

**ANT USA Inc.**

---

Retail Merchandise Planning Software

## **The 3Q Framework™**

### **Three Questions. One Decision Discipline**

**What, Why, and How:**

A Universal Decision Framework as applied to Planning and Replenishment

A White Paper for Retail Planning and Merchandising Leadership

By Dmitry Goykhman

ANT USA Inc

May 2026

# Executive Summary

Most retail planning organizations are drowning in data and starving for clarity. They know something is off — a category is underperforming, a vendor is late, open-to-buy is being spent in the wrong places — but by the time the reporting cycle surfaces it, the window for action has closed. They weren't told. They had to go looking.

This paper introduces the **3Q Framework** — a three-layer decision model built on three questions every retail merchant needs to answer every week: **What** is demanding my attention right now? **Why** is it happening? And **How** do I act on it?

The Framework is not a new invention. It is a disciplined articulation of how the best retail planners already think. What the Framework does is make that thinking systematic, repeatable, and software-supported — so the discipline belongs to the process, not just to the individuals who happen to have it.

The paper covers three things. First, it explains why planning and replenishment are fundamentally different processes that require different decision logic — and why they complement each other when that distinction is respected. Second, it defines each layer of the 3Q Framework precisely: What (exception alerting and signal prioritization), Why (multi-factor root-cause analysis), and How (recommended actions that address both problems and opportunities). Third, it illustrates the framework with a worked example using three competing buy scenarios, showing how 3Q changes which decisions get made and why that matters.

## The 3Q Framework

The 3Q Framework organizes the merchant's decision process into three layers. Each layer answers one question. Each layer depends on the previous one. And the sequence matters — acting without diagnosing, or diagnosing without first knowing where to look, is how organizations stay busy without making progress.

The three layers are:

Layer	Question	What It Does
<b>WHAT</b>	<i>Where do I need to look right now?</i>	Exception alerting and signal prioritization — a curated, automatically delivered set of conditions that demand a merchant's attention today, both problems and opportunities.
<b>WHY</b>	<i>What is causing this condition?</i>	Multi-factor root-cause analysis — the relevant metrics and diagnostics loaded automatically once an exception is surfaced, showing what combination of variables is driving the flagged condition.
<b>HOW</b>	<i>What should I do about it?</i>	Recommended actions — specific, system-suggested steps the merchant can take to address the diagnosed condition, drawn from the planning and analytics engine's output.

# 1. Different Problems That Belong Together

---

Retail inventory management is not one process. It is AT Least three, running in parallel, governed by different logic — and the confusion between them is one of the most persistent sources of planning dysfunction in the industry. Planning, Allocation, and Replenishment form an analytics-decisions-execution loop. In this document we will address the complementary tension between Planning and Replenishment.

## 1.1 Replenishment: the closed loop

Replenishment answers a narrow, well-defined question: given what is selling today, how much of it do I need on the shelf tomorrow? The constraint is availability. The decision cycle is short. The information is largely current. And critically, a wrong replenishment decision is reversible — next week's order corrects this week's mistake.

This makes replenishment well-suited to rules-based automation. Min/max thresholds, rate-of-sale triggers, weeks cover or floorplan targets — these translate directly into system-generated replenishment proposals. The merchant's job is to set the parameters, review the exceptions, and adjust the signal when the underlying assumptions shift. The system (ideally) does the rest. Of course, any system that relies on millions of recurring decisions made on a regular basis is going to deliver imperfect results, but these will be far more productive than manual cycles in most cases.

## 1.2 Planning: the open problem

Merchandise Planning (often referred to as MFP or Inventory Planning) answers a much harder question: given what I know today about trends, vendor lead times, margin targets, open-to-buy constraints, promotional calendars, economy, store openings/closings, merchandising team abilities, weather patterns and seasonal timing, what should I commit to buying over the next three to six months? The decisions are large, made in advance, and effectively irreversible. A wrong forward buy doesn't get corrected next week. It shows up as excess inventory at season end, markdown pressure you didn't budget for, and capital tied up that can't be redeployed.

Planning involves many more variables than replenishment. Discount and markdown cadence, vendor receipt timing, full-price versus markdown mix, sales trend by price point, aged inventory drag, and dozens of other considerations — all these factors interact. There is no single correct variable to optimize. The planner's job is to make large, forward-looking commitments under uncertainty, prioritize where to put scarce open-to-buy dollars, and do it in a way that is defensible when the season closes.

## 1.3 Why they belong together

These two processes are complementary because they operate at different time horizons on the same inventory. Planning SHOULD set the top-down parameters that replenishment executes against. Replenishment generates the sell-through signals that planning needs to validate or revise its forward assumptions. A planning system that ignores replenishment and sales trending data is often flying blind on in-season performance. A replenishment system that ignores the planning context will over-replenish slow categories and under-replenish the ones the plan was built around.

The 3Q Framework applies to both. The questions — What needs attention, Why is it happening, How do I act — are equally valid whether the merchant is reviewing a replenishment exception at 9am or revising a forward buy at the category level. The difference is in the time horizon of the answer, not in the structure of the question.

### The core tension

Replenishment optimizes for availability in the short term. Planning optimizes for capital efficiency over a season. Both require prioritization. Both require diagnosis. Both require action. The 3Q Framework gives merchants a single operating model that works across both.

## 2. The WHAT Layer — Knowing Where to Look

---

The hardest part of planning at scale is not the analysis. It's knowing which categories, stores, or channels need attention today — before the weekly reporting cycle forces the issue.

### 2.1 The failure mode: the full-report trap

Standard reporting gives merchants everything. That is the problem. A weekly category summary with fifty rows, a store performance report with three hundred locations, a vendor receipt tracker with hundreds of open POs — none of these tells you where to start. The merchant who works hardest reads everything and still misses the products that quietly tipped from 82% sell-through to 61% in the last ten days.

Exception reporting addresses this by surfacing only the conditions that cross a defined threshold. But static exception reports have their own failure mode: they are pulled on demand, reviewed on a schedule, and by definition delivered after the fact. The merchant who runs the exception report on Thursday is looking at conditions that may have warranted action on Monday.

### 2.2 Proactive Alerting: bringing the signal to the merchant

The WHAT layer of the 3Q Framework is built on a more active model. Rather than requiring the merchant to go looking, the system should come to the merchant. Conditions that cross a defined threshold — a class tracking below plan by more than 8%, a location over-inventoried relative to its cover target, a vendor receipt delayed past its buffer date — trigger an alert that is automatically delivered to the merchant's workflow.

This is what distinguishes exception alerting from exception reporting. Reporting is a pull mechanism: information is available when requested. Alerting is a push mechanism: information arrives when it matters. The difference in practice is the difference between knowing about a problem and knowing about it in time to do something about it.

Two categories of WHAT conditions deserve equal weight — and in practice, the opportunity side is underserved:

Condition Type	Examples	The Risk of Missing It
<b>Problem signals (remediate and avoid repetition)</b>	Sell-through below plan; inventory above cover target; vendor lateness past buffer threshold; aged inventory accumulating; plan variance widening week over week	Margin erosion, markdown pressure, trapped capital, season-end distress
<b>Opportunity signals (replicate and reinforce)</b>	Category outperforming plan; sell-through rate accelerating; a location running ahead of receipt schedule with capacity to absorb more; a vendor delivering consistently early	Captured upside evaporates. The window to accelerate replenishment, protect OTB, or double down on a trend closes quietly.

The WHAT layer does not diagnose and it does not prescribe. It curates. Its output is a short, prioritized list of conditions — problems and opportunities — that are demanding a merchant’s attention right now. Everything else is subordinated to that list.

**Keeping alerts credible**

The value of the WHAT layer depends on its credibility. An alert system that fires too often, on conditions that don’t require action, trains merchants to ignore it. Threshold calibration — setting the sensitivity for each exception type against each metric — is not a technical task. It is a planning judgment. Getting it right is the precondition for the rest of the framework working.

## 3. The WHY Layer — Diagnosing the Condition

---

A flag without context is noise. Once the WHAT layer surfaces a condition, the merchant needs to understand what combination of variables is driving it before deciding what to do. That is the WHY layer.

### 3.1 Why planning problems involve many variables

Planning performance is almost never attributable to a single variable. A category tracking below plan might be a demand problem (the sell-through rate is down), a receipt problem (inventory hasn't arrived in time for the selling window), a mix problem (the wrong price points or colorways got the weight), a markdown problem (the cadence is too late or too shallow), or some combination of all four. The merchant who acts on one variable without diagnosing the others will fix the symptom and leave the cause intact.

The WHY layer loads the relevant metrics automatically once an exception is surfaced. The merchant doesn't start from a blank spreadsheet. They start with the right questions already framed and the metrics already loaded: sell-through by price point, weeks of supply, receipt timing versus plan, plan variance by location, full-price versus markdown mix.

### 3.2 Good scenarios and bad scenarios both need WHY analysis

The diagnostic obligation applies to both sides of performance. Most organizations are disciplined about diagnosing problems. Far fewer are equally disciplined about diagnosing successes.

When a category outperforms plan, the natural response is satisfaction. The better response is investigation: why did this work? Was it the price point? The timing? The vendor relationship? The fact that a competitor was out of stock? Understanding the cause of a good outcome is the precondition for replicating it. An outperformance that isn't diagnosed is just luck. An outperformance that is diagnosed and understood becomes a repeatable advantage.

The WHY layer therefore serves two functions that look symmetric but operate differently:

- For problem conditions: identify the root cause so the How layer can surface a targeted remediation rather than a generic response.
- For opportunity conditions: identify the success driver so the How layer can surface actions that replicate and reinforce the upside rather than simply acknowledging it.

### 3.3 The variables that matter in forward planning

The following set of variables are one example of primary diagnostic inputs for forward planning WHY analysis. They are not competing objectives. They are lenses on the same underlying performance, each illuminating a different dimension:

Variable	What it explains in a WHY diagnosis
<b>Full-price vs. markdown mix</b>	Blended realized revenue — the single biggest driver of whether planned margin holds
<b>Sell-through rate by price point</b>	Demand signal at the most granular level — reveals whether the problem is product, price, or timing
<b>Vendor receipt timing</b>	Capital tie-up and selling window risk — late receipts miss full-price weeks and force markdown decisions
<b>Markdown depth and cadence</b>	Rate at which trapped capital is being released versus how quickly it should be released to redeploy into better-performing buys
<b>Weeks of supply</b>	Inventory coverage versus expected sell-through — over-coverage is a capital problem, under-coverage is a lost-sales problem
<b>Plan variance trend</b>	Whether the gap between plan and actual is widening or narrowing — trend direction matters as much as current variance level
<b>Aged inventory accumulation</b>	Capital trapped from prior buying decisions that is crowding out current OTB availability

Obviously, different retailers will have a differentiated set of variables to consider. These are often visible in the Category Plan, or a WSSI, or any other multi-KPI report needed by the team's merchants. These considerations apply even in organizations where Pre-Season Planning dominates and replenishment is not possible (think Fast Fashion), or conversely – replenishment dominates and planning is more likely to be expressed in terms of vendor and assortment breadth (think Bird Seed or Hardware).

## 4. The HOW Layer — Recommended Actions

---

The HOW layer is where the framework produces its output: specific, actionable recommendations that the merchant can accept, modify, or override. It is the connection between analysis and decision.

### 4.1 From reporting what happened to recommending what to do

Most reporting and planning systems produce descriptive output. They tell the merchant what happened, or what the current state is. Some produce predictive output: given current trajectory, here is what will happen. The HOW layer goes one step further. It produces recommended actions: here is what you should do about it.

The distinction matters in practice. A merchant who opens a report showing that Category X is 12 points below plan still has to decide what to do. A merchant who opens a system that shows Category X is 12 points below plan and recommends accelerating the markdown on the slowest-moving items in the class while protecting OTB for the top-performing colorways is starting from a different place. The analysis has already been done. The merchant's job is judgment, not excavation.

### 4.2 HOW actions for problem scenarios

Problem scenarios require remediation actions that address the diagnosed cause and reduce the probability of the same condition recurring. The HOW layer for a problem scenario might include:

- Receipt adjustments — canceling, reducing, or rescheduling forward commitments where demand signals have deteriorated.
- Markdown timing acceleration — releasing capital trapped in slow-moving inventory earlier than planned when the cost of holding exceeds the cost of marking down.
- Allocation rebalancing — shifting available inventory between locations based on current sell-through differentials.
- OTB protection — reducing planned forward commitments in underperforming categories to preserve dollars for redeployment.
- Vendor negotiation triggers — flagging specific purchase orders for renegotiation when receipt timing is threatening the selling window.

**The key in the HOW layer is team socialization and buy-in by the merchants and operations members. Recommendations or remediation actions that get ignored or are made too late – destroy the effectiveness of the 3Q Framework.**

### 4.3 HOW actions for opportunity scenarios

Opportunity scenarios require acceleration and reinforcement. The tendency is to leave outperformers alone. The 3Q approach is to pursue the upside actively. Some common methods (out of hundreds of options) are:

- Replenishment acceleration — increasing the frequency or depth of replenishment on items with strong sell-through before the selling window closes.
- OTB reallocation toward outperformers — redirecting uncommitted open-to-buy from weaker categories into the ones that are demonstrably performing.
- Receipt pull-forward — working with vendors to accelerate delivery on items where the demand signal is running ahead of plan.
- Assortment depth protection — holding back planned markdowns on full-price performers longer than the original cadence specified.

The merchant makes the final call on every HOW recommendation. The system's job is to ensure that the right options are in front of the right person at the time the decision needs to be made, not after the window has closed.

#### HOW runs in both directions

The planning community tends to treat the HOW layer as remediation. The full framework treats it as bidirectional: equally focused on capturing upside as on limiting downside. In a constrained OTB environment, the decision to redeploy capital from a slow category into a fast one is as important as the decision to markdown the slow one. Both are HOW actions. Both require a WHY diagnosis first. And neither is visible without the WHAT layer surfacing the signal.

## 5. A Worked Example — 3Q Applied to Three Candidate Buys

The 3Q Framework is most instructive when applied to a concrete planning decision. Consider a merchant managing open-to-buy across three competing buy candidates for the same OTB budget. Each represents a different category profile. The example is constructed so the variables pull in different directions — which is exactly where the framework’s diagnostic value becomes visible.

### 5.1 The scenario

	Buy A — Core Basics	Buy B — Import Fashion	Buy C — Trend Quick-Turn
Cost per unit	\$20	\$30	\$25
Full-price retail	\$50	\$90	\$60
Markdown retail	\$35	\$45	\$40
Units bought	1,000	1,000	1,000
OTB dollars committed	\$20,000	\$30,000	\$25,000
Expected FP / MD mix	80% / 20%	55% / 45%	90% / 10%
Expected revenue	\$47,000	\$69,750	\$58,000
Margin dollars	\$27,000	\$39,750	\$33,000
Capital tie-up (weeks)	14	20	11
Margin per OTB dollar	1.35	1.33	1.32
Margin per OTB \$ per week	<b>0.096</b>	<b>0.066</b>	<b>0.120 ← 1st</b>

## 5.2 3Q applied to the scenario

Here is how the three layers operate on this buying decision:

**WHAT surfaces the decision.** The system flags that OTB allocation is being evaluated across three candidates with materially different capital efficiency profiles. This is not a problem alert — it's an opportunity alert: the merchant has a chance to optimize the allocation before committing. The WHAT layer's job is to ensure that decision doesn't get made by default or by habit.

**WHY reveals what the numbers actually mean.** Ranked by total margin dollars, Buy B looks like the winner at \$39,750. Ranked by margin per OTB dollar alone, all three are nearly tied — between 1.32 and 1.35. But the WHY layer doesn't stop there. It loads vendor receipt timing (Buy B has a 12-week import lead time), markdown exposure (45% of Buy B's units are expected to clear at markdown), and capital tie-up (Buy B's dollars are working for 20 weeks; Buy C's for 11). When all three variables are running simultaneously, Buy C generates \$0.120 of margin per OTB dollar per week. Buy B generates \$0.066. The high-margin item is the least efficient use of constrained capital.

**HOW prescribes the action.** Prioritize Buy C. Protect OTB. Consider reducing Buy B's unit depth to lower capital tie-up, accepting lower total margin in exchange for a faster-cycling dollar that can fund additional cycles of Buy C within the same season. Flag the import lead time on Buy B as a structural risk to the selling window if it does proceed at full depth.

### Why this changes which buys win

A margin-led process would buy more of B and less of C. A turn-led process might prefer A. The 3Q approach runs all three diagnostic lenses simultaneously — and gives the merchant a single prioritized recommendation rather than three competing metrics. The output is not “here is more information.” The output is: here is what to do, and here is why.

## 6. Planning and Replenishment Through the 3Q Lens

The worked example is drawn from forward planning, but the 3Q Framework operates across both planning and replenishment. The decision logic is the same. The time horizon and variable set differ.

Layer	In Forward Planning	In Replenishment
<b>WHAT</b>	Category below plan; OTB over-committed; vendor receipt at risk; a fast-selling class with reorder opportunity	Item at or below min threshold; location over-inventoried versus cover target; fast-mover running below safety stock
<b>WHY</b>	Mix shift, demand change, receipt timing failure, markdown cadence mismatch, aged inventory drag	Demand acceleration, supplier delivery variance, allocation imbalance between locations, seasonal shift in rate of sale
<b>HOW</b>	Receipt adjustment, OTB reallocation, markdown pull-forward, buy depth reduction, assortment protection on outperformers	Replenishment order generation, min/max parameter adjustment, inter-location transfer, rate-of-sale curve recalibration

The complementary relationship between planning and replenishment becomes visible through this table. Planning generates the parameters that replenishment executes against — the assortment, the inventory targets, the receipt schedule. Replenishment generates the performance signals — rate of sale, actual versus planned sell-through, location-level inventory drift — that the planning layer needs to validate and revise its forward assumptions. 3Q is the framework that keeps both processes honest and connected.

## 7. Positioning and Implementation

---

The 3Q Framework is not a technology prescription. It is a decision discipline that any planning organization can adopt, and that the right software should support natively. The implementation principles follow directly from the framework:

- Don't start with the analytics. Start with the WHAT layer. Until the alert credibility problem is solved — until merchants know where to look and are confident the system will tell them — the diagnostic and recommendation layers have no audience. The exception alerting capability, and its calibration, is the precondition.
- Build the bidirectional habit. Most organizations have disciplined problem-response processes. Fewer have equally disciplined opportunity-capture processes. The WHAT layer should be configured to surface both with equal prominence. The WHY and HOW layers should be equally equipped to handle both.
- Keep the WHY layer honest. Diagnosis is only useful if the merchant trusts it. The causal analysis the WHY layer presents should be transparent — the merchant should be able to see which variables are driving the conclusion, not just the conclusion itself. Opaque diagnosis fails at the same moment as opaque forecasting: when the answer is wrong and no one can explain why.
- Make HOW recommendations the starting point, not the ending point. Prescriptive output is a proposal. The merchant's judgment is the decision. The system's job is to do the analysis so the merchant doesn't have to, and to present the right options in context when the window for action is open.
- Validate the framework with a proof of concept on real data. The three-buy example in this paper uses constructed numbers. The same analysis run against a retailer's actual open-to-buy, sell-through, and vendor receipt history will produce a specific, defensible recommendation for that business. That is the only way to calibrate whether the framework is producing useful signal for a given context.

### On technology and trust

Retail merchandising leaders have been burned by “the software does everything” pitches. The 3Q Framework does not aim to replace merchant judgment. It aims to make that judgment faster, better-informed, and less dependent on whether the right person ran the right report at the right time.

## 8. Summary

---

Planning and replenishment are different processes governed by different logic. Planning commits large, forward-looking inventory dollars under uncertainty. Replenishment manages current availability against a short cycle. Both processes require the same three things: knowing where to look, understanding what's driving what you find, and acting on it with a specific response. The 3Q Framework — What, Why, How — is the operating model that makes that sequence systematic.

The WHAT layer solves the signal problem. In a retail business with hundreds of categories, thousands of SKUs, and dozens of vendor relationships, the merchant cannot look everywhere. Proactive alerting — automatically delivered, threshold-calibrated, covering both problems and opportunities — curates the list of conditions that actually require a decision today.

The WHY layer solves the diagnosis problem. Planning performance depends on many factors at once. Sell-through, markdown mix, vendor timing, weeks of supply, and capital tie-up interact in ways that single-metric reporting cannot capture. Loading the full set of relevant metrics automatically, in context, on the flagged condition gives the merchant the causal picture without requiring them to build it from scratch.

The HOW layer solves the action problem. Knowing what is happening and why it is happening is not enough if the response is generic. The HOW layer surfaces specific, system-derived recommendations for both remediation (problems) and reinforcement (opportunities) — drawn from the planning and analytics engine's own output, presented when the decision window is open.

Together, the three layers form a framework that is honest about what it does: it does not replace the merchant. It ensures the merchant's judgment is applied to the right decisions, with the right information, at the right time.

---

**ANT USA Inc.**

Home of The Buyers Toolbox | [www.antusa.com](http://www.antusa.com)