SUPPLY CHAIN ANALYSIS

UrbanSpark Logistics GmbH

End-to-End Supply Chain Optimization

Route Optimization · Inventory Policy Redesign · Stockout Risk Projection

PREPARED BY	Rafael Alencar Oliveira
PROJECT	Independent Portfolio Project
DATE	23.06.2026

Contents

Introduction

This document presents the methodology, findings and recommendations of an end-to-end supply chain analysis conducted for UrbanSpark Logistics GmbH, a fictional B2B wholesale distributor of premium e-bike components based in Munich, Germany. The company and its dataset are explicitly fictional, created as the basis for a portfolio-level supply chain project. The dataset was generated with AI assistance (Claude, Anthropic) and designed to reflect realistic B2B distribution dynamics, including intentional data quality issues introduced for analytical realism. All analytical work was conducted in Microsoft Excel, with Log-Hub used for route optimization and Python for stockout risk modelling.

This document is organized as follows. Section 1 introduces the company, its operational context, and the scope of the analysis. Section 2 establishes the as-is baseline for 2024 across transport, inventory, and stockouts. Section 3 covers the methodology behind the three analytical workstreams (ABC-XYZ segmentation, route optimization, and inventory policy redesign), including the Monte Carlo simulation used to project stockout risk. Section 4 presents the to-be results and the projected improvement in each KPI. Section 5 addresses the key assumptions and limitations of the analysis, noting where results should be interpreted with caution. Section 6 closes with three recommendations for operational improvements beyond the current project's scope.

A companion interactive dashboard presenting the full KPI comparison, route optimization results, inventory policy changes and simulation outputs is available at: github.com/ralenc/urbanspark_supply-chain-project.

1. Project context

UrbanSpark Logistics GmbH is a fictional B2B wholesale distributor of premium e-bike components, built as the basis for this portfolio-level supply chain analysis. It operates a central warehouse in Munich and serves 47 dealer partners across five delivery zones in Southern Germany: Munich, Stuttgart, Nuremberg, Frankfurt and Augsburg. Its portfolio spans 38 active SKUs across six categories (drivetrain, electrics, brakes, wheels, suspension and accessories), sourced from eight tier-1 suppliers.

The underlying dataset was generated with AI assistance (Claude, Anthropic) and is explicitly fictional. Customer names, order volumes, cost figures and operational parameters are all simulated, built to reflect realistic B2B distribution dynamics and including a few intentional data quality issues to keep the analysis grounded. The analysis itself was done in Microsoft Excel, with Log-Hub handling route optimization and Python, also AI-assisted, used for the stockout simulation.

The project covers three interconnected workstreams: route optimization, inventory policy redesign, and stockout risk projection. Each produces its own before/after comparison, and together they roll up into a single KPI baseline.

2. As-Is situation

The as-is analysis established a 2024 baseline across three cost dimensions: transport, inventory and stockouts. Total annual logistics cost was recorded at €482,259, broken down as follows: distribution costs €269,632, inbound freight €4,836, express/emergency shipments €5,160, annual inventory holding cost €6,055, warehouse storage €188,700, estimated ordering cost €5,852, and emergency order cost €1,327.

Transport

The six regular delivery routes were operated using a mixed model: Munich inner and outer ring routes served by own Sprinter fleet, and the four regional zones (Stuttgart, Nuremberg, Frankfurt, Augsburg) contracted to DPD as a 3PL carrier.

A data quality issue was identified during the analysis: as-is route distances in the operational records had been entered as one-way estimates rather than round trips, causing the reported distribution cost to be significantly understated. Distances were corrected by doubling the recorded figures prior to any comparison. It should be noted that for DPD-operated routes specifically, the return leg doesn't actually happen: DPD drivers don't come back to the UrbanSpark warehouse after finishing a delivery run. However, the warehouse was included as the final stop in the Log-Hub input for all routes to ensure a consistent round-trip distance basis for comparison. The round-trip assumption for DPD routes is therefore a modelling convention rather than an operational reality and is documented as such.

Inventory

Stock was held across all 38 SKUs with a total observed tied capital of €33,641 at purchase cost. Safety stock and reorder points had not been systematically reviewed and showed no consistent relationship to demand variability or revenue tier. High-variability items were under-protected and stable high-volume items over-stocked. Reorder quantities appeared to be set arbitrarily, with no evidence of EOQ-based calculation. Annual ordering cost was estimated at €5,852 based on 266 purchase orders across all suppliers. Inventory turnover stood at 7.5x, with an average days sales of inventory (DSI) of 48.7 days.

Stockouts

A total of 33 stockout events were recorded across 2024, affecting 16 distinct SKUs. Emergency order costs resulting from these events totaled €1,327. Notably, most stockout-prone SKUs were Z-classified items (those with high demand variability) regardless of their revenue tier, pointing to the absence of any variability-based safety stock logic in the current policy.

3. Methodology

The project followed a sequential three-workstream structure: route optimization, inventory policy redesign, and stockout risk projection. Each workstream built directly on the outputs of the preceding one. The ABC-XYZ segmentation produced in the inventory preparation phase informed the service level targets used in the safety stock calculation, and the final inventory parameters fed into the Monte Carlo simulation. All calculations were performed in Microsoft Excel unless otherwise stated.

ABC-XYZ Analysis

The portfolio was segmented using a combined ABC-XYZ framework. ABC classification was based on annual revenue contribution, calculated as total units sold in 2024 multiplied by unit selling price, ranked in descending order and expressed as a cumulative percentage. Thresholds were set at 80% for A items, 95% for B items and 100% for C items, yielding 24 A-tier, 9 B-tier and 5 C-tier SKUs. The relatively flat revenue curve (where 24 SKUs were needed to reach 80% of revenue) was noted as a finding in itself, indicating a more evenly distributed portfolio than a classic Pareto structure would suggest.

XYZ classification was applied on top of the ABC layer using the coefficient of variation (CoV = standard deviation / mean) of monthly demand across the 12 months of 2024. Thresholds were set at CoV ≤ 25% (X, stable demand), 25–50% (Y, moderate variability) and > 50% (Z, erratic demand), in line with standard ABC-XYZ methodology. The final ABC-XYZ classification was derived by combining the two single-letter codes (e.g. AZ, BY, CX), producing nine active segments across the 38-SKU portfolio.

Route Optimization

Route optimization was carried out using the Log-Hub Last Mile Lat/Lon tool. GPS coordinates were assigned to all 47 customer locations and the central Munich warehouse (48.199858°N, 11.577069°E). All five delivery zones were submitted in a single calculation using distinct Path IDs, with the warehouse set as both the departure and final stop for every route to ensure consistent round-trip distance output.

The initial Log-Hub run combined all 15 Munich customers into a single route. This result was rejected on operational grounds: the combined average load of approximately 930kg exceeded the Sprinter's 800kg capacity limit. A second targeted calculation was run, splitting the 15 Munich customers into two geographically balanced groups: a North/West cluster (8 customers) and a South/East cluster (9 customers). That allowed Log-Hub to optimize stop sequencing within each group while respecting the vehicle capacity constraint. The capacity validation was performed manually outside of Log-Hub, as the free-tier tool does not support weight constraints natively.

As-is route distances had been recorded as one-way estimates. These were doubled prior to comparison to produce a consistent round-trip basis against Log-Hub's round-trip output. The same carrier rates were applied to optimized distances as to the as-is figures (€2.39/km and €1.82/km for own-fleet Munich routes, and €1.25/km plus an €18 per-run handling fee for DPD routes), so that reported savings reflect distance reduction only, with no rate adjustment assumed. For DPD routes, the contracted rate is purely variable, meaning savings scale directly and proportionally with distance reduction and are fully reliable. For own-fleet Munich routes, the €2.39/km and €1.82/km rates are blended figures that include both variable components (e.g. fuel, tire wear) and fixed components (e.g. driver salary, vehicle depreciation, insurance). Fixed costs do not decrease when a route gets shorter, rather they spread across fewer kilometers, meaning the true effective cost per km on a shorter route is higher than on the original route. Applying the same blended rate to a shorter distance therefore understates the true to-be cost and consequently overstates the saving. The reported Munich route savings should be interpreted as an upper-bound estimate; the actual saving depends on the fixed-to-variable cost split within the blended rate, data for which was not available within the scope of this project.

Inventory Policy Redesign

The inventory policy redesign applied a differentiated (s, Q) framework, continuous review with a fixed reorder point and fixed order quantity, across all 38 SKUs. This matched the existing as-is ordering practice, so the same policy structure was retained as the baseline model. Three parameters were recalculated for every SKU: safety stock, reorder point, and economic order quantity.

Service level targets were differentiated by ABC tier, reflecting the revenue impact of a stockout on each segment: A-tier SKUs were assigned a 98% service level (Z = 2.05), B-tier 95% (Z = 1.65) and C-tier 90% (Z = 1.28).

Safety stock was calculated using the standard formula: `SS = Z × σ_monthly × √(LT/30)`. Where Z is the service level factor, σ_monthly is the standard deviation of monthly demand derived from 2024 order history, and LT is the supplier lead time in days. The LT/30 term converts the lead time into a fraction of the monthly demand window over which the standard deviation was measured, applying the square-root-of-time rule to scale demand variability to the lead time period. In other words, the standard deviation used in this formula (σ_monthly) measures how much demand varies over a 30-day period, since it was calculated from monthly order data. However, what matters for safety stock is not how much demand varies over a generic month, but how much it varies specifically during the lead time. Because variance is additive across independent time periods, standard deviation scales with the square root of time. The term √(LT/30) therefore converts σ_monthly into the equivalent standard deviation over the lead time. If the lead time is shorter than 30 days, demand has less time to deviate and the safety stock requirement is smaller; if it is longer, more uncertainty accumulates and a larger buffer is needed. Results were rounded up to the nearest whole unit, since safety stock represents a physical buffer and rounding down would marginally under-protect against the variability it is designed to cover.

Reorder point was calculated as: `ROP = ⌈(Avg_Monthly / 30 × LT) + SS⌉`. This represents the expected demand during the lead time plus the safety buffer. The same upward rounding convention was applied.

Economic order quantity was calculated using the Wilson formula: `EOQ = √((2 × D × S) / (H × C))`. Where D is annual demand (units), S is the fixed ordering cost per purchase order (€22), H is the annual holding cost rate (18%) and C is the unit purchase cost. Results were rounded to the nearest whole unit. Unlike safety stock and ROP, EOQ is a cost-minimization point rather than a safety buffer, so there is no directional reason to favor rounding up over rounding down.

Seasonal reorder points were calculated separately for SKU-015 (Schwalbe Smart Sam tire) and SKU-016 (Schwalbe Marathon Plus tire), whose demand follows a pronounced seasonal pattern. The year was split into a high season (February–August, 7 months) and a low season (September–January, 5 months), with the February–August window chosen to capture pre-season dealer stocking behavior beginning in February ahead of the main riding season. Seasonal average monthly demand was calculated manually from the 2024 order history for each period, and two separate ROPs were derived (one per season) using the same formula as above with the seasonal average substituted for the annual average. Safety stock was held constant across both seasons.

The (s, Q) policy assumption implies continuous review: stock is monitored in real time, and a reorder is triggered the instant it reaches the reorder point. In practice, UrbanSpark's ordering cadence is partly driven by supplier shipment schedules, which introduces a periodic element the model doesn't capture. This gap is acknowledged as a limitation and addressed in the recommendations section.

Monte Carlo Simulation

Since the relationship between updated reorder parameters and stockout frequency cannot be derived analytically, a Monte Carlo simulation was implemented in Python to project the impact of the redesigned inventory policy on stockout events.

For each of the 38 SKUs, 2,000 independent simulated years were run under both the as-is policy (old ROP and reorder quantity) and the to-be policy (new ROP and EOQ). Each simulated year was constructed as follows: twelve monthly demand values were drawn with replacement from that SKU's actual 2024 monthly order history, a bootstrap resampling technique where months are drawn randomly from the real data and the same month can appear more than once. This preserves each item's actual seasonality and variability without assuming a theoretical distribution. Each monthly value was then broken down into 30 daily demand observations via a Poisson process (λ = monthly demand / 30), introducing realistic day-to-day variation while keeping the monthly total consistent on average. Inventory levels were tracked day by day under the (s, Q) policy: when stock reached the reorder point and no order was already in transit, an order of size Q was placed and arrived after the SKU's lead time. A stockout episode was recorded whenever daily demand exceeded available stock, with consecutive stockout days counted as a single episode. A 30-day burn-in period at the start of each simulated year was excluded from the count to allow the system to reach a representative operating state before measurement began.

4. To-Be Results

Transport

Route optimization produced measurable distance reductions on four of the six regular delivery routes. The most significant improvements were on the Munich runs, where Log-Hub's optimized sequencing cut the combined round-trip distance from 320km to 81.79km across the two routes, driven mainly by more efficient stop ordering within each geographic cluster. Stuttgart and Augsburg also showed meaningful improvements (134.64km and 117.41km saved respectively). Nuremberg was essentially unchanged (4.56km reduction), confirming that the existing manual planning for that zone was already near-optimal. Frankfurt was the one route where optimization produced a marginal deterioration (+11.04km), which comes down to the geographic spread of Frankfurt customers extending as far south as Darmstadt: a genuinely dispersed zone where no meaningfully shorter sequence exists.

Total annual distribution cost decreased from **€269,632** to **€167,934**, a saving of €101,698 per year. Inbound freight (€4,836) and express/emergency costs (€5,160) were held unchanged in the to-be scenario: inbound consolidation was outside the scope of this project, and express costs are expected to decrease as a result of the inventory policy improvements, though precisely quantifying these savings would require observed post-implementation data. Total annual transport cost therefore decreased from **€279,628** to **€177,930**.

Inventory

The redesigned (s, Q) policy produced differentiated changes across the portfolio: some SKUs saw safety stock and reorder points increase (under-protected high-variability items such as SKU-006, SKU-017 and SKU-022), while others decreased (over-stocked items such as SKU-001 and SKU-014). These opposing movements largely offset each other in aggregate, resulting in a modest net reduction in theoretical tied capital from **€37,506** to **€36,709** (−€797). Annual holding cost decreased correspondingly from **€6,751** to **€6,608** (−€143), and annual ordering cost decreased from **€5,852** to **€4,664** (−€1,188), reflecting a reduction from 266 to 212 purchase orders per year as EOQ-based quantities better matched the economic batch size for each SKU.

Inventory turnover improved marginally from 6.7x to 6.9x, and DSI from 54.3 to 53.1 days. These figures are derived from the theoretical average tied capital (Q/2 + SS) applied consistently to both columns, which, as discussed in the limitations section, differs from the observed snapshot value of €33,641 recorded in the inventory records. These modest movements are expected, since the redesign's primary objective was reallocating inventory more intelligently rather than reducing the total outright: holding more where variability is high and revenue impact is significant, and less where demand is stable and stock was excessive. The more meaningful improvement is captured in the stockout projection.

Stockout Risk

The Monte Carlo simulation produced a mean of **19.09** stockout episodes per year under the as-is policy and **4.97** under the to-be policy, representing a relative reduction of **74%**. This was corroborated by a near-identical reduction in total stockout days (23.13 to 6.26, −73%), providing a consistency check across two independent counting methods.

Since the simulated as-is figure (19.09) understates the observed 2024 baseline of 33 events, due to the model simplifications described in the Limitations and assumptions section, the relative reduction was applied to the observed baseline rather than the simulation's own absolute output: 33 × (1 − 0.74) ≈ 8.6, rounded to **9** projected events per year. Emergency order cost and SKUs affected were scaled using the same 9/33 ratio for internal consistency, giving projected values of **€362/year** and **4** affected SKUs respectively.

The SKUs contributing most to the simulated as-is stockout count were almost exclusively Z-classified items: SKU-018 (BZ, 4.34 episodes/year), SKU-022 (AZ, 4.02), SKU-017 (AZ, 3.88) and SKU-007 (AZ, 1.12), regardless of their revenue tier. These same SKUs show the largest improvements under the new policy, with SKU-017 dropping from 3.88 to 0.08 episodes per year. This validates the core premise of the ABC-XYZ-differentiated approach: stockout risk is driven by demand variability, not revenue tier alone, and a policy that ignores the XYZ dimension will systematically under-protect the items most exposed to demand spikes.

Total Annual Logistics Cost

Combining all workstreams, total annual logistics cost decreased from **€482,259** to **€378,263**, a projected reduction of **€103,996** per year (−21.6%). The breakdown by component is as follows:

Cost component	As-is €	To-be €	Δ €
Distribution	269,632	167,934	−101,698
Inbound freight	4,836	4,836	0
Express / emergency transport	5,160	5,160	0
Holding cost	6,751	6,608	−143
Warehouse storage	188,700	188,700	0
Ordering cost	5,852	4,664	−1,188
Emergency order cost	1,327	362	−965
Total	**482,259**	**378,263**	**−103,996**

The dominant saving is in distribution (**97.8%** of the total reduction), driven entirely by route optimization. Inventory policy improvements contribute the remaining **2.2%,** primarily through reduced ordering frequency. Warehouse storage cost is unchanged as no change to the physical warehouse footprint was modelled.

5. Limitations and assumptions

Route Optimization

As-is route distances were recorded as one-way estimates in the operational data. These were doubled prior to comparison to produce a consistent round-trip basis against Log-Hub's round-trip output. This correction is documented explicitly, and the original figures are preserved in the Current_Routes sheet. The doubling assumption implies that the as-is dispatcher's distance estimates reflected a single-direction journey, which is consistent with the Google Maps validation performed during the analysis (Munich–Frankfurt: 392km one-way vs. 410km recorded, accounting for urban customer stop detours).

For own-fleet Munich routes, the blended per-km rates (€2.39 and €1.82) include both variable and fixed cost components. Since fixed costs do not scale with distance, applying the same rate to a shorter optimized route understates the true to-be cost and consequently overstates the saving. The reported Munich route savings should therefore be interpreted as an upper-bound estimate. The true saving depends on the fixed-to-variable cost split within the blended rate, data for which was not available within the scope of this project. This limitation does not apply to DPD routes, where the contracted rate is purely variable.

Log-Hub's free-tier Last Mile tool does not support vehicle capacity constraints natively. The capacity validation for Munich routes (combined load vs. 800kg Sprinter limit) was performed manually outside the tool, and the two-route Munich split was applied as a modelling decision rather than an algorithmically enforced constraint.

Inventory Policy

The as-is tied capital figure used in the KPI comparison (€37,506) is a theoretical average derived from the formula (Q/2 + SS) applied to current reorder quantities and safety stocks across all 38 SKUs. This differs from the observed snapshot value of €33,641 recorded in the Inventory_Costs sheet, which reflects the actual stock value on a specific date. The theoretical average was used in both the as-is and to-be columns to ensure a like-for-like policy comparison: two modelled steady-state averages, rather than an observed snapshot mixed with a modelled projection. Inventory turnover ratio and DSI were derived from the theoretical figure in both columns for the same reason.

The analysis applies a uniform (s, Q) policy, continuous review with fixed reorder point and fixed order quantity, consistently across all 38 SKUs, reflecting current practice. For Z-classified SKUs with highly erratic demand, an (s, S) order-up-to policy may offer additional robustness, as it adapts the order quantity to the actual shortfall below the reorder point rather than always ordering a fixed amount. This refinement was beyond the scope of the current analysis but is noted as a potential improvement, particularly for SKU-007 (AZ), SKU-017 (AZ), SKU-018 (BZ) and SKU-022 (AZ).

The (s, Q) model assumes continuous review: stock is monitored in real time, and a reorder fires instantly when the reorder point is reached. In practice, UrbanSpark's ordering cadence is partly driven by supplier shipment schedules, which introduces a periodic element the model doesn't capture. The gap between continuous-review theory and periodic-shipment reality means actual reorder timing may lag behind the modelled trigger point, slightly increasing effective exposure relative to the calculated safety stock.

The seasonal reorder points calculated for SKU-015 and SKU-016 are based on a single year of demand data (2024). A longer demand history would produce more stable seasonal averages and reduce the risk of the seasonal split being influenced by anomalous months within the observation window.

Monte Carlo Simulation

The simulation draws demand from 12 months of historical data via bootstrap resampling. This limits the model's ability to capture tail demand events (unusually high or low demand occurrences) that didn't occur within the observation window, a known constraint of bootstrap methods applied to short time series. Supplier lead times are modelled as fixed and deterministic; real-world delivery delays aren't represented, so the simulation underestimates true stockout risk under both policies. The projected to-be figure of 9 events per year is therefore a calibrated estimate derived from the simulation's relative reduction applied to the observed baseline, not a guarantee of post-implementation performance. Actual stockout frequency under the new policy should be monitored and compared against this projection once the policy is live.

6. Recommendations

EDI Integration with Key Suppliers

The most impactful operational improvement available to UrbanSpark beyond this project's scope is the integration of an Electronic Data Interchange (EDI) system with its primary suppliers. Under the current model, purchase orders are placed manually, introducing administrative delays between the moment stock reaches the reorder point and the moment the supplier receives the order. This gap effectively extends the functional lead time beyond the supplier's stated production and transit time, eroding the safety stock buffer calculated in this analysis.

EDI integration would enable automated, real-time PO transmission the instant stock hits the reorder point, closing the administrative delay and bringing operational practice closer to the continuous-review assumption providing a basis for the (s, Q) model. A phased rollout is recommended, beginning with the highest-risk, longest-lead-time suppliers: Bosch (21-day lead time, AZ-classified motor) and SRAM/Fox (25-day lead time, AZ/BZ-classified suspension components). These are the items where the gap between modelled and actual reorder timing has the greatest potential impact on stockout risk.

It should be noted that EDI reduces the administrative portion of lead time, not the production or transit portion. Bosch's motor will still take however long it takes to manufacture and ship; EDI simply ensures the order is placed without delay the moment the trigger is reached.

Seasonal ROPs for Tire SKUs

SKU-015 (Schwalbe Smart Sam) and SKU-016 (Schwalbe Marathon Plus) exhibit a pronounced seasonal demand pattern, with average monthly demand roughly twice as high during the high season (February–August) as during the low season (September–January). The annual-average ROP calculated in the main inventory policy underestimates the required buffer during peak months and overestimates it during the off-season. Implementing the two seasonal ROPs derived in this analysis would better align stock levels with actual demand timing and reduce both stockout risk in peak months and excess holding cost in quiet months.

(s, S) Policy for Z-classified SKUs

For the four highest-variability SKUs (SKU-007 (AZ), SKU-017 (AZ), SKU-022 (AZ) and SKU-023 (AZ)) consideration should be given to replacing the fixed (s, Q) policy with an order-up-to (s, S) policy. Under (s, S), when stock falls to the reorder point, the order quantity is not fixed but sized to bring stock back up to a target level S, adapting to how far below the reorder point stock has fallen. For items with erratic, lumpy demand — where a single large order from a customer can push stock well below the reorder point in one step — this adaptive sizing provides a more robust recovery mechanism than a fixed order quantity, which may not be sufficient to restore an adequate buffer in a single replenishment cycle.

Post-implementation Monitoring

The to-be KPI values presented in this analysis are projections derived from models and simulations, not observed outcomes. It's recommended that actual performance be tracked against the projected figures, particularly the stockout frequency target of approximately 9 events per year, once the new policy parameters are implemented. A quarterly review of ROP, safety stock and EOQ values is also recommended, since these are calculated from 2024 demand data and will drift from optimal as demand patterns evolve. The data quality issue identified during this project (one-way distance entries, unadjusted reorder parameters) underlines the risk of allowing operational parameters to go unreviewed for extended periods.