



# Continuous Micro-Optimisation in Retail Operations

---

## Optimising the **Reality** Behind the Plan

Consider **a large omnichannel retailer** operating hundreds of stores alongside one or more e-commerce fulfilment centres.

**Its core systems are familiar:** point-of-sale captures transactions in real time; replenishment systems generate store and DC orders daily or weekly; labour-management tools produce rosters days in advance; and pricing engines define promotions and markdowns.

Together, these systems create **a coherent operating plan.**

They do not, however, manage the minute-by-minute reality of execution.

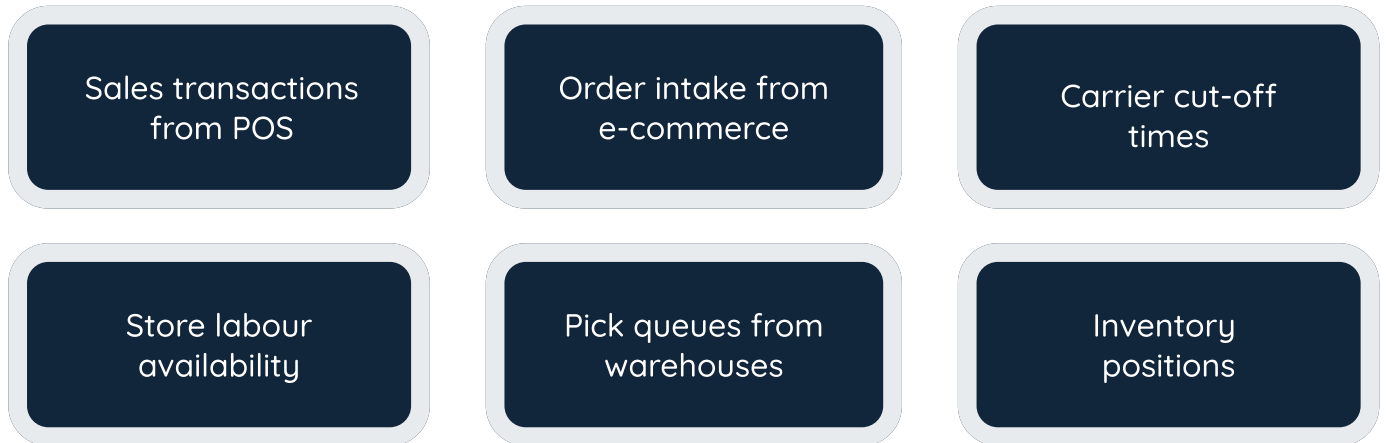
This is where “**Micro-Optimisation**” steps in.

It is not looking to replace this technology stack; instead, it sits alongside it, **acting as a decision** layer that continuously interprets live data and recommends the next best action within the boundaries set by existing plans.



## How does this work ?

Micro-Optimisation is running a series of optimisation models, **ingesting a constant stream of data:**



Any one of these might be diverging from the plan that was originally devised – this might not be “too much of a problem”, but with the power of the optimisation model, a realignment

can be readily implemented to **get the shift back on track.**

This is not a wholesale replanning process, more a fine readjustment.

### Example 1 - Allocating scarce stock

For example, when an online order arrives for a fast-selling item with only a handful of units left in the region, the optimiser evaluates **a small choice set:**

- fulfil from a nearby store,
- ship from a regional DC,
- delay briefly for consolidation,
- or if margins are thin hold stock for higher-value demand.

**The objective function is explicit:** maximise expected contribution margin subject to delivery promises, labour capacity and inventory-protection rules.



**-The result** is a recommendation in seconds, not a revised network plan-

## Example 2 - Responding to changes in footfall

A similar logic applies within the store. As footfall rises unexpectedly in the afternoon, the optimiser may identify that the next 30 minutes represent a risk window for shelf availability in a high-conversion category.

Rather than triggering an emergency replenishment order, **it may recommend a micro-decision**: divert one staff member from backroom tasks to replenish a specific aisle now, because the marginal sales impact outweighs the labour cost. The labour-management system remains intact; the optimiser merely proposes **a short-lived task reallocation**.

Replenishment planning, too, benefits from this approach. Traditional systems generate orders at fixed cadences, assuming that deviations can wait.

**Micro-optimisation fills the gap between cycles.** It may detect that a store is likely to stock out of a promoted item within the next two hours and recommend a cross-store transfer or a partial release from safety stock already on site. The decision is local, reversible and constrained by policy, yet it **materially improves on-plan execution**.



Instead of triggering full replanning, micro-decisions **adjust execution in real time** — one shelf, one shift, one action at a time



## Example 3 - Pricing & Promotions

Pricing and promotions are affected in subtler ways. Instead of redesigning offers, the optimiser can throttle or prioritise their exposure.



If a promotion is driving demand faster than anticipated in a particular catchment area, the system may temporarily suppress its visibility online or in-app for that location, protecting availability for higher-margin channels elsewhere. Again, **the pricing system defines the rules**; optimisation governs their moment-by-moment application.

# Why **Micro-Optimisation** works in the retail environment?

There are five characteristics, we believe, for micro-optimisation to be of real value:



Retail operations provide multiple challenges which are ideally suited to micro-optimisation. What distinguishes this model is not its ambition but its restraint. Each optimisation problem is small: dozens of orders, a few locations, a narrow time window. Each of these represents real decisions that local teams face on a daily basis, often having

to rely on intuition and guesswork to solve them and respects the constraints imposed by merchandising, replenishment and labour plans.

And each is designed to be forgotten once executed, replaced immediately by the next decision.



Over time, **the cumulative effect is significant.** Stockouts are avoided not by holding more inventory, but by allocating it more intelligently. Labour productivity improves not through wholesale re-rostering, but through timely micro-adjustments.

Customer promises become more reliable because they are **constantly recalculated against reality**, not against yesterday's assumptions.

## Continuous Optimisation, One Decision at a Time

In this sense, micro-optimisation turns retail execution into a continuously **self-correcting system**. Planning remains essential, but advantage accrues to those retailers that can also optimise the gaps between plans — one transaction, one shelf, and one decision at a time.



**Jonathan Nicholas**  
Managing Director,  
Customer Success, Sophus



**Katerina Petta**  
Network Design Capability  
Builder, The Middle