Dynamic Inventory Allocation for Seasonal Merchandise at Dillard's

Junxuan Li, Alejandro Toriello, He Wang
H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology

Seth Borin, Christina Gallarno
Dillard's Inc.

To appear in INFORMS Journal on Applied Analytics (Interfaces)

We consider how to allocate inventory of seasonal goods in a two-echelon distribution network for Dillard's Inc., a large department store chain in the U.S. Our objective is to allocate products with limited inventory from a distribution center to multiple retail stores over the selling season to maximize total sales revenue. Under the assumption that the true demand distributions are available to the retailer, we develop an effective dynamic inventory allocation heuristic. We further consider a more realistic and challenging setting for seasonal goods, where demand distributions are unknown to the retailer, and propose two "learning-while-doing" extensions of our inventory allocation heuristic; these policies update demand distribution estimates in a rolling horizon using censored point-of-sales data. We evaluate the performance of the policies using simulation on Dillard's historical sales data. Dillard's Inc. has incorporated the proposed policy into their current replenishment methodology and has been using the policy to set order levels for its seasonal merchandise.

Key words: inventory allocation, demand learning, seasonal goods, dynamic control.

Introduction and Background

Dillard's Inc. is a large department store chain in the U.S. with approximately 300 stores and $6.5 billion annual revenue in 2018. In this paper, we present an approach to allocate inventory of seasonal merchandise (e.g., fashion goods, holiday decorations) from Dillard's distribution centers to its retail stores. The specific merchandise groups that are considered in this study account for 20% to 30% of total revenue at Dillard's.

Inventory management for seasonal goods is particularly challenging, since these goods usually have short selling seasons, high demand uncertainty, and little salvage value at the end of their life cycles. In addition, seasonal goods often have long supply lead times and must be ordered from suppliers well before the selling season begins. In many cases, a product cannot be replenished during the selling season after its initial order quantity is determined. Therefore, the allocation of this limited inventory among retail locations can have a significant impact on the product's sales, and these inventory allocation decisions are quite important for companies like Dillard's.

For seasonal merchandise, most sales occur within a two to three week timeframe, and the uncertainty of the demand is often higher than the average demand (i.e., the coefficient of variation of demand is larger than 1). For instance, cold weather boots generally do not begin selling until cold weather sets in for an area, but the peak is nearly always the weeks between Thanksgiving and Christmas. Some cold weather boots may sell one pair of a size at most in any store during this time period; other styles could sell up to 30 or more pairs of a single size at a store in the same time horizon. Christmas decorations are another example where the selling season is extremely short and the sales are highly uncertain at a product-store level. In this case, we are unlikely to see the exact same product in the assortment year-over-year. These facts lead to great challenges in demand forecast and inventory allocation.

Figure 1 shows Dillard's store locations. Dillard's nationwide distribution network is divided into six regions. Stores in the same region are served by a regional distribution center (DC). Transshipment between the regions rarely happens in Dillard's distribution network, so we consider each region separately. We focus on one of the six regions in this study.

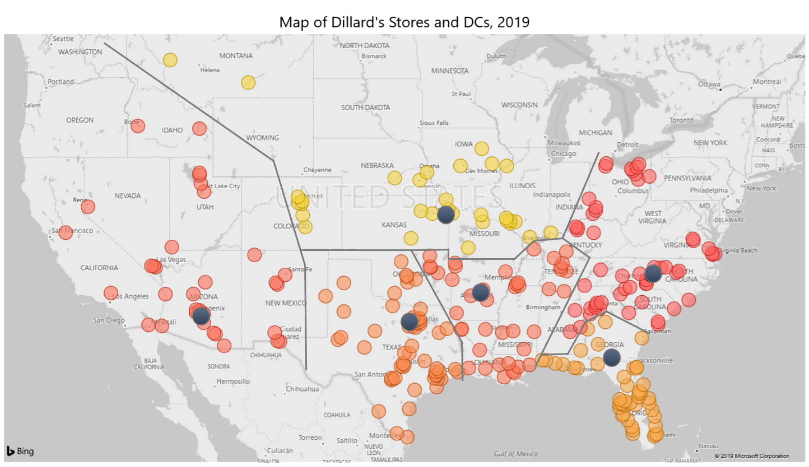

Figure 1 **Dillard's retail store locations. The nationwide distribution network is divided into six regions. Locations of distribution centers within each region are marked by the dark circles.**

The merchants at Dillard's work with vendors and manufacturers to determine products' order quantities into the DCs. In many cases, orders must be placed six to twelve months prior to delivery. The order quantity decisions are based on vendor supply, store financial plans and assortments, long term forecasts, shipping costs, vendor lead time, product life, strategic initiatives, and marketing campaigns. Therefore, when a product's selling season starts, inventory available at a DC is often fixed and cannot be replenished.

Inventory shipments are made throughout the selling season from the DC to the stores on a weekly basis, while unsatisfied customer demand is lost. After the product is marked down to its clearance price at the end of the selling season, no more shipments are made for it. Our objective is to determine the shipment quantities from the regional DC to the stores during the selling season, in order to maximize total expected revenue obtained from the product during its life cycle. Because the DC supplies many products to retail stores, shipments are on a fixed weekly schedule and can be considered a sunk cost; actual shipment decisions and quantities do not incur a cost. Furthermore, the quantities are quite small compared to the shipping capacity, and we can thus reasonably assume they are uncapacitated.

When compared to one-time allocations at the start of the selling season, making multiple inventory shipments throughout the season provides several benefits. First, demand forecasts for seasonal goods based on previous years' sales are usually inaccurate, so early sales in the season can help improve the forecasts. If there are multiple shipment opportunities, the retailer can adjust shipment quantities to the stores using these updated forecasts. Second, the inventory at the DCs provides a "risk pooling" effect: after the initial inventory shipment and any random demand fluctuations early in the season, subsequent shipments can help re-balance the inventory level at different stores in order to reduce lost sales. Based on this observation, we consider a dynamic, data-driven approach to allocating seasonal good inventory to retail stores.

Related Literature

Below, we discuss two streams of related literature. The first stream focuses on inventory management in a two-echelon distribution system with one warehouse/DC and multiple retailers, assuming the demand distribution is known. Examples include Eppen and Schrage (1981), Federgruen and Zipkin (1984), Jackson (1988), Jackson and Muckstadt (1989), McGavin et al. (1993), Graves (1996) and Axsater et al. (2002). Caro and Gallien (2010) considered multi-stage SKU-level inventory allocation with an additional constraint specifying that a product must be taken off the shelf if any of its major sizes is stocked out. The main difference between our problem and this stream of literature is that we assume the true SKU-level demand distributions are unknown.

The second stream of literature studies demand learning in multi-period inventory systems. Eppen and Iyer (1997) and Agrawal and Smith (2013) proposed Bayesian learning methods. The issue of demand censoring (unobserved lost sales) has been discussed in detail in Huh and Rusmevichientong (2009), Huh et al. (2011), Besbes and Muharremoglu (2013), and Jain et al. (2014). Fisher and Raman (1996), Fisher and Rajaram (2000), and Fisher et al. (2001) have also studied inventory management and demand learning for fashion goods.

The demand learning method proposed in our paper is motivated by ideas in Eppen and Iyer (1997), Jain et al. (2014), and Gallien et al. (2015). In particular, the two-stage model in Gallien et al. (2015) applied to Zara is related to the model in our paper, but with two major differences: first, Gallien et al. (2015) update demand estimation using a known A/F (actual demand vs. forecast demand) ratio distribution assuming no lost sales in the first period, while our demand learning methods use Poisson regression based on weekly censored sales data; second, Gallien et al. (2015) solve the two-period inventory allocation problem approximately using a greedy knapsack algorithm, while we solve the two-period inventory problem using a cut generation method.

Problem Description

We focus on one of Dillard's distribution regions, which contains 41 retail stores served by one distribution center (DC). The stores receive shipments from the DC, but cannot return inventory to the DC or transship it to other stores. Therefore, we treat the DC and its associated retail stores as a two-echelon distribution network; see Figure 2. The decision maker chooses when and how much of a product's inventory at the DC to ship to each store, based on the current inventory level, customer demand forecast, and observed past sales.

Our study focuses on Dillard's exclusive products; these are owned by Dillard's and produced many months before the selling season. An exclusive product typically has a predetermined price schedule and a primary selling window of 12-16 weeks. As we discussed previously, peak demand may occur in a narrower window of 2-3 weeks during the selling season.

Products are stored at the DC before the selling season starts. Then, after the season begins, products are shipped from the DC to each retail store. Due to long production lead times, the DC itself cannot be replenished during the season.

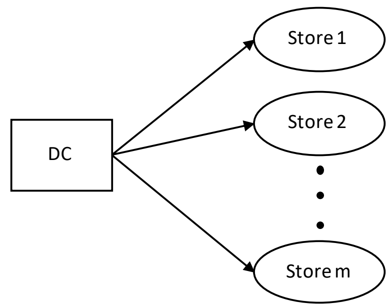

Figure 2 Illustration of a two-echelon supply chain.

At Dillard's, the life cycle of a seasonal item is determined by its markdown pricing schedule (usually fixed). The item is first offered at its full retail price, and then is marked down to 35%–40% off. We treat the duration of the full price and the first markdown price as the primary selling season of the product, which typically lasts 12–16 weeks. At the end of the primary selling season, if there is any inventory left, the item is marked down again at 60%–70% off for clearance; we assume this clearance price is the product's salvage value.

Shipments occur on a weekly basis, so we define each time period as one week. Orders are placed at the same time every week, i.e., when each period begins. The company has trucks that are scheduled to run from the DC to the store locations every week, so we assume transportation costs and ordering costs are sunk and do not consider them in the model. The lead time from the DC to the stores is usually 1-2 days. Similarly, exclusive product inventory is owned by Dillard's, so we treat holding costs as sunk and do not consider them in the model. We also do not consider shipment capacity restrictions, as the order quantities for a product are quite small relative to the total shipment capacity, which potentially includes many products going to several stores. Figure 3 shows the timeline of order placement and shipment. If there is insufficient inventory at a store, any unmet demand is lost. Unsold items at the end of selling season are salvaged at the clearance price. Our objective is to design an inventory control policy at the SKU level that manages shipments from the DC to the retail stores throughout the selling season, in order to maximize total expected revenue.

This inventory management problem requires information about weekly demand. We consider both a setting in which distributional information about weekly demand is known, and a more realistic setting for seasonal goods in which this distribution is not

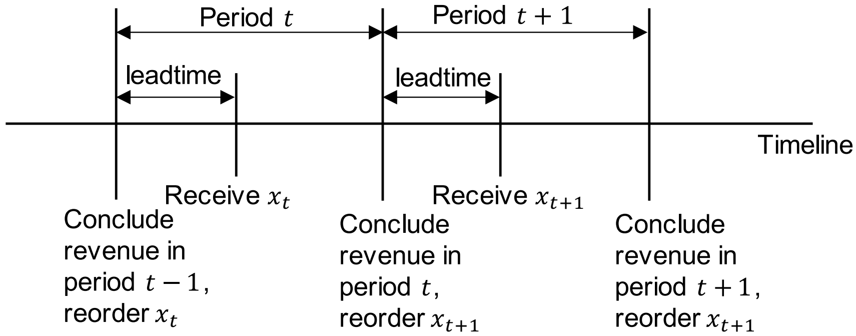

Figure 3 **Timeline of inventory order placement.**

known. When the demand distribution is unavailable, we assume that the decision maker is equipped with information technology that enables them to observe point-of-sales (POS) data, and thus potentially leverage it to learn or update a demand distribution during the season.

Inventory Management with Known Demand Distribution

We start by considering an idealized situation where demand distributions for all store locations and SKUs are known to the decision maker. We develop heuristics for this case that we later build on for the more complex and more realistic situation that involves unknown distributions.

Assuming the demand distribution is known, the problem can be formulated as a finite horizon dynamic program, where every period corresponds to one week of the selling season. The system state of the model tracks the current week as well as the previous week's ending inventory levels at the DC and retail stores. The actions are inventory replenishment quantities for each store; since inventory cannot be replenished at the DC, feasible replenishment quantities cannot exceed the DC's remaining inventory. The week's reward is the revenue from all store sales during that week; the terminal reward is the salvage value of the inventory left at the stores when the season ends. We include a mathematical formulation of the dynamic program in the Appendix.

However, even in the idealized situation with known demand distributions, the dynamic programming formulation poses significant computational challenges. Assuming every retail store can receive inventory, a distribution system with 41 stores yields at least 2^{41} (more than a trillion) possible states. Therefore, as solving the exact solution to the dynamic programming formulation is intractable, we examine two heuristic policies, a Ship-Once (SO) policy and a Mean Demand Heuristic (MDH) policy.

The Ship-Once (SO) Policy

The SO policy is a heuristic that allocates all of the DC's inventory to the stores at the beginning of the selling season. Under this policy, the inventory allocation problem reduces to a single-period newsvendor model, where the entire season's demand for each store is aggregated. In a single-period setting, it is optimal to ship all the DC's inventory to the stores. The resulting inventory level at each store is the sum of its initial inventory plus the DC's shipment. If the store's realized demand is higher than its inventory, the excess demand becomes lost sales; if the realized demand is lower than the inventory, unsold inventory is salvaged at the clearance price. This single-period problem can be formulated as a multi-location revenue maximizing newsvendor model, with a constraint specifying that the total amount of inventory shipped to the store should not exceed the DC's inventory. The formal mathematical definition of the SO policy can be found in the Appendix.

When implementing the SO policy, for each store location we first compute its aggregate demand by adding the weekly demand rates for the entire season. Since Dillard's uses markdown pricing, in the single-period model we treat the product's price at a store as the average of each week's price weighted by that week's demand rates. We then decide the shipment quantity to each store by solving the single-period inventory allocation model specified by the SO policy.

The Mean Demand Heuristic (MDH) Policy

The SO policy does not allow multiple shipments during the season. Multiple shipment opportunities provide a *risk-pooling effect*, as subsequent shipments can help re-balance the inventory level at different stores caused by random demand fluctuations. To exploit the opportunity afforded by multiple shipments, we propose a dynamic inventory allocation heuristic, which we call the Mean Demand Heuristic (MDH) policy. This policy allows shipment from the DC to the stores throughout the season. The MDH policy explicitly considers the fact that shipment decisions not only affect inventory levels of the current week, but also have an impact on the inventory levels of subsequent weeks.

The main idea of the MDH heuristic is to reduce the multi-period inventory allocation problem into a *two-stage* problem; the first stage contains the current period, and the second stage contains all subsequent periods. We make two further simplifications in order to solve the two-stage model efficiently: First, we solve the second-stage sub-problem approximately using the SO policy; second, we use mean demand values rather than full demand distributions to calculate the inventory levels at the end of the first stage, which gives the policy its name. These simplifications yield a two-stage stochastic optimization model that we solve with a cutting-plane algorithm; we include technical details in the Appendix.

The outputs of the MDH policy are shipment quantities for all stores in the first week. We then apply this policy in a rolling horizon to decide shipment quantities in every subsequent week. As an example, consider a selling season of 12 weeks. At the start of week 1, the MDH policy partitions the time horizon into two stages, the first containing week 1, the second containing all remaining weeks. We determine shipment quantities for week 1 using the model's solution. After the actual demand in week 1 is realized, we observe the inventory levels at the stores and the DC at the beginning of week 2. We then apply the MDH policy again, this time including week 2 in the first stage, and weeks 3 to 12 in the second. We apply the two-stage model in a rolling horizon framework until reaching the last week, where we simply apply the SO policy to determine shipment quantities.

Inventory Management with Unknown Demand Distribution

In practice, a demand forecast available at the start of the season may be inaccurate. This is especially true for goods with short life cycles (e.g. fashion products), where consumer tastes vary from year to year and forecasts are usually obtained using sales data from similar products sold in previous years. Therefore, sales in the early part of the season can provide valuable information to improve the forecasts.

This motivates us to incorporate learning into the inventory decision model by updating demand distributions as we observe actual sales data. The model framework is illustrated in Figure 4. Given an initial estimation of demand distributions, the inventory control model is solved for each SKU to determine shipments to each store. Then, a new batch of sales data is observed and added to the sales history database. We update demand distribution estimates by applying a weighted regularized Poisson regression, described in the next section. We then repeat the loop in Figure 4 until the season ends. Based on this framework, we propose demand learning policies that extend the MDH policy.

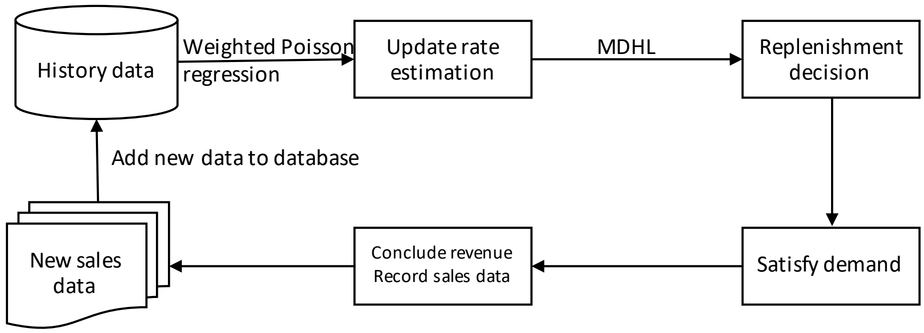

Figure 4 **Policy dynamics of decision dynamics incorporating learning**

Mean Demand Heuristic with Learning (MDHL) Policy

In a similar fashion to the MDH policy, the MDHL policy solves a two-stage inventory allocation model in a rolling horizon during the selling season. At the beginning of each week, MDHL includes newly observed sales data from the previous week into the historical demand data set, and updates the demand forecast using a weighted regularized Poisson regression. The policy then solves a two-period inventory model identical to the one in the MDH policy, with the updated demand forecast. The formal definition of the MDHL policy is given in the Appendix.

Demand Learning from Sales Data

Demand learning from week to week is a critical part of the MDHL policy. We next detail two forecasting methods for demand learning. We assume that weekly demand follows a Poisson distribution, which we find has a good fit to Dillard's data. We also assume that demands in different weeks and locations are independent. Based on these assumptions, we propose two approaches to estimating Poisson distributions from sales data, a classical maximum likelihood estimation (MLE) method and a regularized Poisson regression method. A key challenge for learning demand from sales data is *demand censoring*; when inventory at stores cannot meet customer demand, lost sales occur and are usually not observable. Both methods proposed are able to tackle the challenge of demand censoring, with different treatment on stock-out events. We conclude this section with a numerical validation of these methods.

Maximum Likelihood Estimation

Assuming the decision maker has access to daily sales volume and stock-out information, we can formulate the log-likelihood function for a Poisson demand distribution. Stock-out indicators are used to derive the conditional probability of demand realization given an aggregated demand rate, and the log-likelihood function is the logarithm of this conditional probability. One can show that the log-likelihood function is concave with respect to the unknown aggregated demand, and therefore can be solved efficiently. The complete procedure is included in the Appendix.

When estimating demand using maximum likelihood estimation, we are faced with a data sparsity issue, as most SKU's only have a few units sold per week. In other words, sales volume is too low at the weekly-SKU level for demand prediction. Therefore, instead of predicting demand separately for each SKU and week, we make a prediction of aggregated demand for the entire season and all SKU's. Then, given the aggregated forecast, we use seasonality factors suggested by Dillard's to get a disaggregated demand forecast for each week and SKU.

Regularized Poisson Regression

Poisson regression assumes that the daily demand rate of a SKU in a given store is represented by an exponential function in the form of $e^{\beta^\top x}$, where x is a vector containing information such as week, day of week, SKU, and store, and β is an unknown vector of parameters. Similar to the MLE method, we can formulate the log-likelihood function for the Poisson regression model and solve it efficiently. Stock-out indicators serve as signals of missing data in the regression. Since there are multiple unknown parameters in the Poison regression model, we add a LASSO regularization term to the log-likelihood maximization problem to avoid over-fitting. We include technical details in the Appendix.

Compared to MLE, regularized Poisson regression has two advantages. First, it does not require the assumption that the fraction of demand for each SKU-store-weekday combination is known. Second, in regularized Poisson regression, we can assign different weights to data in different weeks; this turns out to be convenient if we believe that recently observed data should have higher weights in demand estimation.

Comparing Demand Learning Methods

We test the two demand learning methods using actual sales data from Dillard's. The test includes nine products, six clusters of stores, and two price levels (full price and first markdown price).

We first compare some statistical properties for the two methods. We consider the following metrics:

- *Mean absolute residual*: This quantity measures the absolute difference between predicted daily demand and actual daily demand (when there is no stock-out). A smaller mean absolute residual means better prediction accuracy.
- *Mean/variance ratio*: This quantity measures the ratio between the mean and variance of estimated demand. Since a Poisson distribution has the same mean and variance, if actual demand follows the Poisson distribution, the ratio should be close to 1.
- *Mann-Whitney test*: We use the Mann-Whitney test to check if demand forecast from our prediction methods follows the same distribution as the actual demand data. (For each method, we generated 1,080 testing samples.) If the predicted demand and the actual demand are statistically different, the Mann-Whitney test will reject the null hypothesis. Therefore, fewer rejections means better prediction accuracy.

The testing results are listed in Table 1. Although both methods have small mean absolute residual, the regularized Poisson regression's residual is smaller. Both methods yield mean/variance ratios close to 1, indicating that it is reasonable to assume demand follows a Poisson distribution. Finally, in terms of the Mann-Whitney test, both methods have relatively few rejections, with regularized Poisson regression requiring fewer than classical MLE. Based on these results, we conclude that both methods produce reasonably accurate demand forecasts, with the regularized Poisson regression performing slightly better than MLE.

Table 1 Statistical Metrics of the Two Demand Prediction Methods

Demand estimation method	Mean absolute residual	Mean/Variance Ratios	Mann-Whitney Test
Classical Poisson MLE	0.012	1.1	38 rejects out of 1080
Regularized Poisson Regression	0.011	0.9	21 rejects out of 1080

In the two figures below, we focus on a particular product and two stores, labeled A and B. We use data for the first nine weeks of the season, when the product is sold at full retail price. In Figure 5, we plot the probability mass function (PMF) of the demand distribution estimated using the MLE method, as well as the PMF of the actual sales data (normalized by seasonality factors). Figure 6 shows the same plot for the regularized

Poisson regression method. The two methods produce almost identical results, and the predicted PMF is close the actual PMF.

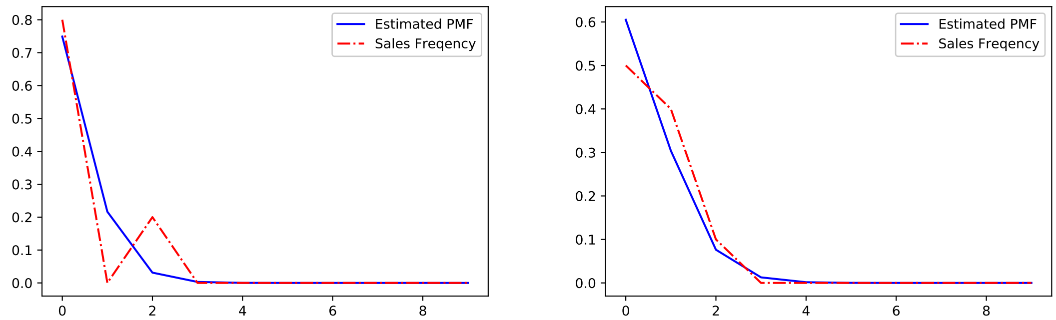

(a) Store A, MLE. (b) Store B, MLE.

Figure 5 **Classical MLE demand estimation visualization**

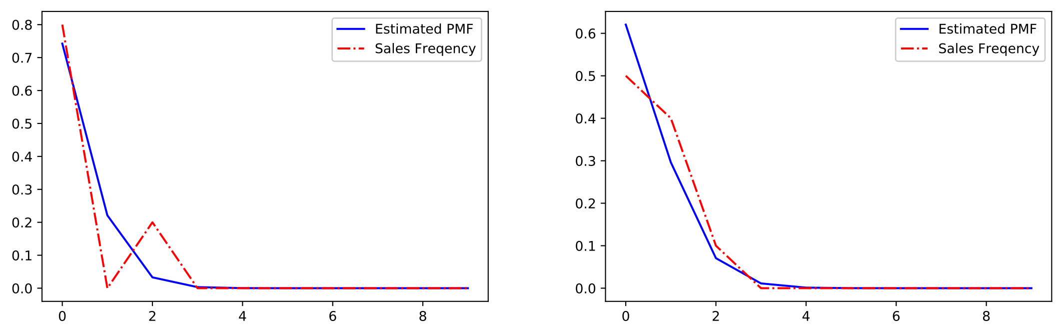

(a) Store A, Poisson Regression (b) Store B, Poisson Regression

Figure 6 **Regularized Poisson regression demand estimation visualization**

Implementation

Before we implemented the proposed models in Dillard's inventory management pipeline, various tests were conducted to demonstrate the potential values of different algorithms. These tests helped the implementation team in Dillard's to best utilize these models based on their current business system. The testing scripts and results generated in the preliminary tests helped different engineering teams, e.g. system engineers, database engineers, data scientists and software engineers, to improve current inventory control and sales systems to make these advanced models applicable to Dillard's current business platform. In this section, we briefly introduce preliminary tests, and then explain how Dillard's implemented the proposed methodologies in their business.

Preliminary Tests

Various preliminary tests were conducted to evaluate the proposed polices based on Dillard's sales data. We provide here a high-level overview of the tests to help the readers understand the implementation at Dillard's, while detailed test procedures and results are included in the Appendix.

Using a data set collected from Dillard's point-of-sale system, we generate a probabilistic demand model with a given price markdown process during the selling season. As average daily sales per SKU and store are often very low, we aggregated sales data by week and by product, and then generate a demand model using regularized Poisson regression. A demand forecast for each day, SKU and store can be obtained by disaggregation.

The first set of tests assumed the true demand distribution is known. We evaluated the SO and MDH policies using sampled testing scenarios and calculated expected revenues obtained by the two policies. Based on the test results, MDH outperformed SO with a significant increase in expected revenues, proving the benefit of using multiple shipments throughout the season.

The second set of tests explored more realistic scenarios in which the demand forecast available at the beginning of the season is inaccurate. The true demand rates for all stores were randomly perturbed to be distinct from the original demand forecast model. We evaluated three inventory control policies, SO, MDH, and MDHL, based on sampled testing scenarios, and calculated expected revenues obtained by these policies. We observed that the MDHL policy had the best performance among the three heuristic policies.

Field Implementation

Dillard's has already developed, tested, and implemented its own internal demand learning forecast based on a large range of variables. At the start of this work, Dillard's was in the process of developing a short-term forecast for seasonal merchandise. Since the completion of the paper, Dillard's has incorporated the MDHL policy into their current replenishment methodology and has been using the policy, with slight modifications, to set order levels for seasonal DC-replenished merchandise. Dillard's benefits from the MDHL policy for its continuous improvement of its short-term forecasts, while the initial demand inputs for the MDHL model are derived using an internal forecast.

Dillard's uses their own existing process to assimilate new merchandise into the replenishment system as well as their own end-of-life allocation process to flush residual merchandise at the end of the season, so the MDHL policy serves as a middleman, making allocation decisions throughout the body of the season. For newly received products, Dillard's has an established method for determining inventory positions early in the product life. For end-of-life merchandise, Dillard's requires additional flexibility in their process in order to allow management to change replenishment strategies mid-season (adding new stores, removing stores, adjusting product lifespans, and so forth). Such changes in strategy are not easily anticipated and thus cannot and should not be accounted for in the MDHL model. This process handles short-term, seasonal, and long-term merchandise, allowing for a consistent end-of-life process across the business.

To incorporate the MDHL policy into Dillard's replenishment system, Dillard's first worked with its buyers to standardize aspects of their replenishment strategy. A challenge faced by the implementation team was that not all fields in the dataset could be explicitly defined or documented; for example, the duration of "markdown periods" varies not only across products but also across years. For cases like these, the implementation team made initial assumptions and then fine-tuned these as information was revealed. The team then collected data (using SQL), scripted the model (using Python), and tested and validated its output. This process consisted of several iterations of parameter tuning, troubleshooting code, and boosting run time efficiency. Compared to the SQL data collection and Python processing components, the inventory allocation stage of the process, which solves the optimization model using Gurobi, uses the majority of the run time. The team finally had to translate the output of the model, which is expressed in terms of *unit transfers*, to that of *targeted inventory positions*, so as to fit with Dillard's existing replenishment systems.

Since April 2019, the MDHL policy has been operating on approximately 1,500 SKUs of merchandise across all stores, which translates to 4,300 SKU-DC combinations or 126,000 SKU-store pairs. For the initial rollout of the policy, the implementation team worked with Dillard's merchandising managers to select test products, and then ensured that the model was behaving as expected. Dillard's measured and communicated the success of replenishment methods to management and buyers through store in-stock rates, revenue, and sell-through rate. With the success of the initial rollout and positive feedback from Dillard's executives, the policy will be expanded to all seasonal DC-replenished merchandise. These products account for approximately 15,000 SKUs, which translates to 40,000 SKU-DC combinations or 1.2 million SKU-stores. In terms of revenue, the model is currently used on merchandise accounting for 6% of total revenue, and will expand to additional merchandise groups with a total scope of 20% to 30% of revenue.

References

Agrawal N, Smith SA (2013) Optimal inventory management for a retail chain with diverse store demands. *European Journal of Operational Research* 225(3):393–403.

Axsater S, Marklund J, Silver EA (2002) Heuristic methods for centralized control of one-warehouse, n-retailer inventory systems. *Manufacturing & Service Operations Management* 4(1):75–97.

Besbes O, Muharremoglu A (2013) On implications of demand censoring in the newsvendor problem. *Management Science* 59(6):1407–1424.

Caro F, Gallien J (2010) Inventory mangement of a fast-fation retail network. *Operations Research* 58(2):257–273.

Chen L, Mersereau A, Wang Z (2017) Optimal merchandise testing with limited inventory. *Operations Research* 65(4):968–991.

Eppen G, Iyer A (1997) Improved fashion buying with Bayesian updates. *Operations research* 45(6):805–819.

Eppen G, Schrage L (1981) Centralized ordering policies in a multi-warehouse system with leadtimes and random demand. *Multi-Level Production/Inventory Control Systems: Theory and Practice*, 51–69 (North Holland, Amsterdam: L. Schwarz, ed.).

Federgruen A, Zipkin P (1984) Approximations of dynamic, multilocation production and inventory problems. *Management Science* 30(1):69–84.

Fisher M, Rajaram K (2000) Accurate retail testing of fashion merchandise: Methodology and application. *Marketing Science* 19(3):266–278.

Fisher M, Rajaram K, Raman A (2001) Optimizing inventory replenishment of retail fashion products. *Manufacture & Service Operations Management* 3(3):230–241.

Fisher M, Raman A (1996) Reducing the cost of demand uncertainty through accurate response to early sales. *Operations Research* 44(1):87–99.

Gallien J, Mersereau A, Garro A, Mora AD, Vidal MN (2015) Initial shipement decisions for new products at zara. *Operations Research* 63(2):269–286.

Graves SC (1996) A multiechelon inventory model with fixed replenishment intervals. *Management Science* 42(1):1–18.

Huh WT, Levi R, Rusmevichientong P, Orlin JB (2011) Adaptive data-driven inventory control with censored demand based on kaplan-meier estimator. *Operations Research* 59(4):929–941.

Huh WT, Rusmevichientong P (2009) A nonparametric asymptotic analysis of inventory planning with censored demand. *Mathematics of Operations Research* 34(1):103–123.

Jackson PL (1988) Stock allocation in a two-echelon distribution system or "what to do until your ship comes in". *Management Science* 34(7):880–895.

Jackson PL, Muckstadt JA (1989) Risk pooling in a two-period, two-echelon inventory stocking and allocation problem. *Naval Research Logistics (NRL)* 36(1):1–26.

Jain A, Rudi N, Wang T (2014) Demand estimation and ordering under censoring: Stock-out timing is (almost) all you need. *Operations Research* 63(1):134–150.

McGavin EJ, Schwarz LB, Ward JE (1993) Two-interval inventory-allocation policies in a one-warehouse n-identical-retailer distribution system. *Management Science* 39(9):1092–1107.

Appendix 1: Dynamic Inventory Allocation Model

Let $\boldsymbol{s}_t = (s_t^0, s_t^1, ..., s_t^m)$ be the vector of SKU inventory at the end of period t, whose components are inventories at the DC and m stores respectively. Let D_t^i be random demand for store i in period t, which follows distribution F_t^i. We formulate multi-location multi-period inventory allocation problem as finite horizon dynamic programming problem. Let $\boldsymbol{s}_0 = (s_0^0, s_0^1, ..., s_0^m)$ be initial SKU inventory for the DC and stores, where s_0^0 is total SKU amount available for the entire selling season, and $s_0^i = 0$ for $i = 1, ..., m$. Let $\boldsymbol{s}_{t-1}$ be the state of period t, state space can be then expressed as $\mathcal{S} = \{(s^0, s^1, ..., s^m) : \sum_{i=0}^m s^i \leq s_0^0\}$. Replenishment amount $\boldsymbol{x}_t = (x_t^1, ..., x_t^m)$ is then the action in period t, with action space $\mathcal{A}_t = \{\boldsymbol{x}_t = (x_t^1, ..., x_t^m) : \sum_{i=1}^m x_t^i \leq s_{t-1}^0\}$. Denote r_t^i be selling price in store i and period t, we define reward function for each period,

$$R_t = \begin{cases} \sum_{i=1}^m r_1^i \cdot \min\{D_1^i, s_0^i + x_1^i\}, & t = 1 \\ \sum_{i=1}^m r_t^i \cdot (\min\{D_{t,L}^i, s_{t-1}^i\} + \min\{D_{t,R}^i, (s_{t-1}^i - D_{t,L}^i)^+ + x_t^i\}), & t = 2, ..., T \end{cases}$$

where $D_{t,L}^i$ is the demand occurred during lead time, and $D_{t,R}^i$ is the demand after the replenishment in store i and period t, e.g. $D_t^i = D_{t,L}^i + D_{t,R}^i$. Then optimality equation of T period dynamic programming problem

$$V_1(\boldsymbol{s}_0) = \max_{\boldsymbol{x}_1 \in \mathcal{A}_1} \mathbb{E}\left[\sum_{i=1}^m r_1^i \cdot \min\{D_1^i, s_0^i + x_1^i\} + V_2(\boldsymbol{s}_1)\right] \tag{1a}$$

$$V_t(\boldsymbol{s}_{t-1}) = \max_{\boldsymbol{x}_t \in \mathcal{A}_t} \mathbb{E}\left[\sum_{i=1}^m r_t^i \cdot (\min\{D_{t,L}^i, s_{t-1}^i\} + \min\{D_{t,R}^i, (s_{t-1}^i - D_{t,L}^i)^+ + x_t^i\}) + V_{t+1}(\boldsymbol{s}_t)\right] \tag{1b}$$

with state transaction function

$$s_t^i = \begin{cases} (s_0^i + x_1^i - D_1^i)^+ & t = 1 \\ ((s_{t-1}^i - D_{t,L}^i)^+ + x_t^i - D_{t,R}^i)^+ & t = 2, ..., T. \end{cases}$$

and boundary condition $V_{T+1}(s_T) = \sum_{i=1}^m c^i s_T^i$. Hence, solving (1a) will give us a sequence of optimal actions $\{\boldsymbol{x}_t^*\}$ and optimal overall expected revenue V_1^*.

Appendix 2: Preliminary Stochastic Models

In this section, we consider two preliminary models, a multi-location single-period model, and a multi-location two-period model. These models will be used as subproblems for the heuristics defined in Appendix 4.

Multi-location single-period inventory allocation model

We consider a single period model with one DC and m stores for a SKU. For store $i = 1, \cdots, m$, we denote the price by r^i, salvage value by c^i, lead time by L^i, and initial inventory by s_0^i. The total inventory available at the DC is s_0^0. Our decision is to assign x^i units to each store in order to maximize the expected revenue. The revenue maximization problem can be formulated as

$$\max_{x \in \mathbb{N}^m} \mathbb{E}\left[\sum_{i=1}^{m} r^i \cdot (\min\{D_L^i, s_0^i\} + \min\{D_R^i, (s_0^i - D_L^i)^+ + x^i\}) + c^i \cdot ((s_0^i - D_L^i)^+ + x^i - D_R^i)^+\right] \tag{2}$$
$$s.t. \ \sum_{i=1}^{m} x^i \le s_0^0,$$

where D_L^i is the demand occurred during replenishment lead time, and D_R^i is the demand after the replenishment in store i.

The objective function of (2) may be non-concave. As in Gallien et al. (2015), we approximate (2) by allowing backlogging during the lead time, resulting in the following approximation model:

$$\max_{x \in \mathbb{N}^m} \mathbb{E}\left[\sum_{i=1}^{m} r^i \cdot \min\{D^i, s_0^i + x^i\} + c^i \cdot (s_0^i + x^i - D^i)^+\right] \tag{3}$$
$$s.t. \ \sum_{i=1}^{m} x^i \le s_0^0,$$

where $D^i = D_L^i + D_R^i$ is the random demand for store i. Since $r^i \ge c^i$, one can show that the objective function of (3) is concave, and the problem can be solved using Lagrangian relaxation method. The solution to the single-period model (3) is used in the SO policy and in the last period for the MDH policy.

Multi-location two-period inventory allocation model

We consider a two-period model for each SKU with periods $t = 1, 2$. At the beginning of period 1, we have s_0^0 units of inventory in the DC and s_0^i units of inventory in store i. For store i in period t, the selling price is r_t^i, the replenishment quantity is x_t^i and inventory at the end of period t is s_t^i. We denote the inventory available in the DC after the first replenishment by s_1^0. Clearance price at the end of period 2 is c^i. Our decision is to assign quantities x_t^i $(t = 1, 2)$ to each store in each period in order to maximize the expected revenue.

We assume that reorder decisions are made at the beginning of each period, and the items arrive after some lead time. Let $D_{1,L}^i$ be the demand occurred during the first period lead time, and $D_{1,R}^i$ be the demand occurred from the first replenishment to end of first period, and $D_{2,L}^i$, $D_{2,R}^i$ corresponding demand for the second period. The problem is formulated as

$$\max_{x_1 \in \mathbb{N}^m} \mathbb{E}\left[\sum_{i=1}^{m} r_1^i \cdot (\min\{D_{1,L}^i, s_0^i\} + \min\{D_{1,R}^i, (s_0^i - D_{1,L}^i)^+ + x_1^i\}) + V_2(\boldsymbol{s}_1)\right] \tag{4}$$

$$
\begin{aligned}
s.t.\ & \sum_{i=1}^{m} x_1^i \le s_0^0, \\
& s_1^0 = s_0^0 - \sum_{i=1}^{m} x_1^i, \\
& s_1^i = ((s_0^i - D_{1,L}^i)^+ + x_1^i - D_{1,R}^i)^+, \qquad i = 1, ..., m.
\end{aligned}
$$

In the objective function, the first term represents the revenue in the first period. The second term, $V_2(\boldsymbol{s}_1)$, is the optimal expected revenue in the second period given DC inventory s_1^0 and store inventory s_1^i at the end of period 1, defined as

$$
\begin{aligned}
V_2(\boldsymbol{s}_1) = \max_{\boldsymbol{x}_2 \in \mathbb{N}^m} \mathbb{E} \Bigg[\sum_{i=1}^{m} r_2^i \cdot (\min\{D_{2,L}^i, s_1^i\} + \min\{D_{2,R}^i, (s_1^i - D_{2,L}^i)^+ + x_2^i\}) \\
+ c^i((s_1^i - D_{2,L}^i)^+ + x_2^i - D_{2,R}^i)^+ \Bigg] \qquad (5) \\
s.t.\ \sum_{i=1}^{m} x_2^i \le s_1^0.
\end{aligned}
$$

Similar to the single period model, by allowing backlogging during the lead time, we get

$$
\begin{aligned}
\max_{\boldsymbol{x}_1 \in \mathbb{N}^m} \mathbb{E} & \left[\sum_{i=1}^{m} r_1^i \cdot \min\{D_1^i, s_0^i + x_1^i\} + \hat{V}_2(\boldsymbol{s}_1) \right] \qquad (6) \\
s.t.\ & \sum_{i=1}^{m} x_1^i \le s_0^0, \\
& s_1^0 = s_0^0 - \sum_{i=1}^{m} x_1^i, \\
& s_1^i = (s_0^i + x_1^i - D_1^i)^+, \qquad i = 1, ..., m
\end{aligned}
$$

where $D_t^i = D_{t,L}^i + D_{t,R}^i$ represents the random demand for store i in period t, and $\hat{V}_2(\boldsymbol{s}_1)$ is the approximated value-to-go function for period 2:

$$
\begin{aligned}
\hat{V}_2(\boldsymbol{s}_1) = \max_{\boldsymbol{x}_2 \in \mathbb{N}^m} \mathbb{E} & \left[\sum_{i=1}^{m} r_2^i \cdot \min\{D_2^i, s_1^i + x_2^i\} + c^i \cdot (s_1^i + x_2^i - D_2^i)^+ \right] \qquad (7) \\
s.t.\ & \sum_{i=1}^{m} x_2^i \le s_1^0.
\end{aligned}
$$

We propose a mean demand heuristic to solve (6) by making two simplifications. First, we relax the constraints $s_1^i \ge 0$, which means we allow back-order at the end of period 1. Second, we replace the inventory level $s_1^i = s_0^i + x_1^i - D_1^i$ by its expected value, $s_1^i = s_0^i + x_1^i - d_1^i$, where d_1^i is the mean demand for period 1 in store i. Note that we only use mean demand to approximate the inventory level s_1^i, but we still use the actual demand distribution to compute the profits in the first period. These two simplification steps lead to the following problem:

$$
\begin{aligned}
\max_{\boldsymbol{x}_1 \in \mathbb{N}^m} \mathbb{E} & \left[\sum_{i=1}^{m} r_1^i \cdot \min\{D_1^i, s_0^i + x_1^i\} + \hat{V}_2(\boldsymbol{s}_1) \right] \qquad (8) \\
s.t.\ & \sum_{i=1}^{m} x_1^i \le s_0^0,
\end{aligned}
$$

$$s_1^0 = s_0^0 - \sum_{i=1}^{m} x_1^i,$$
$$s_1^i = s_0^i + x_1^i - d_1^i, \qquad i = 1, ..., m.$$

We propose a cutting plane method to solve (8). This method defines a "master" problem that approximates (8) and repeatedly add "cuts" to refine it. Intuitively, starting from unbounded $\hat{V}_2$ in the initial "master" problem, by adding each cut, which is a linear function of s_1, the upper bound of $\hat{V}_2$ is gradually refined until no profit improvement is observed. The reason of applying such cutting plane method is due to the fact that $\hat{V}_2$ is a compound implicit function of x_1. While expliciting the relation between x_1 and $\hat{V}_2$ is essentially impossible, this cutting plane method allows us to approximately characterize this relation using a set of linear constraints, so that (8) is tractable. To be specific, given an initial value of first stage decisions $\boldsymbol{x}_1$, we calculate the inventory level $s_1^i = s_0^i + x_1^i - d_1^i$, then solve the dual problem of the second period problem (7), given by

$$\min_{\alpha,\beta,\rho,\eta,\gamma} \sum_{i=1}^{m} r_2^i \cdot \mathbb{E}[D_2^i] + s_1^0\alpha + \sum_{i=1}^{m} \sum_{d_2^i \in \mathcal{D}_2^i} [(d_2^i - s_1^i)\eta_{d_2^i}^i + (s_1^i - d_2^i)\gamma_{d_2^i}^i] \tag{9}$$
$$\alpha + \sum_{d_2^i \in \mathcal{D}^i} \eta_{d_2^i}^i - \sum_{d_2^i \in \mathcal{D}_2^i} \gamma_{d_2^i}^i \geq 0 \qquad i = 1, ..., m,$$
$$\beta^i \leq -r_2^i, \qquad i = 1, ..., m,$$
$$\rho^i \leq c^i, \qquad i = 1, ..., m,$$
$$-p_{d_2^i}^i \cdot \beta^i + \eta_{d_2^i}^i \geq 0 \qquad i = 1, ..., m, d_2^i \in \mathcal{D}_2^i,$$
$$-p_{d_2^i}^i \cdot \rho^i + \gamma_{d_2^i}^i \geq 0 \qquad i = 1, ..., m, d_2^i \in \mathcal{D}_2^i,$$
$$\alpha \geq 0,$$
$$\eta_{d_2^i}^i \leq 0, \qquad i = 1, ..., m, d_2^i \in \mathcal{D}_2^i.$$
$$\gamma_{d_2^i}^i \leq 0, \qquad i = 1, ..., m, d_2^i \in \mathcal{D}_2^i.$$

where $\mathcal{D}_2^i$ is the support of second period demand. As function $\hat{V}_2(\boldsymbol{s}_1)$ is concave, the dual solution of (9), i.e $(\alpha, \beta, \rho, \eta, \gamma)$, provides a gradient of $\hat{V}_2(\cdot)$ at $\boldsymbol{s}_1$; moreover, the linear function associated with this gradient, i.e.

$$\sum_{i=1}^{m} r_2^i \cdot \mathbb{E}[D_2^i] + s_1^0\alpha + \sum_{i=1}^{m} \sum_{d_2^i \in \mathcal{D}_2^i} [(d_2^i - s_1^i)\eta_{d_2^i}^i + (s_1^i - d_2^i)\gamma_{d_2^i}^i] \tag{10}$$

at given $\boldsymbol{s}_1$, forms an upper bound $\hat{V}_2(\cdot)$. We add linear function (10) (i.e. cut) into the following master problem,

$$\max_{\boldsymbol{x}_1 \in \mathbb{N}^m} \mathbb{E}\left[\sum_{i=1}^{m} r_1^i \cdot \min\{D_1^i, s_0^i + x_1^i\} + \tilde{V}(\boldsymbol{s}_1)\right]$$
$$s.t. \sum_{i=1}^{m} x_1^i \leq s_0^0,$$
$$s_1^0 = s_0^0 - \sum_{i=1}^{m} x_1^i,$$

$$s_1^i = s_0^i + x_1^i - d_1^i, \qquad i = 1, ..., m,$$
$$\tilde{V}(\boldsymbol{s}_1) \leq \text{all cuts generated by (9).}$$

We repeat this process and continue adding cuttings until the solution to the master problem converges.

Appendix 3: Demand Learning Methods

Maximum likelihood estimation (MLE) for Poisson distributions

Suppose at a given price level, the aggregated normalized weekly demand in a given cluster follows a Poisson distribution with parameter θ. Then for the ith SKU in jth week on the kth day in the sth store, the daily demand follows a Poisson distribution with rate $\theta_{ijks} = \theta \cdot p_i \cdot SF_j \cdot \eta_k \cdot \phi_s$, where p_i is the SKU proportion for a given style, SF_j is the weekly seasonality factor, η_k is the in-week seasonality factor, and ϕ_s is the store proportion. We assume that weekly seasonality factors SF_j and in-week seasonality factors η_k are known for a given product category. In our implementation, they is calculated from historical data using sales fractions on an aggregated product category. An example of in-week seasonality factor is shown in Table 2. Similarly, we assume that ϕ_s can be calculated from historical data. In addition, we assume the SKU proportions are known or can be properly represented by the proportion of SKU sales among aggregated product sales over all stores during the entire selling season. Note that these assumptions are made to simplify the demand learning process, since we may focus on estimating a single number θ, which is treated as an unknown parameter in the MDHL algorithm. It is also consistent with Dillard's current practice.

Table 2 Fraction of sales by the day of the week for a sample product.

Day	Sun.	Mon.	Tue.	Wed.	Thu.	Fri.	Sat.
Sales Fraction (%)	11.9	10.9	11.1	12.8	13.2	17.6	22.5

Let y_{ijks} be the sales and let o_{ijks} be the stock-out indicator of ith SKU in jth week on the kth day for the sth store, e.g. $o_{ijks} = 0$ if in stock and $o_{ijk} = 1$ if out-of-stock at the beginning of the day. Then for a cluster of S stores with I SKU's and J weeks, we can formulate the likelihood function as

$$L(\theta) = \prod_{i=1}^{I}\prod_{j=1}^{J}\prod_{k=1}^{7}\prod_{s=1}^{S}\left[\frac{(\theta \cdot p_i \cdot SF_j \cdot \eta_k \cdot \phi_s)^{y_{ijks}}\, e^{-\theta \cdot p_i \cdot SF_j \cdot \eta_k \cdot \phi_s}}{y_{ijks}!}\right]^{1-o_{ijks}},$$

and its logarithm is

$$\ell(\theta) = \log\theta\left[\sum_{i=1}^{I}\sum_{j=1}^{J}\sum_{k=1}^{7}\sum_{s=1}^{S}(1-o_{ijks})y_{ijks}\right] - \theta\left[\sum_{i=1}^{I}\sum_{j=1}^{J}\sum_{k=1}^{7}\sum_{s=1}^{S}(1-o_{ijks})p_i \cdot SF_j \cdot \eta_k \cdot \phi_s\right]$$
$$+\sum_{i=1}^{I}\sum_{j=1}^{J}\sum_{k=1}^{7}\sum_{s=1}^{S}(1-o_{ijks})\left[y_{ijks}\log(p_i \cdot SF_j \cdot \eta_k \cdot \phi_s) - \log(y_{ijks}!)\right].$$

We can derive the MLE of θ by setting the first order derivative of $\ell(\theta)$ to zero,

$$\hat{\theta} = \frac{\sum_{i=1}^{I}\sum_{j=1}^{J}\sum_{k=1}^{7}\sum_{s=1}^{S}(1-o_{ijks})y_{ijks}}{\sum_{i=1}^{I}\sum_{j=1}^{J}\sum_{k=1}^{7}\sum_{s=1}^{S}(1-o_{ijks})p_i \cdot SF_j \cdot \eta_k \cdot \phi_s},$$

so $\hat{\theta}_{ijks} = \hat{\theta} \cdot p_i \cdot SF_j \cdot \eta_k$, and the max likelihood is $L(\hat{\theta})$.

Regularized Poisson regression

Unlike the MLE method given above, here we assume that the mean demand of a SKU on a given date in a given store can be represented using an exponential function. If this product has K SKU's and the cluster has S stores,

$$\lambda(x) = e^{\beta_0 + \beta_1 x_1 + \beta_2 x_2 + \sum_{i=3}^{2+(K-1)} \beta_i x_i + \sum_{j=K+2}^{K+2+(S-1)} \beta_j x_j},$$

where x is a feature vector that describes the day, week, SKU and store; x_1 is the in-week seasonality factor, x_2 the weekly seasonality factor. The $K-1$ following indicator variables $x_3, x_4, ..., x_{2+(K-1)}$ are used to indicate the SKU, and the subsequent $S-1$ indicator variables $x_{K+2}, x_{K+3}, ..., x_{K+2+(S-1)}$ are used to indicate the store. To train parameters β for a cluster at a price level, we use the data set that belongs to this cluster and price; for a data set with N historical records (where we only use data points with positive inventory at the beginning of the day), we can formulate the likelihood function as

$$L(\beta) = \prod_{n=1}^{N} \frac{\left(e^{\beta^\top x^n}\right)^{y^n} e^{-e^{\beta^\top x^n}}}{y^n!},$$

and its logarithm as

$$\ell(\beta) = \sum_{n=1}^{N} \left[y^n \beta^\top x^n - e^{\beta^\top x^n} - \log(y^n!) \right].$$

We solve a Lasso version of $\ell(\beta)$ with a regularization term to avoid over fitting, and we maximize the log-likelihood with a regularization penalty α,

$$\max_{\beta} \ell(\beta) + \alpha \|\beta\|.$$

Python has an open source package for Poisson regression in *statsmodels.discrete.discrete_model.Poisson*, and details about this package can be found at `http://www.statsmodels.org/dev/generated/statsmodels.discrete.discrete_model.Poisson.html`. After we obtain β, we can derive a prediction for a particular SKU/store/day by plugging in the corresponding feature vector into λ.

Weighted regularized Poisson regression

We use a weighted version of the regularized regression model above as a part of the MDHL heuristic policy. In a similar fashion, suppose we have a historical data set with N_1 records and current season sales data with N_2 records, and we would like to give higher weight to the current season; specifically, we apply weight $w > 1$ to the N_2 records from the current season. Then we solve the problem

$$\max_{\beta} \sum_{n=1}^{N_1} \left[y^n \beta^\top x^n - e^{\beta^\top x^n} - \log(y^n!) \right] + w \sum_{n=1}^{N_2} \left[y^n \beta^\top x^n - e^{\beta^\top x^n} - \log(y^n!) \right] + \alpha \|\beta\|.$$

When w is integral, we can use w copies of the current sales data and combine this with the historical set to get $N_1 + wN_2$ records. We can use the same Python package, *statsmodels.discrete.discrete_model.Poisson*, to obtain an estimate of β.

Appendix 4: Inventory Management Policies

In this section, we define Ship-Once (SO), Mean Demand Heuristic (MDH), and Mean Demand Heuristic with Learning (MDHL).

To describe these policies, we define the following notation. Consider a product with K SKU's. The selling season of this product contains T weeks. Let S be the number of stores. For each $k \in [K]$, $t \in [T]$, and $s \in [S]$, let d^k_{ts} be the rate of Poisson demand distribution for SKU k in week t in store s, and let r^k_{ts} be the corresponding retail price. Let s^k_{DC} be the inventory of SKU k at the DC when the season begins.

Ship-Once (SO) Policy

The SO policy is detailed in Algorithm 1. It uses the multi-location single-period model (3) described in Appendix 2.

Algorithm 1 SO policy

Repeat the following steps for $k = 1, ..., K$:

1: Compute aggregated store demand rates for the T weeks: $d^k_s = \sum_{t=1}^{T} d^k_{ts}$;

2: Compute average store prices for the T weeks, weighted by the weekly demand rates: $r^k_s = \frac{\sum_{t=1}^{T} r^k_{ts} d^k_{ts}}{d^k_s}$;

3: Calculate allocation decision $x^k = \{x^k_s : s \in [S]\}$ for SKU k by solving the multi-location single-period model, using the prices and demands computed in Steps 1 and 2, DC inventory s^k_{DC}, and clearance price c.

Mean Demand Heuristic (MDH) Policy

Let s^k_{ts} be the inventory level of SKU k in store s at the start of week k. The MDH policy is detailed in Algorithm 2. MDH uses the single-period model and the two-period model described in Appendix 2.

Mean Demand Heuristic with Learning (MDHL) policy

The MDHL policy is detailed in Algorithm 3. MDHL use the multi-location single-period model and the two-period model in Appendix 2, and the regularized weighted Poisson regression described in Appendix 3.

Appendix 5: Preliminary Tests Details

Data Description

We use a data set from a region with one distribution center (DC) serving 41 retail stores, and covering nine fashion products. Each product has seven to nine SKU's, representing different sizes of the product. Our data include the following information:

- retail prices posted on each day;
- daily sales volume and inventory information at the SKU level;
- weekly seasonality factors, which account for sales trends during the season.

Algorithm 2 MDH policy

Repeat the following steps for $k = 1, ..., K$:

1: Set s^{k1} as the zero vector; set s_{DC}^{k1} as the initial inventory level at the DC.
2: **for** $i = 1, ..., T$, **do**
3:     **if** $i < T$ **then**
4:         Compute period-1 and period-2 store demand rates, $d_s^{k1} = d_{is}^k$, $d_s^{k2} = \sum_{t=i+1}^T d_{ts}^k$;
5:         Compute period-1 and period-2 average retail price, $r_s^{k1} = r_{is}^k$, $r_s^{k2} = \frac{\sum_{t=i}^T r_{ts}^k d_{ts}^k}{\sum_{t=i}^T d_{ts}^k}$;
6:         Calculate allocation decision $x^k = \{x_s^k : s \in [S]\}$ for SKU k by solving the multi-location two-period model, with price vectors r^{k1} and r^{k2} from Step 5, demand rate vectors d^{k1} and d^{k2} from Step 4, initial store inventory vector s^{ki}, DC inventory s_{DC}^{ki}, and clearance price c.
7:         Observe realized sales volume z_{is}^k in each store; update store starting inventory $s_s^{k,i+1} = s_s^{ki} - z_{is}^k$; update DC inventory $s_{DC}^{k,i+1} = s_{DC}^{ki} - \sum_{s=1}^S x_s^k$.
8:     **else**
9:         Calculate allocation decision $x^k = \{x_s^k : s \in [S]\}$ for SKU k by solving multi-location single-period model, with price vector r_T^k, demand rate vector d_T^k, initial inventory vector s^{kT}, DC inventory s_{DC}^{kT}, and clearance price c.
10:     **end if**
11: **end for**

During the 2016-2017 season, the products were offered in these stores at their full prices in the first 10 weeks; then, the first markdown prices (at 35% to 40% off) were applied in the next 4 weeks. After that, the second markdown prices (at roughly 65% off) were applied to the remaining weeks for clearance. In the data, we observe that a majority of the stores had been actively replenished until the 14th week. Therefore, we assume that all the products have a selling season consisting of 14 weeks.

Demand Model Estimation

As with many fashion goods, for the products in this study we observe extremely low average daily sales per SKU and store. Therefore, estimating demand models for each separate SKU is difficult. To tackle this challenge, we aggregate sales data by *week* and by *product*. We generate a probabilistic demand model using regularized Poisson regression, as described in the previous section. We then disaggregate the model to obtain a demand forecast model for each *day*, *SKU*, and *store*.

Testing Inventory Policies with Known Demand Distributions

We first conduct numerical tests for a situation where the true demand distribution is known to the decision maker. We generate a demand model using regularized Poisson regression, and sample testing scenarios from this model for calculating expected revenues attained by different inventory policies. The test was performed on a personal laptop with a 1.80 GHz quad-core CPU. On average, the SO policy requires about six seconds to compute inventory allocations for each product (including all SKU's); recall that the SO policy only makes inventory allocation decisions at the beginning of the season, so it requires less computation.

Algorithm 3 MDHL policy

For $k \in [K]$, Initialize s^{k1} as zero vector, s^{k1}_{DC} be the initial DC inventory, set d^k_{ts} as initial demand rate estimation, history reference set is $\mathcal{H}'$, new season information set $\mathcal{H} = \emptyset$.

for $i = 1, ..., T$, **do**

  for $k = 1, ..., K$, **do**

    if $i < T$ **then**

      Compute period-1 and period-2 store demand rates, $d^{k1}_s = d^k_{is}$, $d^{k2}_s = \sum_{t=i}^T d^k_{ts}$;

      Compute period-1 and period-2 store retail price, $r^{k1}_s = r^k_{is}$, $r^{k2}_s = \frac{\sum_{t=i}^T r^k_{ts} d^k_{ts}}{\sum_{t=i}^T d^k_{ts}}$;

      Calculate allocation decision $x^k = \{x^k_s : s \in [S]\}$ for SKU k by solving multi-location two-period model, with price vector r^{k1} and r^{k2}, demand rate vector d^{k1} and d^{k2}, initial store inventory vector s^{ki}, and s^{ki}_{DC} DC inventory and clearance price c.

      Observe sales z^k_{is} in each store, update store starting inventory $s^{k,i+1}_s = s^{ki}_s - z^k_{is}$, update DC inventory $s^{k,i+1}_{DC} = s^{ki}_{DC} - \sum_{s=1}^S x^k_s$.

    else

      Calculate allocation decision $x^k = \{x^k_s : s \in [S]\}$ for SKU k by solving multi-location single-period model, with price vector r^k_T, demand rate vector d^k_T, initial inventory vector s^{kT}, and s^{kT}_{DC} DC inventory and clearance price c.

    end if

    collect new sales data $\mathcal{I}^k_i$;

  end for

  Update new season information set $\mathcal{H} = \mathcal{H} \cap \{\cap_{k \in [K]} \mathcal{I}^k_i\}$;

  Update demand rates using history reference set $\mathcal{H}'$ and new season information set $\mathcal{H}$.

end for

The MDH policy requires one to two minutes to compute inventory allocations for each product (including all SKU's) per week, and the decisions are made in a rolling horizon for the 14-week season.

A comparison of expected revenues achieved by the SO and MDH policies is presented in Figure 7. The left and right boxplots correspond to average revenue achieved by the SO and MDH policies, respectively. The sample mean within each box is represented by a triangle. The blue lines are upper bounds on the maximum expected revenue that can be achieved by any policy; this upper bound is calculated as the expected revenue that could be attained if all stores' demand realized at the DC, so there wouldn't be any lost sales caused by allocating inventory to the wrong locations.

Testing Inventory Policies with Unknown Demand Distributions

We next report the results of more realistic tests in which the demand forecast available at the beginning of the season is inaccurate. We treat the demand model estimated from Dillard's sales data as the retailer's forecast, and consider three scenarios:

1. The actual demand rates can be either higher or lower than the forecast.

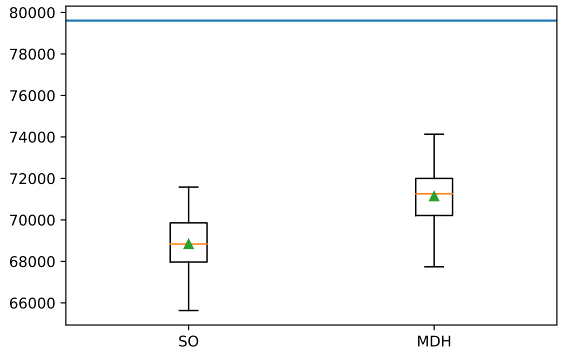

(a) Product: 1

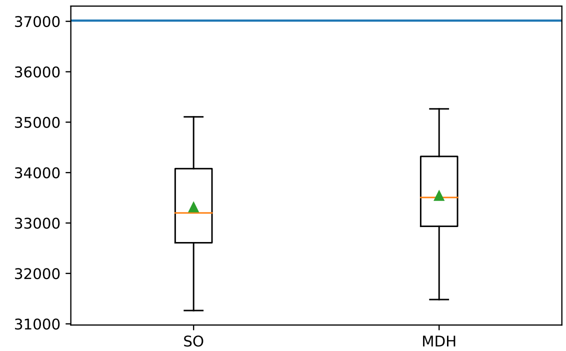

(b) Product: 2

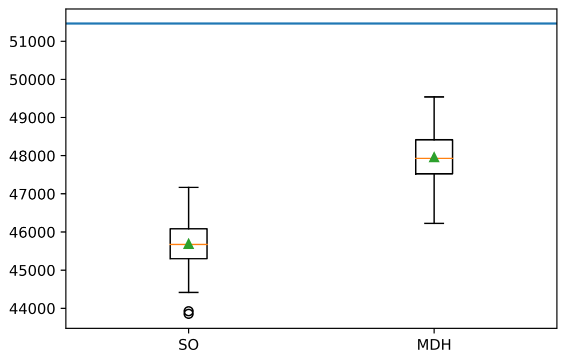

(c) Product: 3

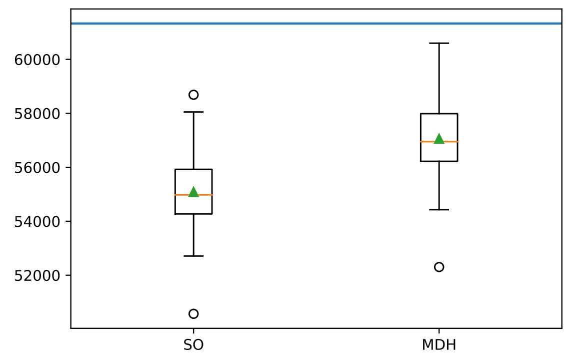

(d) Product: 4

Figure 7 Performance of SO and MDH policies with known demand distributions.

2. The actual demand rates are higher than the forecast for all stores. This represents a scenario where a product is more popular than expected.
3. The actual demand rates are lower than the forecast for all stores. This represents a scenario where a product is less popular than expected.

In our experiments, when the true demand is lower than expected, it lies between 50% and 100% of the forecast value. When the true demand is higher than expected, it lies between 100% and 200% of the forecast value.

We test four inventory control policies, SO, MDH, and MDHL, running experiments on a laptop with a 1.80 GHz quad-core CPU. On average, the SO policy requires about 6 seconds to compute inventory allocations for each product (including all SKU's). The MDH and MDHL policies require 1-2 minutes to compute inventory allocations for each product per week. For the MDHL policy, the demand learning step takes less than 10 seconds per period, so their computation time is roughly the same as the MDH policy, which does not involve demand learning.

In Figure 8, we present the test results for two representative products, one with relatively high demand, and a second with relatively low demand, based on sampled demand scenarios. Each row of the figure displays results for one of the three cases. In addition to showing boxplots for the four heuristic policies, we consider a benchmark given by the MDH policy using the true demand model, denoted by MDHt. Since the

MDHt policy uses the true demand distribution, it serves as a benchmark for all three MDH based policies (MDH and MDHL), which do not have access to the true demand distribution. As before, the blue lines in the figures are upper bounds on the maximum expected revenue, assuming all demand realizes at the DC.

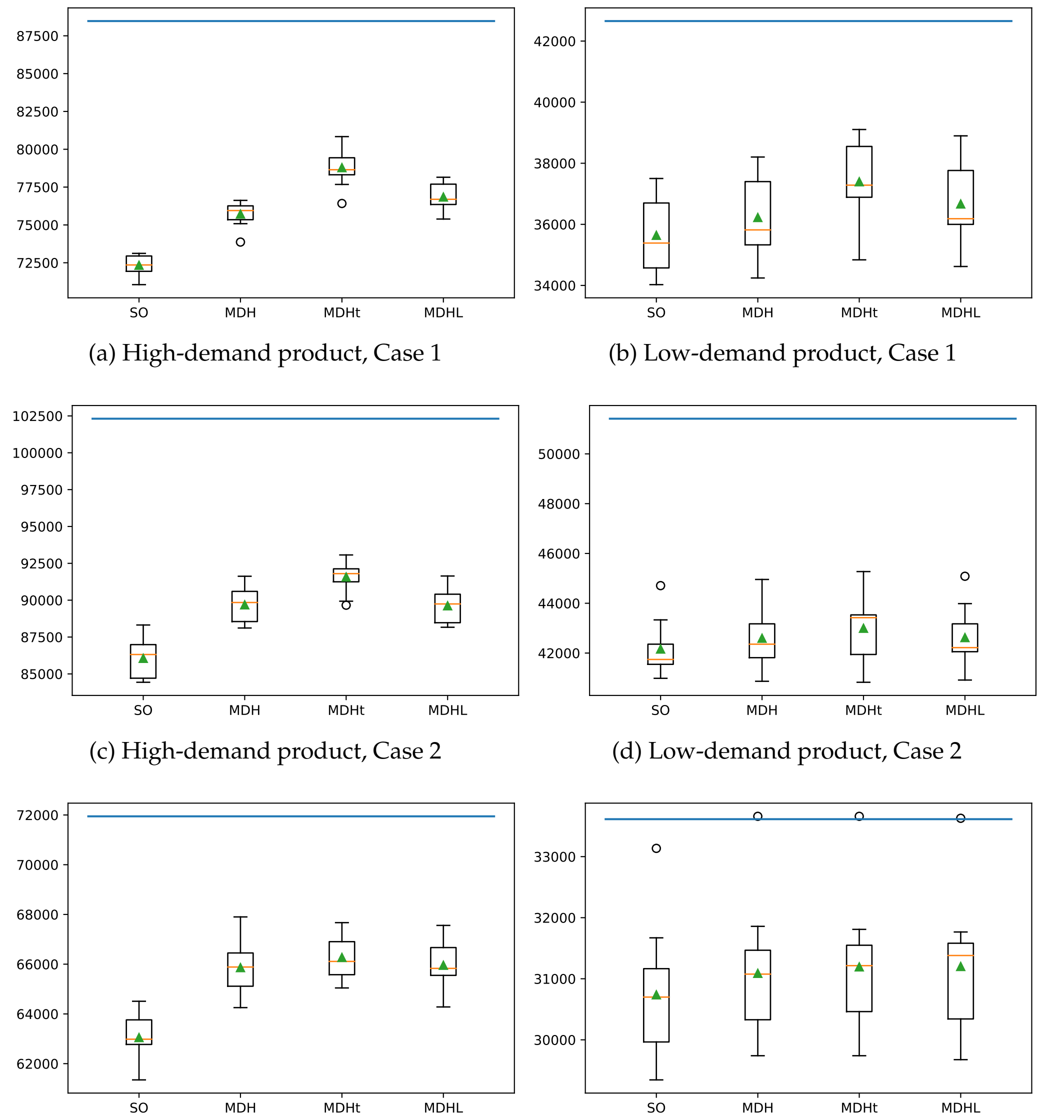

(a) High-demand product, Case 1

(b) Low-demand product, Case 1

(c) High-demand product, Case 2

(d) Low-demand product, Case 2

(e) High-demand product, Case 3

(f) Low-demand product, Case 3

Figure 8 **Test results for experiments with inaccurate forecasts.**

We make the following observations from these results. First, the MDHL policy has the best performance among the three heuristic policies in most cases.

Second, MDHL provides the most benefit over other policies in case 1, where the initial forecasts overestimate demand at some locations and underestimate it in others. This is intuitive, because having a mixture of underestimation and overestimation is where learning demand and adjusting inventory allocations dynamically can provide the largest benefit. In contrast, when the true demand rates are lower than the forecast in all locations (Case 3), MDHL affords relatively small benefit compared to the other policies. In this case, overestimating demand leads to overstocking; since there is sufficient inventory to satisfy all customer demand, the revenue is insensitive to how inventory is allocated among stores.