Drug Launch Inventory Simulation:
Finished Goods Inventory Projection for New Product Introduction

by

Won Suh Kim

and

Gianmarco Morea

Submitted on May 8, 2026 in Partial Fulfillment of the Requirements for
the Degree of Master of Applied Science in Supply Chain Management
at the Massachusetts Institute of Technology

Abstract

To support best-in-class pharmaceutical product launches, supply chain planners must reconcile rapidly shifting demand assumptions with lead times that require production months before regulatory approval. Traditionally, this gap compels static inventory over-buffering based on Months On-Hand (MOH) targets to mitigate the risks of supply shortages and lost market share. However, this heuristic-based approach often leads to supply inertia, trapping excess capital and increasing expiry risk.

This project develops a standalone deterministic simulation tool designed to project finished goods inventory across safety stock, cycle stock, and pipeline inventory categories. The tool was applied to a case study of tolebrutinib, a Bruton's tyrosine kinase inhibitor under development for multiple sclerosis indications, to demonstrate its application and estimate the projected benefits of transitioning to a statistical inventory policy. A key contribution is the shift from MOH planning to a statistical methodology that calculates safety stock using Mean Absolute Percentage Error (MAPE)—leveraging historical data from similar on-market products—and target cycle service levels.

The tool is designed to support a more dynamic, cross-functional Integrated Business Planning (IBP) process at Sanofi. Rather than anchoring the process to raw commercial inputs, the tool allows stakeholders to assign weights to various demand scenarios to compile a risk-adjusted aggregate forecast. Any gaps between this aggregate forecast and the commercial team's high-demand scenario are accompanied by immediate visibility into potential lost revenue and profit, facilitating real-time strategic alignment. Furthermore, stakeholder groups can propose live adjustments to input variables, such as service levels and manufacturing lead times, to instantaneously assess inventory impacts. In the case study, the tool projects inventory reductions in excess of 75% relative to MOH heuristics, with corresponding improvements in working capital efficiency.

Capstone Advisor: Dr. Jaime Macias Aguayo
Title: Postdoctoral Associate, Center for Transportation and Logistics

Capstone Co-Advisor: Dr. Maria Jesús Saénz
Title: Executive Director, MIT Supply Chain Management Master's Program

Table of Contents

1 Introduction

1.1 Background

Sanofi is a French biopharmaceutical company comprising three major business units: General Medicines, Vaccines, and Specialty Care (Sanofi, 2024a). It employs over 100,000 people worldwide and generated over €41 billion in revenue in 2024 (Sanofi, 2024c). Its Specialty Care business unit is further broken down into five therapeutic areas: Immunology, Neurology, Oncology, Rare Blood Disorders, and Rare Diseases (Sanofi, 2024a). Sanofi Specialty Care focuses on producing best-in-class—often first-in-class—therapies, such as Dupixent, which boasted eight approved indications at the time of writing (Sanofi, 2024b). This unit's portfolio encompasses most of the company's research and development (R&D) blockbuster pipeline, aligning with the company's goal of becoming an immunology powerhouse powered by AI at scale (Sanofi, 2024a).

1.2 Motivation

Profit margins in the pharmaceutical industry are typically quite high when measuring revenue vs. Cost of Goods Sold (COGS). However, considering that over 95% of compounds fail during development, R&D expenses for drug discovery can add up to billions of euros of capital (National Center for Advancing Translational Sciences, 2019). Therefore, it is imperative to maintain optimal cost efficiency in supply chain operations to maximize cash leverage. Prior to FDA approval and the commercial launch of a new drug, Sanofi builds up finished goods to support forecasted demand. Manufacturing lead times for pharmaceuticals are often measured in months and years, making a make-to-stock manufacturing model the most appropriate to avoid backorders, lost sales, and unmet patient demand. As undesirable as those outcomes are, stockpiling finished goods inventory also presents challenges, including immobilizing capital in assets and creating the risk of expiry and obsolescence.

1.3 Problem Statement and Key Questions

Sanofi's Specialty Care unit identified a gap in its inventory management playbook, particularly regarding new product launches. As it considered multiple demand scenarios and draft inventory projections, there appeared to be a risk of redundant buffers across its supply chain. For example, if it implemented an optimistic demand plan while also increasing safety stock to capture a higher service level, it would buffer two parts of its supply chain without a clearly consequential return on investment. In short, it hypothesized that it was building excessive inventory for new product launches and saw an opportunity to reduce capital investment. To address this problem, we developed InventEdge, an inventory simulation tool designed to simulate the expected inventory outcomes of transitioning from a MOH-based policy to an optimized statistical periodic review policy, serving as a complement to Sanofi's existing Enterprise Resource Planning (ERP) and supply chain planning software.

The questions to be answered include:

- Which simulation methodology is most appropriate for projecting Sanofi's new drug inventory?
- What are the expected inventory outcomes of adopting a statistical periodic review policy in place of MOH heuristics, and can this transition reduce over-buffering while maintaining target service levels?

1.4 Project Goals and Expected Outcomes

This project's goal was to develop an inventory simulation tool that projects inventory for Sanofi's new product launches while maintaining service levels following FDA approval and commercial launch. Per

Sanofi's specifications, our initial focus was to support the launch of tolebrutinib within the United States. The tool, named InventEdge by the project's core team comprising the authors and their Sanofi partners, streamlines scenario planning for finished-goods inventory and enables proactive responses to major demand events, such as delays in the New Drug Application (NDA) process. Based on the Association for Supply Chain Management's (ASCM) 2011 safety stock methodology, InventEdge simulates the expected finished goods inventory outcomes of a statistical periodic review policy, providing projections that can be leveraged as inputs in the company's Integrated Business Planning (IBP) process (King, 2011). Inventory is presented as the sum of safety stock, cycle stock, and pipeline inventory to better quantify the capital investment in finished goods inventory. The tool is both customizable and scalable, allowing Sanofi and future partners to iterate on it for all assets across their full life cycles. For that purpose, documentation was incorporated into the tool's code as part of the deliverable. This project contributes to Sanofi's broader program goal of delivering significant supply chain cost reductions over the next few years.

2 State of the Practice

The key objective was to develop a robust and user-friendly simulation tool that enables Sanofi planners to scenario plan inventory for new SKUs by incorporating the latest information, such as changes to demand scenarios, for user inputs. To build the mathematical foundation for this simulator, we investigated two primary streams of research: classical inventory management policies and the application of simulation for supply chain systems.

2.1 Inventory Management Theory

Inventory refers to the materials and goods necessary for a company to sustain business operations and meet customer demands. When effectively managed, it grows the bottom line by maximizing order fulfillment while simultaneously minimizing inventory levels and associated costs, such as holding costs and obsolescence risks. Excessive inventory reduction, however, introduces stock-out risk, which compromises service levels and customer satisfaction.

This problem is compounded in high-competition markets such as pharmaceuticals, where new product launches have multiplied from dozens in 2015-2019 to over 100 in 2020-2025 (IQVIA, 2025). IQVIA, a leading provider of analytics for pharmaceutical organizations, confirmed that less than 20% of drugs manage to recover from poor launches during the first six months of commercial sales (IQVIA, 2025). In that competitive landscape, each Sanofi drug launch runs the risk of losing significant market share if it suffers supply shortages early in its product lifecycle. For companies like Sanofi that primarily deal in blockbuster drugs that have the potential to drive billions of euros in revenue, the revenue and patient risk associated with supply shortages are too significant to ignore, making optimal inventory management imperative.

Based on its location in the value stream, inventory can typically be classified into three buckets: raw material (RM), work-in-process (WIP), and finished goods (FG) (APICS, 2025). It can be further organized into three functional categories, as shown in Figure 1:

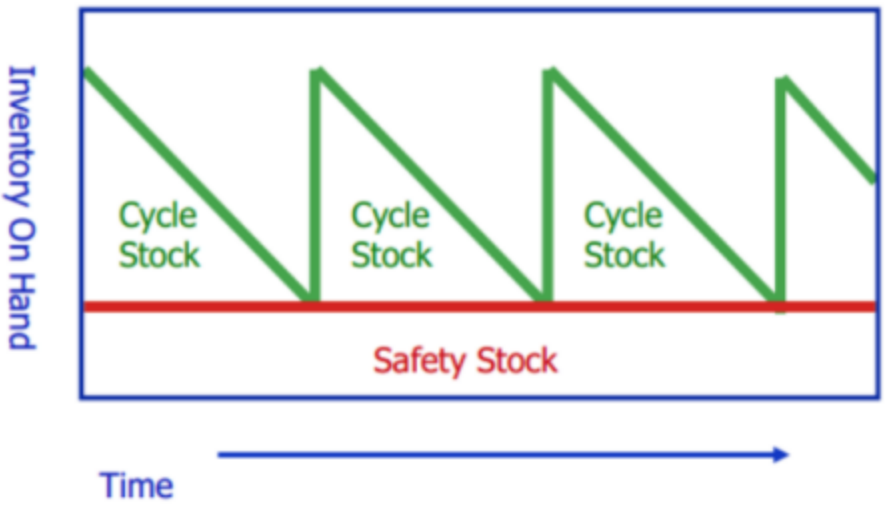

Figure 1: "Inventory chart: Depiction of functional inventory classifications" (Caplice & Ponce, 2022)

- **Cycle Stock**: Inventory to meet demand between replenishments
- **Safety Stock**: Buffer inventory that supports target service levels against variability
- **Pipeline Inventory**: In-process inventory that is not readily available to fulfill orders

A company's balance sheet distinguishes between WIP and FG to capture labor and overhead costs, but functional inventory classifications are defined by supply chain teams to track opportunity costs and determine inventory management strategy (Caplice & Ponce, 2022). As customers take inventory from company shelves, supply chain teams follow inventory policies to determine how much is ordered and when. These policies are embedded in modern ERP and inventory management systems, often consisting of the five shown in Figure 2:

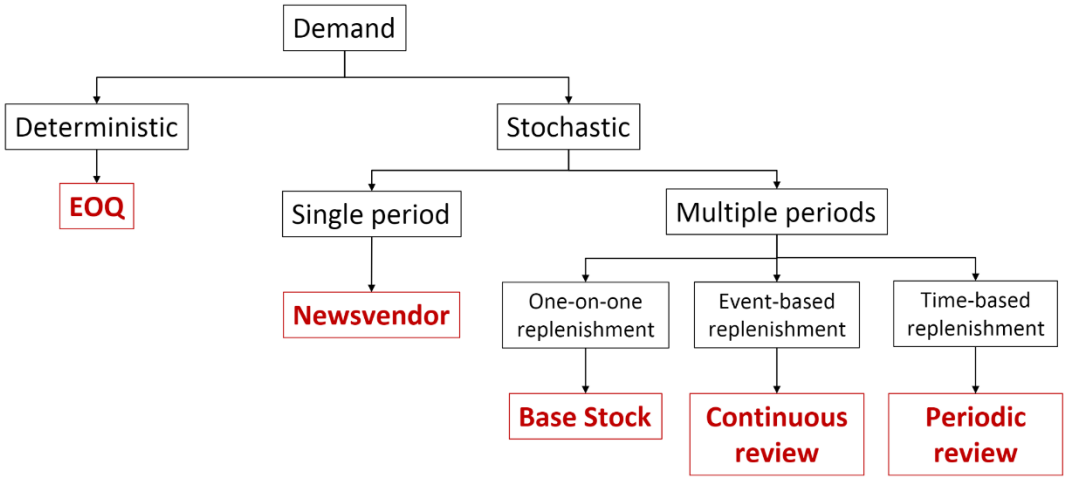

Figure 2: Classical Inventory Policies (Silver et al., 2017)

As shown above, these policies differ in three dimensions: nature of demand, time horizon, and replenishment type. What they have in common is the existence of trade-offs. Economic Order Quantity

(EOQ) balances fixed and variable costs, Newsvendor balances inventory above and below the target level, and the remaining three policies balance costs and service level (Caplice & Ponce, 2022). To understand how to optimize the trade-offs between cost and service level to support Sanofi's drug launches, we narrowed our focus to the three multiple-period models:

- **Base Stock**
  Each item is assigned a target stocking level, and replenishment occurs one-to-one as demand comes in. This policy is well-suited when high responsiveness is required (Caplice & Ponce, 2022).

- **Continuous Review**
  Continuous Review policies are commonly implemented as two-bin Kanban systems—a visual, pull-based inventory system in which two bins are used per item, and the emptying of the first bin triggers replenishment while the second bin supplies demand. These policies require continuous monitoring of the inventory position (Silver et al., 2017). The most important policies within the Continuous Review family are the (s, Q) and (s, S) models. When the inventory position falls below the reorder point s, an order is triggered. In the (s, Q) model, the order quantity Q is fixed, while in the (s, S) model, the order quantity varies to raise the inventory position up to level S. The (s, Q) policy is more suitable for batch production environments, due to the alignment of fixed quantities (Silver et al., 2017).

- **Periodic Review**
  Like the Continuous Review policy, this is also essentially a two-bin system. However, as the name suggests, the inventory position is reviewed at fixed time intervals (Silver et al., 2017). Within the Periodic Review family, the most common models are the (R, S) and (R, s, S) policies. The inventory position is reviewed every R time periods. In the (R, S) model, an order is placed to replenish the inventory up to level S. This policy assumes unlimited capacity, negligible ordering costs, and that demand will be fully met with no lost sales. In an environment with high ordering costs, the (R, s, S) model is more appropriate, as an order is placed only if, at the review time, the inventory position is below s (Silver et al., 2017).

Service levels can be defined by different metrics, including performance-based metrics and cost-based metrics. A performance-based approach would leverage metrics, such as Cycle Service Level (CSL) and Item Fill Rate (IFR), to set a safety factor to meet a target service level (Caplice & Ponce, 2022). A cost-based approach finds a safety factor value that minimizes costs over a time horizon (Caplice & Ponce, 2022). Examples of cost-based metrics include Cost of Stock Out Event (CSOE) and Cost of Item Short (CIS). Both approaches result in a safety factor that subsequently defines a safety stock to capture demand variability over lead time. Because many companies do not collect sufficient data on forecast errors, that variability is commonly represented by the standard deviation of demand over lead time, which assumes the forecast is the mean demand (Caplice & Ponce, 2022).

2.2 Simulation Applications

Simulation is the process of modeling a real-world system, such as a supply chain, and conducting experiments to understand the system's behavior and scenario plan (MIT Center for Transportation & Logistics, 2023). The types of simulation can be categorized by three primary dimensions: the certainty of input parameters, the time dimension, and the timing of system state changes (Law, 2015).

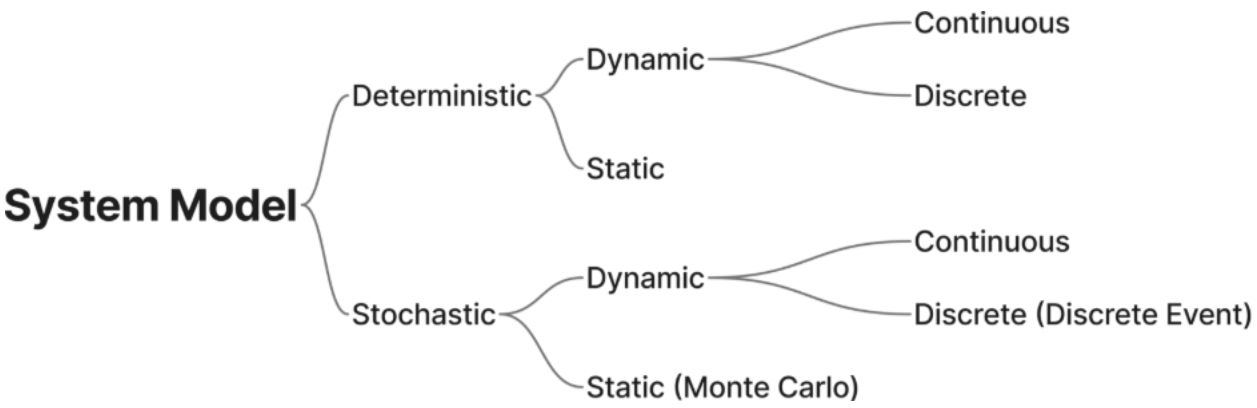

Figure 3: Simulation Types

The branches of simulation types are illustrated in Figure 3, and the decision criteria for defining a simulation model are outlined below:

- **Deterministic vs. Stochastic**: Deterministic models have constant inputs, so they systematically repeat the result. Stochastic models involve at least one random variable, meaning different simulations will produce disparate outputs (Thiede, 2012).

- **Dynamic vs. Static**: Dynamic models display change over time, whereas static models, such as Monte Carlo simulation, only update at one point in time (Thiede, 2012).

- **Continuous vs. Discrete**: Continuous models update the system continuously over time, whereas Discrete models only do so at discrete time intervals (Thiede, 2012). There are also combinations of these models, known as combined discrete-continuous models, which simulate how discrete events affect a continuous variable or vice-versa (Thiede, 2012).

These simulation classifications help distinguish four paradigms, as shown in Figure 4:

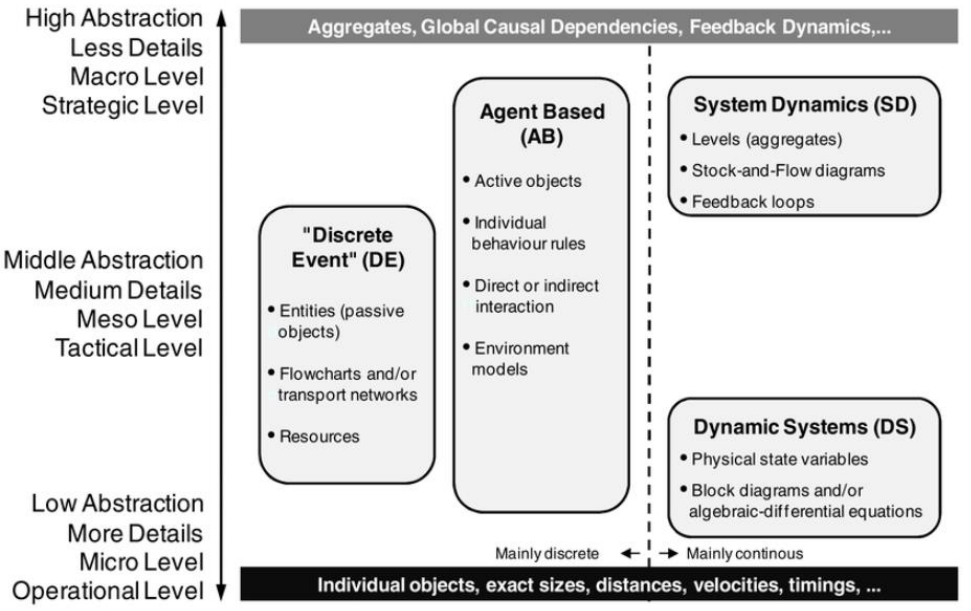

Figure 4: Simulation Paradigms (Thiede, 2012) Adapted From (Borshchev & Filippov, 2004)

The four simulation paradigms can be mapped in dimensions of level of abstraction and time.

- **Discrete-Event (DE)**: A simulation paradigm that is stochastic, dynamic, and discrete. The system instantaneously updates at discrete points in time, represented by events. DE is appropriate when modeling environments with an average level of detail, such as warehouse operations (MIT Center for Transportation & Logistics, 2023).

- **Agent-Based (AB)**: A specialized form of DE simulation (DES) in which agents interact with each other and their environment (MIT Center for Transportation & Logistics, 2023). An agent is an entity acting based on its individual behavioral rules, thus affecting the state of the system (MIT Center for Transportation & Logistics, 2023). A typical application of AB simulation (ABS) is traffic flow simulation, where agents are represented by vehicles (Law, 2015).

- **Systems Dynamics (SD)**: A continuous simulation that focuses on modeling systems at an aggregate, macro level, and emphasizes feedback loops and time delays (Law, 2015). This approach is particularly appropriate for non-linear systems with complex feedback structures (Bala et al., 2017).

- **Dynamic Systems (DS)**: A continuous simulation paradigm on the opposite end of the spectrum compared to SD regarding abstraction and details (Thiede, 2012). Often implemented in the MATLAB programming language, DS simulations typically operate using physical state variables and algebraic equations (Thiede, 2012).

Regardless of simulation type or paradigm, simulation methods should not be confused with optimization methods. Simulation is intended to evaluate the reaction of a system to different sets of parameters; it will not output an optimal solution (Tersine, 1988). This makes it a logical method for answering “what if” questions (MIT Center for Transportation & Logistics, 2023). Other advantages especially relevant to Sanofi’s use case are control over inputs, time expansion, and ease of communication with decision makers (MIT Center for Transportation & Logistics, 2023).

2.3 Demand Forecasting and Planning Assumptions

Demand forecasting is the most critical input to inventory planning and remains a significant source of uncertainty in pharmaceutical drug launches. Traditional forecasting methods that rely on historical sales data are inapplicable, necessitating scenario-based approaches based on market estimates. Sanofi's commercial team develops multiple demand scenarios reflecting different market conditions, competitive dynamics, and adoption trajectories. These scenarios represent expert judgments and planning assumptions rather than probabilistic forecasts of actual market demand.

For the purposes of this project, the model uses a deterministic monthly demand plan derived from weighted commercial scenarios provided by Sanofi stakeholders. By assigning weights to various demand scenarios—low, base, and high—planners create a risk-adjusted aggregate forecast that reflects their collective view of the most likely market outcome. This weighted demand serves as a planning assumption, representing what stakeholders believe the market will look like given current information,

not a stochastic prediction of actual realized demand. Once implemented in production, actual demand is expected to vary around this planning assumption.

To account for this demand variability, the model incorporates historical forecast accuracy from analogous, on-market products. Specifically, the model uses Mean Absolute Percentage Error (MAPE)—the average magnitude of forecast errors from comparable products—to quantify forecast deviations and drive the safety stock calculation. This approach bridges the gap between the deterministic planning assumptions used in the IBP process and the stochastic reality of actual market demand. By leveraging historical forecast accuracy, the tool provides a more statistically grounded measure of demand variability than traditional Months On-Hand (MOH) heuristics, enabling planners to set safety stock targets that reflect realistic forecast error ranges rather than arbitrary multiples of monthly demand.

3 Methodology

This section outlines our tool's methodology, justifies its selection, and details the required data. It also covers the step-by-step logic and mathematical formulas used to simulate launch inventory levels over a user-defined time horizon.

3.1 Selection of Simulation Approach

Following our research and discussions with the Sanofi team, we concluded the following:

1. The inventory category of interest is Finished Goods.
2. Launch drug inventory control is performed through a Periodic Review inventory policy, with a monthly IBP process being leveraged to make inventory management decisions.
3. Inventory output by the simulation is broken down into all three functional inventory classifications—cycle stock, safety stock, and pipeline inventory—and the Unit of Measurement (UOM) is both euros and units.
4. The service level metric is performance-based to support launch excellence as requested by Sanofi. Specifically, we use CSL to calculate a safety factor (King, 2011).
5. MAPE was selected to represent forecast error, per Sanofi's request, as it is the metric used internally. This ensures the safety stock calculation is grounded in the same measure of forecast accuracy already embedded in Sanofi's planning processes, facilitating direct comparability with existing performance benchmarks and minimizing friction during IBP discussions.
6. While software integration and data pipelines were not required for this standalone tool, they may be required upon implementation in an enterprise cloud environment. This was a sponsor request to enable future integration, scaling, and development should our value proposition meet or exceed expectations.

These conclusions collectively point to a deterministic simulation approach. The monthly IBP cycle and finished goods focus establish a fixed, recurring planning horizon that does not require real-time demand simulation. Demand variability is captured through the safety stock calculation using historical MAPE data, eliminating the need for stochastic demand modeling. A deterministic approach also aligns with the standalone, interpretable nature of the tool, enabling stakeholders to directly trace inventory projections to their underlying planning assumptions.

The tool operates in discrete daily time steps, aggregating outputs into monthly summaries aligned with Sanofi's IBP cycle. Rather than simulating demand variability, the model projects inventory using a

deterministic monthly demand plan derived from weighted demand scenarios—representing a planning assumption rather than actual demand realization. Demand variability is captured through the safety stock calculation, which leverages historical forecast accuracy (MAPE) from comparable products, not through stochasticity in the demand forecast itself.

3.2 Data Collection and Preparation

The variable data will consist of user-defined parameters as defined in Table 1:

Group	Input Parameter	Definition
Demand	Scenario Weights (H/B/L)	Probability weights for high, base, and low demand scenarios, summing to 100%
Demand	MAPE (%)	Mean Absolute Percentage Error of the forecast, used to calculate safety stock
Demand	Service Level (%)	Target cycle service level; drives the safety factor Z
Financials	ASP (€)	Average Selling Price per unit
Financials	COGS (€)	Cost of Goods Sold per unit
Policy	Review Period R (Months)	Frequency of inventory review, aligned with the IBP cycle
Policy	Lead Time L (Months)	Replenishment lead time estimated by user
Policy	MOQ (Units)	Minimum Order Quantity; orders are rounded up to the nearest multiple
Initial Conditions	On-Hand	Stock on hand at the start of the simulation horizon
Initial Conditions	On-Order	Units in transit at the start of the simulation horizon
Initial Conditions	Backorders	Unfilled demand at the start of the simulation horizon
Date Range	Start / End Month	Slices the uploaded forecast horizon to define the simulation window

Table 1: User Inputs

Demand forecast data for the high, base, and low scenarios is provided by the commercial team and must be input into the tool for it to function. With these inputs defined, the tool is ready to execute the inventory projection across the user-specified simulation horizon.

3.3 Simulation Logic and Inventory Model

Figure 5 illustrates the general workflow of the tool:

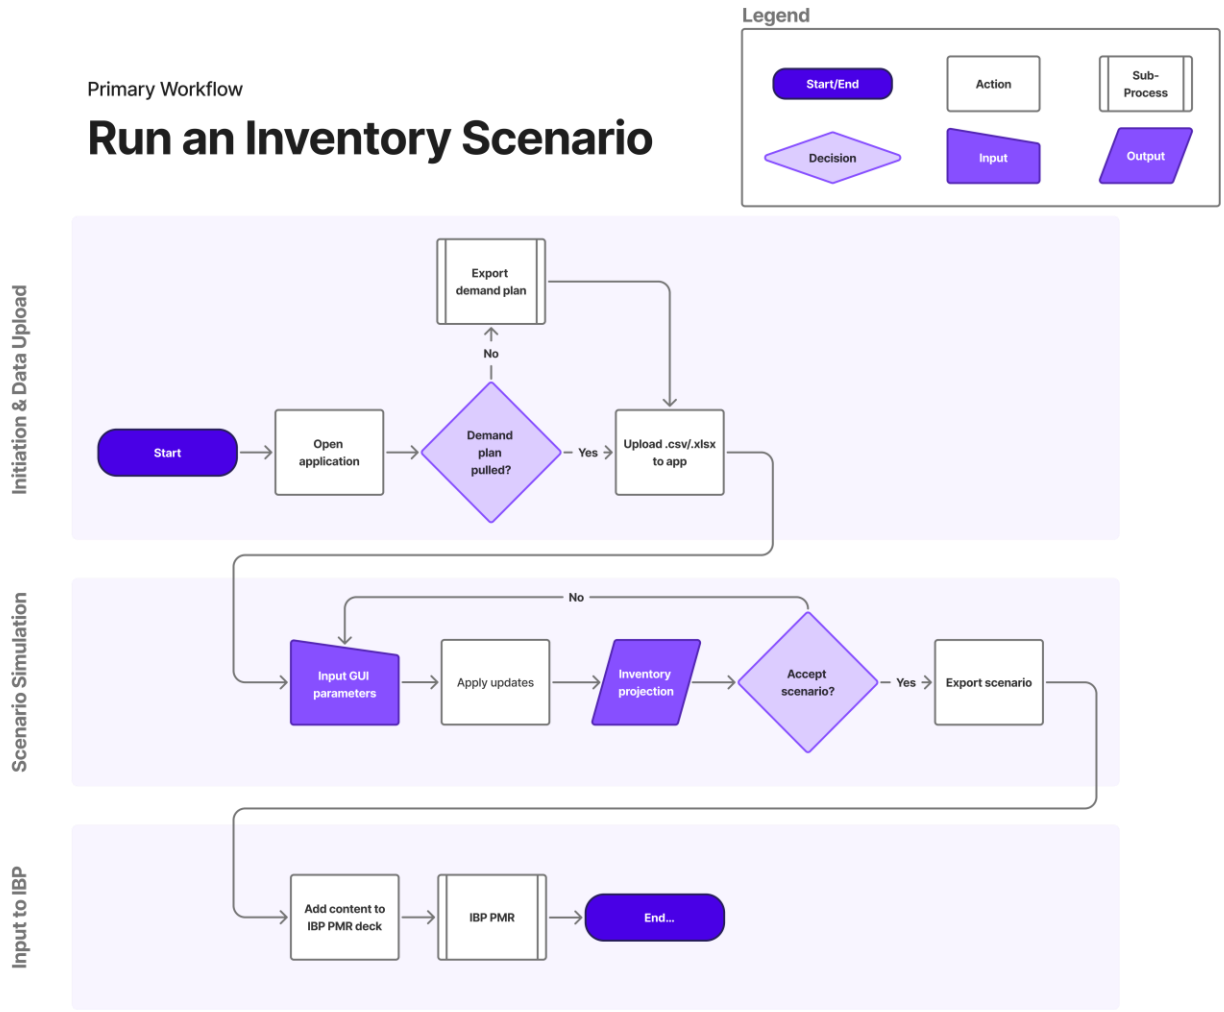

Figure 5: Simulation Tool Workflow

The inventory projection will break down inventory into three buckets: safety stock, cycle stock, and pipeline inventory. The calculations will be defined by the following equations:

Z: Safety Factor
σ_{DL+R}: Standard Deviation of Demand over Lead Time and Review Period
D: Forecasted Demand (Months)
L: Lead Time (Months)
R: Review Period (Months)

Safety Stock (SS) is buffer inventory held to accommodate demand variability or supply delays. The following equation calculates SS as a safety factor, which represents the probability of not stocking out over a given lead time, multiplied by demand variability over the lead time and inventory review period.

$$SS = Z * \sigma_{DL+R}$$

Cycle Stock (CS) is inventory held to meet forecasted demand between inventory replenishments.

$$CS = D * (L + R)$$

Pipeline Inventory (PI) is inventory that is on order but is still in transit or production.

$$PI = D * L$$

Total projected inventory is calculated as

$$Total\ Inventory = SS + CS + PI$$

Ending Inventory is the projected inventory position after all receipts and demands are settled.

$$Ending\ Inventory = Beginning\ Inventory + Receipts - Demand$$

Inventory will also be represented in a euro value that has been calculated from COGS inputs. Tool outputs will be leveraged through Sanofi's IBP process to compare scenarios, make recommendations, and facilitate launch inventory decision-making.

To demonstrate the model's application and evaluate its projected inventory outcomes, we applied it to a case study of tolebrutinib, a Bruton's tyrosine kinase inhibitor that is under development for multiple sclerosis indications and targeted for commercial launch in the United States.

4 Results and Discussion

This section presents the inventory projection outputs of InventEdge applied to the tolebrutinib case study, followed by a discussion of their financial and operational implications. Results are organized around the comparison between the statistical periodic review policy and the MOH heuristic approach across a 20-month post-launch horizon. The section concludes with the limitations of the analysis.

4.1 Results

Using the tolebrutinib case study parameters, we output the following inventory projections:
The transition from MOH-based planning to a statistically grounded approach yielded substantial inventory reductions across the tolebrutinib launch horizon. Applied to a 20-month, post-launch period (January 2027 - August 2028) and evaluated across three commercial demand scenarios—low, base, and high—the statistical periodic review policy projects a material reduction in finished goods inventory requirements relative to the MOH heuristic approach. Figure 6 presents the monthly inventory profile of the InventEdge output, broken down by average cycle stock, safety stock, and pipeline inventory across the simulation horizon.

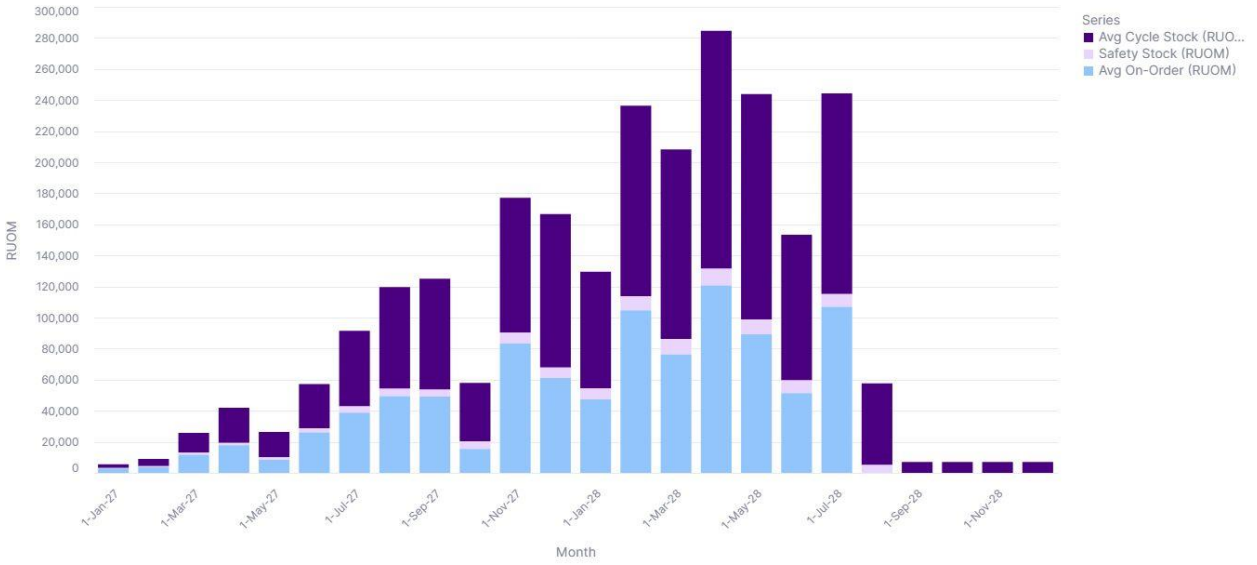

Figure 6: InventEdge Monthly Inventory Projection for Tolebrutinib Launch

Over the comparable 20-month period, the statistical periodic review policy is projected to reduce average total finished goods inventory by 77.7% relative to the MOH heuristic approach—from approximately 2,961,901 UOM to 660,944 UOM per month. Safety stock, which represents the largest driver of excess inventory under the legacy approach, is projected to be reduced by 70% on average, from approximately 2,148,514 UOM to 645,439 UOM, reflecting the shift from arbitrary MOH multiples to a safety stock target anchored in historical MAPE data from analogous products. Table 2 summarizes the key comparative metrics across the two approaches.

Metric	MOH Heuristic	InventEdge
Avg Monthly Total Inventory (UOM)	2,961,901	660,944
Avg Monthly Total Inventory (€)	€4,176,281	€931,932
Avg Safety Stock Target (UOM)	2,148,514	645,439
Peak Total Inventory (UOM)	7,210,590	1,180,187
Total Inventory Reduction	—	77.7%
Safety Stock Reduction	—	70.0%
Peak Inventory Reduction	—	83.6%
Cumulative Capital Avoidance (20 months)	—	€64.9M

Table 2: Summary Comparison for Jan 2027 through Aug 2028

4.2 Implications

The implementation of the tool carries significant implications for Sanofi's strategic and operational planning processes, extending beyond the immediate tolebrutinib case study to inform enterprise-wide inventory management practices.

IBP Transformation

The simulation tool is designed to transform Sanofi's IBP process from a static workflow into a dynamic, interactive decision-making environment. Traditionally, the IBP cycle operates as a monthly cadence in

which commercial teams submit demand forecasts, supply chain teams calculate inventory requirements using MOH heuristics, and finance teams evaluate capital implications in separate, asynchronous workstreams. This sequential approach introduces latency between assumption changes and their downstream impacts, often resulting in decisions based on outdated information. The simulation tool aims to enable real-time strategic alignment by providing instantaneous visibility into the inventory and capital implications of changes to demand assumptions. During IBP sessions, stakeholders can adjust scenario weights, modify service level targets, or update lead time assumptions and immediately observe the resulting inventory projections.

Cross-Functional Decision-Making Enhancement

The tool facilitates cross-functional collaboration by establishing a common analytical foundation that translates high-level data into actionable, accountable targets. By providing a functional inventory breakdown—categorizing safety stock, cycle stock, and pipeline inventory—the tool offers intervention points for different stakeholder groups. For instance, manufacturing teams can focus on lead time reduction to decrease pipeline inventory, demand planning teams can improve forecast accuracy to reduce safety stock requirements, and commercial teams can refine demand assumptions to optimize cycle stock. This shared quantitative framework allows supply chain teams to clearly articulate the trade-offs between service levels and capital investment, while commercial teams can better understand the inventory implications of their demand assumptions. Ultimately, this transparency replaces generalized directives to "reduce inventory" with a shared understanding that reduces the adversarial dynamics often found in IBP reconciliation discussions.

Financial and Operational Benefits

The transition from MOH heuristics to statistical safety stock methodology generates quantifiable financial benefits. For the tolebrutinib launch, the projected target inventory reduction of approximately 77.7% over the first 20 months post-launch represents a significant improvement in working capital efficiency. At a unit value of €1.41 per tablet, the cumulative capital avoidance for tolebrutinib alone is estimated at approximately €64.9M over the time horizon. It should be noted that these figures represent directional planning projections rather than precise predictions of realized post-launch inventory performance. Actual outcomes will depend on how closely post-launch demand aligns with the weighted commercial scenarios and reference product MAPE inputs used in the model.

Beyond capital efficiency, the statistical periodic review policy reduces expiry risk by aligning projected inventory levels more closely with expected demand. Under the MOH heuristic approach, peak projected inventory reached over 7 million UOM in late 2028—a level that, given tolebrutinib's shelf life constraints, would expose a meaningful portion of stock to expiry risk. In comparison, the periodic review policy projects a peak of approximately 1.2 million UOM in July 2028, an 83.6% reduction in peak inventory exposure. This alignment between projected inventory and demand-driven targets substantially reduces the risk of write-offs from near-expiry stock, representing a meaningful reduction in potential inventory obsolescence costs for tolebrutinib alone, with compounding benefits as the methodology is extended to additional launch assets.

4.3 Limitations

While the simulation tool demonstrates significant value in the tolebrutinib case study, several limitations constrain its current applicability and should be considered when interpreting results.

Data Quality and Availability Constraints

The statistical safety stock methodology assumes that a comparable on-market product can serve as a reliable proxy for MAPE. This assumption is necessary because no historical forecast accuracy data exists for a pre-launch drug; a reference product from a similar therapeutic area is therefore the best available approximation of expected demand variability. This assumption may not hold for truly first-in-class drugs with no comparable precedents, where reference product selection requires subjective judgment that could introduce bias, or where the reference product has insufficient historical data to represent the forecast variability of a mature product lifecycle. Beyond reference product selection, MAPE itself carries metric-level limitations. As a ratio, it does not directly estimate the standard deviation of demand required by the safety stock formula, and it can become unstable during periods of near-zero demand that may occur early in a launch before market uptake accelerates. These constraints are acknowledged but do not undermine its selection, given that it is Sanofi's internal standard for measuring forecast error and provides a practical, interpretable approximation of forecast uncertainty suitable for the IBP planning environment.

Model Assumptions and Simplifications

The inventory control formulas applied are based on a Normal-demand distribution assumption (Caplice & Ponce, 2022), which is appropriate for products with aggregated, smooth demand patterns. It may not hold for products with highly skewed demand, such as those driven by seasonality or episodic treatment regimens. The model additionally assumes lead time to be a fixed parameter with known variability. This is reasonable given that pharmaceutical manufacturing lead times are structured and plannable within a monthly IBP cycle; it may not hold when unpredictable disruptions occur, such as regulatory inspections, capacity constraints, or raw material shortages. Finally, the model assumes that theoretically optimal inventory levels are achievable in practice, abstracting away supply-side constraints such as manufacturing capacity limitations or quality release variability. This simplification is appropriate for strategic planning purposes but may not reflect operational reality during early launch periods when production processes are still being optimized.

Case Study and Evaluation Constraints

The model assumes that the tolebrutinib case study parameters—including commercial forecasts, lead times, and service level targets—are representative of a plausible real-world launch scenario. This assumption is appropriate given that all inputs were derived directly from Sanofi's IBP process and reflect the best available planning data at the time of the study. It may not hold once post-launch actuals become available, at which point retrospective comparison would provide a more complete assessment of model accuracy. Results should therefore be interpreted as projections rather than confirmed outcomes.

The comparison with traditional MOH methodology assumes that historical MOH targets would have been applied unchanged throughout the launch period. This is a necessary simplification for establishing a meaningful baseline against which the statistical approach can be evaluated. It may not hold in practice, as experienced planners may have made ad hoc adjustments that partially captured the inefficiencies identified by the statistical approach, potentially narrowing the observed performance gap. Finally, the model assumes that its structure and parameterization are transferable to other products and markets. This is a reasonable starting point for a generalizable planning tool, though quantitative outputs are specific to tolebrutinib's commercial assumptions and supply chain parameters. New input parameters would be required before applying the tool to other drugs.

5 Recommendations

Sanofi's Specialty Care unit faced a systematic risk of inventory over-buffering at new product launches, driven by a reliance on MOH heuristics that failed to account for statistical demand variability. This project developed InventEdge, a deterministic simulation tool that models the expected inventory outcomes of transitioning to a statistical periodic review policy. Applied to the tolebrutinib launch, the tool projects an average total finished goods inventory reduction in excess of 75% relative to MOH heuristics, with corresponding improvements in working capital efficiency and a projected reduction in peak inventory exposure of 83.6%. The following recommendations address how Sanofi should operationalize this transition and extend the tool's capabilities.

5.1 Change Management

The transition from MOH heuristics to statistical safety stock methodology represents a shift in planning philosophy that requires deliberate change management. The Sanofi Specialty Care Global Supply Chain Planning team should align with their Center of Excellence (CoE) team on tool ownership, maintenance, user training, and methodology governance. A Finance team business partner should also be identified and included as a core stakeholder to ensure alignment of capital efficiency objectives. Within Sanofi Specialty Care's Global Supply Chain team, the IBP team already has strong relationships with both these groups, making them well-positioned to lead change management. This step will be instrumental in emphasizing the collaborative nature of the IBP process and the shared accountability for inventory outcomes.

5.2 Future Work

The current simulation tool provides a foundation for continued development and enhancement. The following future work streams are recommended to extend the tool's capabilities and value.

"Time to Recover" Metric

A critical enhancement is the incorporation of supply chain resilience metrics, specifically "Time to Recover" (TTR), which quantifies the duration required to restore target inventory levels following a disruptive event. The current model assumes steady-state replenishment and does not capture the dynamic recovery behavior following supply interruptions, demand spikes, or quality events. Future development should integrate TTR calculations that consider manufacturing capacity constraints, safety stock depletion rates, and expediting options. This enhancement would enable planners to evaluate not only the probability of stock-out but also the duration and severity of potential supply gaps. TTR metrics would be particularly valuable for products with limited manufacturing flexibility or single-source supply chains.

Drug Substance and Drug Product

The current tool's focus on finished goods inventory captures only a portion of the total supply chain capital investment. Future development should extend the simulation to include Drug Substance (DS) and Drug Product (DP) inventory stages, enabling end-to-end inventory simulation. This expansion would require modeling of multi-stage lead times and their interdependencies; incorporation of manufacturing constraints and batch size economics; and integration of quality release variability, which is particularly significant for drug products. An end-to-end model would enable Sanofi to simulate inventory positioning across the supply chain, potentially identifying opportunities to hold buffer inventory at earlier stages where unit costs are lower.

Global Markets

The current tool only operates at the single-country level. Future development should extend the model to support multi-market analysis that incorporates regional demand variation, transfer pricing, duty implications, regulatory constraints on international product movement, and lead time variation for regional distribution. A global capability would enable Sanofi to identify opportunities for inventory pooling and risk sharing across markets. In turn, it would be possible to reduce total global inventory requirements while maintaining regional service levels.

Total Supply Chain Cost

The current simulation tool focuses primarily on inventory value projections, providing visibility into capital investment requirements across safety stock, cycle stock, and pipeline inventory. The tool also provides stock-out cost visibility by quantifying potential lost revenue and profit when the risk-adjusted aggregate forecast falls short of high-demand scenarios. This capability enables stakeholders to understand the financial implications of conservative inventory positioning during IBP sessions.

Future development should expand upon this foundation to incorporate total supply chain cost optimization, including:

- **Purchasing Costs**: Cost per item or total landed cost for acquiring product
- **Ordering Costs**: Setup costs and administrative expenses associated with replenishment
- **Holding Costs**: Storage, insurance, obsolescence, and cost of capital beyond the current inventory value calculations
- **Shortage Costs**: Expansion beyond lost revenue to include expediting expenses, customer satisfaction impacts, market share loss, and long-term brand damage when demand exceeds available inventory
- **Quality Costs**: Write-offs, rework, and disposal costs for expired or damaged inventory

This expanded cost model would enable Sanofi to optimize service levels based on a wider array of cost considerations beyond potential lost revenue alone. For example, the improved tool could identify scenarios where accepting a slightly higher stock-out risk (e.g., 96% vs. 98% service level) would generate net savings when reduced holding costs exceed the expected total cost of occasional stock-outs. Such analysis would be particularly valuable for products with high holding costs relative to stock-out penalties, or conversely, for critical therapies where stock-out costs far exceed inventory carrying costs.

Implementation of total cost optimization would elevate the tool from an inventory projection and revenue risk system into a complete decision support platform for inventory optimization across Sanofi's entire product portfolio.

References

APICS. (2025). *CSCP exam content manual (ECM)* (Version 5.3). APICS.

Bala, B. K., Arshad, F. M., & Noh, K. M. (2017). *System Dynamics*. Springer Texts in Business and Economics. https://doi.org/10.1007/978-981-10-2045-2

Borshchev, A., & Filippov, A. (2004). From System Dynamics and Discrete Event to Practical Agent Based Modeling: Reasons, Techniques, Tools. International Conference of the System Dynamics Society.

Caplice, C., & Ponce, E. (2022). *MITx MicroMasters Program in SCM key concepts*. MIT Center for Transportation & Logistics.

IQVIA. (2025). *Launch Excellence IX: Embracing change and unlocking efficiency* [White paper]. https://www.iqvia.com/locations/emea/library/white-papers/launch-excellence-ix

King, P. L. (2011). Crack the code: Determine optimal safety stock levels using fundamental equations. *APICS Magazine*, 33–36.

Law, A. (2015). *Simulation Modeling and Analysis*. McGraw-Hill Higher Education.

MIT Center for Transportation & Logistics. (2023, November 7). *Applications of simulation in supply chain management* [Video]. YouTube. https://www.youtube.com/watch?v=ZAU4YQMakUY

National Center for Advancing Translational Sciences. (2019). *New therapeutic uses*. National Institutes of Health. https://ncats.nih.gov/research/research-activities/ntu

Sanofi. (2024a). *Annual report 2024*. https://www.sanofi.com

Sanofi. (2024b). *Dupixent® (dupilumab) product information*. https://www.dupixent.com

Sanofi. (2024c). *Full-year 2024 financial results*. https://www.sanofi.com

Silver, E. A., Pyke, D. F., & Thomas, D. J. (2017). *Inventory and Production Management in Supply Chains* (4th ed.). CRC Press.

Tersine, R. J. (1988). *Principles of Inventory and Materials Management* (3rd ed.). North-Holland.

Thiede, S. (2012). *Energy Efficiency in Manufacturing Systems*. Springer Science & Business Media.