RETAIL CASE STUDY / INVENTORY AND REPLENISHMENT

The Order Is the Decision

Replacing rule-based replenishment with an economic ordering model: censored demand, priced error, and replay validation

EXECUTIVE CONCLUSION

The forecast was not the problem. The order was. Rules that treat running out and overbuying as equally costly misprice every item in the catalog.

A large online retailer replenished a fast-changing catalog with weeks-of-cover rules and service-level settings. This note describes the economic ordering model that replaced them: demand estimated as a distribution and corrected for stockout censoring, an order-up-to policy derived from the two costs of error, discrete supplier rules handled exactly, and a replay of historical decisions before launch. Around that record it sets out the methodology an operations research team needs: censored estimation, the lost-sales caveat to base-stock theory, coupled constraints, decision-focused forecast evaluation, off-policy replay and its pitfalls.

$13M+
measured financial impact

2 costs
priced per item: shortage and overage

Censored
demand corrected before estimation

Replay
validated before any live order

Work led by Adam DeJans Jr. in an earlier industry role, not a Bit Bros client engagement. The $13M+ is measured financial impact as reported by the project; Bit Bros has not independently audited it. The retailer is not named. Worked examples and exhibits use synthetic numbers.

The replenishment problem

Each item i is reviewed every R periods and replenished with lead time L. At review the planner sees inventory position IP_i (on hand plus on order, less backorders where they exist) and places an order $q_i \geq 0$ that must respect case packs, supplier minimums and budget. Demand over the protection interval, $D_i = D_i(R{+}L)$, is uncertain, often intermittent, and observed only through sales. Online, much unmet demand is lost rather than backordered: a customer who finds an item unavailable buys elsewhere or buys a substitute.

The incumbent process set a target in weeks of cover or a service level for each category, then ordered up to it. Those targets are easy to explain and identical in spirit for every item, which is precisely their weakness: a high-margin item that sells out loses far more than it costs to hold another unit, while a low-margin, bulky item is the reverse.

THE DECISION

At each review, how much of each item should be ordered so that the expected cost of the two possible mistakes, running out and overbuying, is smallest, subject to the supplier's ordering rules and any shared budget?

Notation and inventory dynamics

Symbol	Meaning
$R,\ L$	Review period and replenishment lead time, in periods
$I_t,\ IP_t$	On-hand inventory and inventory position at the start of period t
q_t	Order placed in period t, arriving in period $t{+}L$
$D_t,\ s_t,\ A_t$	Demand, recorded sales and units available in period t
$b,\ h$	Cost of one unit short; cost of holding one unit for one period
$\tau = b/(b{+}h)$	Critical ratio
S	Order-up-to level; under base-stock $q_t = (S - IP_t)^+$
$F_X,\ \bar{F}_X$	Distribution function of X and its tail, $\bar{F}_X(x) = \Pr\{X > x\}$

The two regimes differ in their state equations. With backorders, on-hand inventory may go negative and

$$I_{t+1} = I_t + q_{t-L} - D_t, \qquad IP_{t+1} = IP_t + q_t - D_t,$$

so inventory position is a sufficient summary of the pipeline. With lost sales,

$$A_t = I_t + q_{t-L}, \quad s_t = \min(D_t, A_t), \quad I_{t+1} = (A_t - D_t)^+,$$

and $(D_t - A_t)^+$ units are lost. Inventory position no longer summarizes the system: the state is the whole vector $(I_t, q_{t-L}, \ldots, q_{t-1})$, which is the root of the lost-sales difficulty in Section 2. The same equation, $s_t = \min(D_t, A_t)$, is the root of the censoring problem in Section 1.

1. Sales are not demand

Stockouts censor the history the model learns from

Observed sales are $s_t = \min(D_t, A_t)$, where A_t is what was available. Whenever an item runs out, the record shows what was sold, not what was wanted. A model fitted to raw sales therefore underestimates demand exactly where the business has under-stocked, and the error is not noise: it is a bias with a consistent sign, and it grows as the in-stock rate falls.

The bias feeds itself. Orders set from an understated estimate produce more stockouts, which censor the next history further. Exhibit 1 traces that loop for two items with identical demand (negative binomial, mean 20, variance 120 per period) and different economics. Each cycle re-estimates mean demand from censored sales and resets the order-up-to level at the item's critical ratio; the plotted path uses expected sales, the long-history limit.

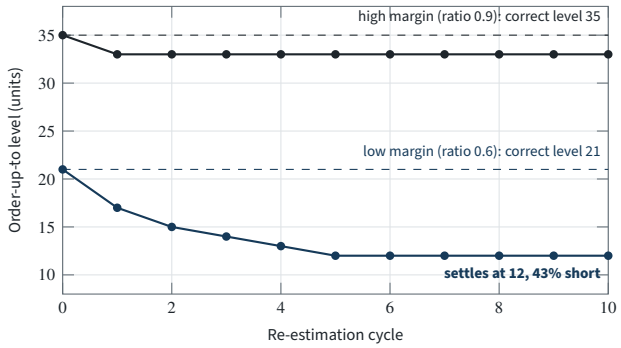

Exhibit 1. Order-up-to level when mean demand is re-estimated each cycle from censored sales (solid) against the correct level (dashed). Synthetic: negative binomial demand, mean 20, dispersion 4; expected censored sales.

The high-margin item barely moves, because its target sits far above the mean and stockouts are rare. The low-margin item collapses to 43% below its correct level and stays there, because it never again shows the demand it lost. With a finite year of weekly history the loop is worse, not better: across 200 simulated runs the naive estimate drives the low-margin item to 10.8 units on average after ten cycles, while a censored likelihood estimate holds it at 21.3. Correcting for censoring matters most on precisely the items a uniform rule already under-serves.

Estimate demand with a censored likelihood

The parametric fix treats stocked-out periods as censored observations. For a demand model $f(\cdot \mid \theta)$, the likelihood over periods t is

$$\mathcal{L}(\theta) = \prod_{t:\, s_t < A_t} f(s_t \mid \theta) \prod_{t:\, s_t = A_t} \Pr\{D \geq A_t \mid \theta\},$$

so a stocked-out period contributes the probability that demand was *at least* what was sold. For the negative binomial with mean μ and dispersion k,

$$f(d \mid \mu, k) = \frac{\Gamma(d{+}k)}{\Gamma(k)\, d!} \left(\frac{k}{k{+}\mu}\right)^k \left(\frac{\mu}{k{+}\mu}\right)^d, \qquad \mathrm{Var}(D) = \mu + \frac{\mu^2}{k},$$

and the censored log-likelihood $\ell(\mu, k) = \sum_{t \notin C} \log f(s_t) + \sum_{t \in C} \log \bar{F}(A_t - 1)$, with C the stocked-out periods, is maximized numerically over two parameters.

The EM algorithm imputes what the cap hid

Expectation-maximization treats the true demand of censored periods as missing data. The E-step replaces each censored observation by its conditional expectation given the current estimate; the M-step refits as if demand had been observed. For Poisson demand both steps are closed form, because $\mathbb{E}[D\,\mathbf{1}\{D \geq a\}] = \lambda \Pr\{D \geq a - 1\}$:

$$\hat{d}_t = \lambda \frac{\Pr\{D \geq A_t - 1\}}{\Pr\{D \geq A_t\}} \quad (t \in C), \qquad \lambda \leftarrow \frac{1}{T} \sum_{t=1}^{T} \hat{d}_t,$$

with $\hat{d}_t = s_t$ for uncensored periods. Each iteration cannot decrease the likelihood, and the fixed point is the censored maximum-likelihood estimate.

Algorithm 1. EM for censored Poisson demand

Require: sales $s_{1..T}$, censored set C, caps A_t, tolerance ε
$\lambda \leftarrow \frac{1}{T}\sum_t s_t$ ▷ naive start
repeat
  for $t = 1, \dots, T$ **do**
    $\hat{d}_t \leftarrow s_t$ if $t \notin C$, else $\lambda\, \bar{F}(A_t - 2)/\bar{F}(A_t - 1)$
  end for
  $\lambda' \leftarrow \lambda$; $\lambda \leftarrow \frac{1}{T}\sum_t \hat{d}_t$
until $|\lambda - \lambda'| < \varepsilon$
return λ

On a synthetic year of weekly history with true mean 20 and stock capped at 18, so that 70% of weeks stock out, the naive mean averages 17.1 (bias −2.9) across 2,000 replications, while EM averages 20.1 (bias +0.1). The price of removing the bias is variance: the EM estimate's standard deviation is 0.81 against 0.24, because censored weeks carry less information. Heavily censored items therefore need more history, or pooling, to estimate well. Exhibit 3 repeats the experiment across stock caps: the naive bias grows steadily with the censoring rate, reaching −31% when almost every week stocks out, while EM stays within a few percent of the truth at the cost of a wider spread.

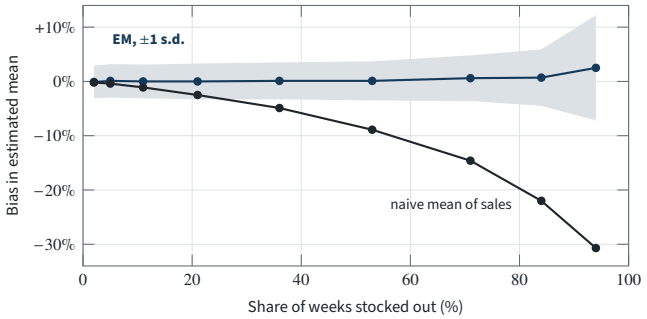

Exhibit 3. Bias of the estimated mean against the share of censored weeks: naive mean of sales (dark) and EM estimate with a band of one standard deviation (blue). Synthetic: Poisson demand, mean 20, one year of weekly history, stock caps from 14 to 30, 1,000 replications each.

Kaplan–Meier when the family is unknown

When stockout levels vary across periods, the distribution can be recovered without assuming a family. With e_j the number of uncensored periods whose sales equal j and r_j the number of periods with $s_t \geq j$, the product-limit estimator is

$$\hat{\bar{F}}(d) = \prod_{j \leq d} \left(1 - \frac{e_j}{r_j}\right).$$

It is consistent only up to the largest cap ever observed: beyond that level the data are silent, and a parametric tail must be grafted on.

Three refinements matter in practice. Where in-stock time within a period is recorded, the period is only partly censored, and scaling demand by the in-stock fraction (or modeling arrivals within the period) uses that information. Sparse items should borrow strength from similar items through hierarchical or empirical-Bayes shrinkage. In the gamma–Poisson case, a category prior $\lambda_i \sim \text{Gamma}(\alpha, \beta)$ and n_i periods of demand totalling y_i give the posterior mean

$$\mathbb{E}[\lambda_i \mid y_i] = \frac{\alpha + y_i}{\beta + n_i} = w_i \frac{y_i}{n_i} + (1 - w_i)\frac{\alpha}{\beta}, \qquad w_i = \frac{n_i}{n_i + \beta},$$

a weighted average of the item's own rate and its category's, with α and β estimated from the category by the method of moments. With a category mean of 6 per week worth two weeks of data ($\alpha = 12$, $\beta = 2$), an item averaging 2 per week over four weeks is estimated at 3.33; after a year, at 2.15. And stockouts push some demand onto substitutes, which inflates their apparent demand; substitution-aware estimation, or at minimum flagging the affected periods, prevents a stockout in one item from distorting the forecast of another.

Demand model	When it fits
Poisson	Low-volume items with variance close to the mean
Negative binomial	Over-dispersed demand, the common case online
Zero-inflated	Many periods with no demand at all, then bursts
Empirical or quantile	Large-volume items with enough history to skip a parametric family

Protection-interval demand

The policy needs the distribution of demand over $R{+}L$, not over one period. With independent per-period demand of mean μ and variance σ^2 and a random lead time independent of demand, write $N = R{+}L$ and $D = \sum_{t=1}^{N} D_t$. Conditioning on N and applying the laws of total expectation and total variance,

$$\begin{aligned}\mathbb{E}[D] &= \mathbb{E}\big[\mathbb{E}[D \mid N]\big] = \mu\,\mathbb{E}[N],\\ \operatorname{Var}(D) &= \mathbb{E}\big[\operatorname{Var}(D \mid N)\big] + \operatorname{Var}\big(\mathbb{E}[D \mid N]\big)\\ &= \sigma^2\,\mathbb{E}[N] + \mu^2 \operatorname{Var}(L).\end{aligned}$$

Lead-time variability often dominates the second term for imported goods, and ignoring it understates safety stock in exactly the items with the longest, least reliable supply.

2. Price both mistakes

The base-stock benchmark

Let b be the cost of one unit short and h the cost of holding one unit for one review period. If unmet demand is backordered, the optimal periodic-review policy is base-stock: order up to S each review, with

$$S^* = \min\left\{S :\ F_{D(R+L)}(S) \ \geq\ \frac{b}{b+h}\right\},$$

the multi-period analogue of the newsvendor critical ratio. The derivation is short. Writing D for protection-interval demand, the expected cost of ordering up to S is

$$G(S) = h\,\mathbb{E}[(S - D)^+] + b\,\mathbb{E}[(D - S)^+],$$

and since $\frac{d}{dS}\mathbb{E}[(S-D)^+] = F(S)$ and $\frac{d}{dS}\mathbb{E}[(D-S)^+] = -\bar{F}(S)$,

$$G'(S) = h\,F(S) - b\,\big(1 - F(S)\big) = (b{+}h)F(S) - b,$$

which is increasing, so G is convex and $G'(S^*) = 0$ gives $F(S^*) = \tau$. For integer demand the first difference plays the same role: $\Delta G(S) = G(S{+}1) - G(S) = (b{+}h)F(S) - b$, and S^* is the smallest S with $\Delta G(S) \geq 0$. Convexity is what makes the search, and the discrete extensions below, cheap. The service level is therefore an output, computed item by item from two numbers finance can sign, not an input set once per category.

The lost-sales caveat

Online, much unmet demand is lost, and with a positive lead time base-stock is then no longer optimal: the optimal order depends on the whole pipeline, not only on inventory position, and computing it exactly is intractable for realistic lead times. The research literature offers two useful anchors. Base-stock policies are asymptotically optimal as the lost-sales penalty grows relative to holding cost, which covers most high-margin items, and constant-order policies become asymptotically optimal as lead times grow long. The practical method is to use base-stock with the ratio $b/(b+h)$ as a starting point and tune the level per item, or per segment, by simulation against the lost-sales dynamics. Formally this is sample average approximation: draw N demand paths $\omega_1, \dots, \omega_N$ from the corrected model, simulate the lost-sales recursion under each candidate level, and choose

$$\hat{S} = \arg\min_{S \in \mathbb{Z}_+} \frac{1}{NT}\sum_{n=1}^{N}\sum_{t=1}^{T}\Big[h\,I_{t+1}(S;\omega_n) + b\,\big(D_t(\omega_n) - A_t(S;\omega_n)\big)^+\Big].$$

Using the same paths for every candidate (common random numbers) makes differences between neighboring levels precise even when each level's cost is noisy. The simulated cost is typically unimodal in S, though under lost sales this is not guaranteed, so a local search started at the base-stock level is checked against a coarse grid.

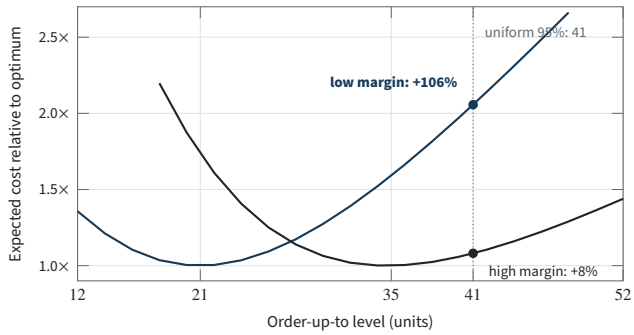

Exhibit 2. Expected cost of error against the order-up-to level, relative to each item's optimum: low-margin item (ratio 0.6, optimum 21) and high-margin item (ratio 0.9, optimum 35), same demand as Exhibit 1. A uniform 95% cycle-service target sets both at 41.

Exhibit 2 prices the uniform target directly. Both items share demand, so a 95% cycle-service rule gives both an order-up-to level of 41. For the high-margin item that costs 8% more than its optimum. For the low-margin item it more than doubles the expected cost of error. The same rule is close to right for one item and badly wrong for its neighbor, which is why item-level economics beat any category-level target.

Cost component	What it captures	Owner
Lost margin	Contribution forgone on a lost sale	Finance
Lost future demand	Customers who do not return after a stock-out, where measured	Commercial
Holding and capital	Storage, handling and cost of capital per unit-period	Finance
Markdown risk	Expected loss when excess must be cleared	Merchandising

Discrete supplier rules and coupled constraints

A continuous S is rarely orderable. With case packs of size m, convexity of the expected cost in the order-up-to level means the best feasible quantity is one of the two pack multiples bracketing the continuous optimum, so rounding is exact with two evaluations. Minimum order quantities and vendor-level minimums couple items. With $z_{ik} = 1$ if item i takes the k-th feasible quantity q_{ik} (including $q_{i0} = 0$), $w = 1$ if the vendor is ordered, fixed ordering cost K and unit price p_i,

$$\begin{aligned} \min\ & Kw + \sum_{i,k} G_i(IP_i{+}q_{ik})\, z_{ik} \\ \text{s.t.}\ & \sum_k z_{ik} = 1, \qquad \sum_{k\geq 1} z_{ik} \leq w \quad \forall i, \\ & \sum_{i,k} p_i q_{ik} z_{ik} \ \geq\ M w, \qquad z, w \in \{0,1\}. \end{aligned}$$

The linking constraint $\sum_{k\geq 1} z_{ik} \leq w$ is essential: without it, setting $w = 0$ would relax the minimum while still allowing orders. Instances of this size typically solve in a fraction of a second per vendor. A shared budget B across many vendors is better handled by Lagrangian relaxation. Dualizing $\sum_i p_i q_i \leq B$ with multiplier $\lambda \geq 0$ gives

$$\mathcal{L}(\lambda) = \sum_i \min_{q_i} \left\{ G_i(IP_i{+}q_i) + \lambda p_i q_i \right\} - \lambda B,$$

which separates by item: each item solves its own problem with unit cost raised by λp_i, equivalent to a critical ratio of $(b_i - \lambda p_i)/(b_i + h_i)$. The dual function is concave and spend is non-increasing in λ, so bisection on λ finds the price at which spend meets B. Complementary slackness, $\lambda^*\big(\sum_i p_i q_i^* - B\big) = 0$, says the budget carries a price only when it binds; λ^* is then the marginal value of one more dollar of budget. With integer quantities a small duality gap can remain, closed by a greedy pass over the items with the best marginal return. Where several items share a vendor's fixed cost, joint-replenishment or can-order policies decide which items to add when one triggers an order.

3. Judge the forecast by the decision

A probabilistic forecast feeds this policy through its quantiles, so its evaluation should match. The pinball loss at level τ is

$$\rho_\tau(u) = u\left(\tau - \mathbf{1}\{u < 0\}\right), \qquad u = D - S,$$

and its expectation is minimized at the τ-quantile of D. More than that, at $\tau = b/(b{+}h)$ it *is* the decision cost: checking the two cases $D \geq S$ and $D < S$ shows

$$G(S) = (b{+}h)\,\mathbb{E}\big[\rho_\tau(D - S)\big].$$

A forecast scored by pinball loss at the item's own critical ratio is therefore scored in dollars, up to a constant. Calibration is checked with the probability integral transform, whose values should be uniform if predicted quantiles mean what they say; and interval coverage is tracked by segment. Point-forecast accuracy metrics such as MAPE are a weak guide here, since two forecasts with the same error can imply very different costs once the asymmetry between shortage and overage is priced.

4. Replay before a single live order

Replay is off-policy evaluation

Before launch, the policy was replayed against historical periods and compared with the decisions actually taken, both priced against the same costs. Statistically, this is off-policy evaluation: estimating how a new policy would have performed using data generated by the old one. A credible replay obeys four disciplines.

Discipline	Practice
No hindsight	Freeze the data available at each decision date and replay from it; one leak discredits the result
Simulate, do not reuse sales	Evaluate candidate orders against reconstructed demand, because recorded sales were capped by the old orders
Avoid grading your own homework	The demand model that sets orders should not alone generate the demand that scores them; use held-out periods and stress the estimate
Report a range	State what was excluded and present the value as an interval, not a point

Two estimator families exist. Importance weighting reweights logged outcomes by $\pi_{\text{new}}(a \mid x)/\pi_{\text{old}}(a \mid x)$ and is unbiased, but it requires the old policy to have randomized over the actions the new one takes; deterministic historical orders give it no support, and near-deterministic ones give it enormous variance. The direct method instead simulates the new policy through a model of demand and dynamics; its variance is low but it inherits every error in the model. Doubly robust estimators combine the two. Replenishment replays are almost always direct-method estimates, which is why the disciplines below matter.

The second and third disciplines interact. Recorded sales cannot reveal what a larger order would have sold, so a sound replay draws demand from the corrected distribution and simulates inventory forward. But if the same model both chooses and scores the orders, the replay inherits its errors and flatters it. Scoring under several plausible demand models, and on periods withheld from fitting, shows how much of the estimated value survives model error. A replay is strong evidence; a randomized live comparison across items is stronger, and is the natural next step after launch.

REFERENCE DECISION FLOW

Sales and availability → censored demand model → protection-interval distribution → priced order-up-to level → discrete and coupled rules → replay, then release

5. What decides whether the model is used

Make the economics legible to planners and finance

Each recommended order should carry its reason: the demand distribution, the two costs, the implied critical ratio and the rule that set the final quantity. Planners can then challenge an input rather than override an output, and finance can see why two similar items receive different service levels.

Record overrides, then learn from them

Overrides should remain possible and be logged with a reason code. Over time the record separates overrides that carry information the model lacks, such as a promotion it has not seen, from those that merely restate the old rules, and the former can feed back into the model's inputs.

Monitor the decision, not only the forecast

Calibration by segment, realized cost against the cost the model predicted, in-stock rates on low-margin items, and the share of periods censored belong on the operating review. A rising censoring rate is the earliest sign that the loop in Exhibit 1 is returning.

6. Limitations

The method assumes demand is roughly stationary within an estimation window, which promotions, launches and seasonality violate; those require explicit covariates or separate models. Substitution and halo effects across items are hard to identify from censored data. Lost future demand, the most important part of shortage cost for some items, is rarely measured well and should be stated as an assumption with a sensitivity range rather than buried in a constant.

7. What the project reported

The ordering model launched into production and delivered more than $13 million in measured financial impact, as reported by the project. Just as important, the economics behind each order became explicit, giving planners and finance a shared basis for discussing trade-offs. The figure has not been independently audited by Bit Bros, and the contributions of censoring correction, priced costs and process change were not separated.

8. Lessons that transfer

Fix censoring before tuning the forecast

A more accurate model trained on capped sales learns the cap more accurately. Correcting the history is usually the largest single improvement available, and it costs little.

The order, not the forecast, is what earns or loses money

Forecast accuracy matters only through its effect on the order. Score forecasts at the quantile the policy uses.

Uniform targets fail most on low-margin items

They are the items whose correct service level is lowest, whose cost a uniform target inflates most (Exhibit 2), and whose censoring loop is most destructive (Exhibit 1).

A USEFUL EXECUTIVE DIAGNOSTIC

For your lowest-margin items, how was the demand estimate corrected for the weeks they were out of stock, and what service level do their economics actually justify?

Further reading

Conrad, S. A. Sales data and the estimation of demand. *Operational Research Quarterly* 27(1), 1976.

Nahmias, S. Demand estimation in lost sales inventory systems. *Naval Research Logistics* 41(6), 1994.

Kaplan, E. L. and Meier, P. Nonparametric estimation from incomplete observations. *Journal of the American Statistical Association* 53(282), 1958.

Zipkin, P. Old and new methods for lost-sales inventory systems. *Operations Research* 56(5), 2008.

Huh, W. T., Janakiraman, G., Muckstadt, J. A. and Rusmevichientong, P. Asymptotic optimality of order-up-to policies in lost sales inventory systems. *Management Science* 55(3), 2009.

Goldberg, D. A., Katz-Rogozhnikov, D. A., Lu, Y., Sharma, M. and Squillante, M. S. Asymptotic optimality of constant-order policies for lost sales inventory models with large lead times. *Mathematics of Operations Research* 41(3), 2016.

Dempster, A. P., Laird, N. M. and Rubin, D. B. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society B* 39(1), 1977.

Gneiting, T. and Raftery, A. E. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association* 102(477), 2007.

Kleywegt, A. J., Shapiro, A. and Homem-de-Mello, T. The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization* 12(2), 2002.

Dudik, M., Langford, J. and Li, L. Doubly robust policy evaluation and learning. *ICML*, 2011.

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

About the case and Bit Bros

This paper draws on replenishment work led by **Adam DeJans Jr.** in an earlier industry role. It was not a Bit Bros client engagement. The retailer's identity, proprietary models, data and operating parameters are not disclosed, and the mathematics is simplified to explain the method; it is not production code. Bit Bros is a boutique applied economics and optimization firm that helps operating teams price their decisions and build the systems that make them repeatable.

Adam DeJans Jr.
Technical Managing Director, Bit Bros LLC
Decision scientist specializing in inventory, optimization, simulation and decisions under uncertainty.

Correct the history. Price the error. Then order.