.* **2021**, *7*, e623. <https://doi.org/10.7717/peerj-cs.623>.
46. Babiloni, E.; Cardós, M.; Guijarro, E. On the exact calculation of the mean stock level in the base stock periodic review policy. *J. Ind. Eng. Manag.* **2011**, *4*, 194–205. <https://doi.org/10.3926/jiem.2011.v4n2.p194-205>.
47. Riezebos, J.; Zhu, S.X. Inventory control with seasonality of lead times. *Omega* **2020**, *92*, 102162. <https://doi.org/10.1016/j.omega.2019.102162>.
48. Kommenda, M.; Burlacu, B.; Kronberger, G.; Affenzeller, M. Parameter identification for symbolic regression using nonlinear least squares. *Genet. Program. Evolvable Mach.* **2020**, *21*, 471–501. <https://doi.org/10.1007/s10710-019-09371-3>.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.