Analytics for an Online Retailer: Demand Forecasting and Price Optimization

Kris Johnson
Operations Research Center, Massachusetts Institute of Technology, krisdj@mit.edu

Bin Hong Alex Lee
Engineering Systems Division, Massachusetts Institute of Technology, binhong@mit.edu

Murali Narayanaswamy
VP Pricing and Operations Strategy, Rue La La, mnarayanaswamy@ruelala.com

Philip Roizin
Chief Financial Officer, Rue La La, proizin@ruelala.com

David Simchi-Levi
Engineering Systems Division, Department of Civil & Environmental Engineering and the Operations Research Center, Massachusetts Institute of Technology, dslevi@mit.edu

Jonathan Waggoner
Chief Operating Officer, Rue La La, jwaggoner@ruelala.com

Company Background & Motivation

We present our work with an online retailer, Rue La La, as an example of how a retailer can use its wealth of data to optimize pricing decisions on a daily basis. Rue La La is in the online fashion sample sales industry, where they offer extremely limited-time discounts ("flash sales") on designer apparel and accessories. According to IBISWorld (2012), this industry emerged in the mid-2000s and by 2012 was worth approximately 2 billion USD, benefiting from an annual industry growth of approximately 50% over the last 5 years; some of Rue La La's competitors in this industry include companies such as Gilt Groupe, HauteLook, and Beyond the Rack. Many companies in the industry also have brick and mortar stores, whereas others like Rue La La only sell products online.

Upon visiting Rue La La's website (www.ruelala.com), the customer sees several "events", each representing a collection of for-sale products that are similar in some way. For example, one event might represent a collection of products from the same designer, whereas another event might represent a collection of men's sweaters. At the bottom of each event, there is a countdown timer informing the customer of the time remaining until the event is no longer available; events typically last between 1-4 days.

When a customer sees an event he is interested in, he can click on the event which takes him to a new page that shows all of the products for sale in that event; each product on this page is referred to as a "style". Finally, if the customer likes a particular style, he may click on the style which takes him to a new page that displays detailed information about the style, including which sizes are available; we will refer to a size-specific product as an "item". The price for each item is set at the style level, where a style is essentially an aggregation of all sizes of otherwise identical items. The price does not change throughout the duration of the event given the short event length.

With regards to Rue La La's operations, merchants procure items from designers who typically ship the items immediately to Rue La La's warehouse. On a frequent periodic basis, merchants identify opportunities for future events based on available styles in inventory, customer needs, etc. When the event starts, customers place orders, and Rue La La ships items from its warehouse to the customers. When the event ends or an item runs out of inventory, customers may no longer place an order for that item. If there is remaining inventory at the end of the event, then the merchants will plan a subsequent event where they will sell the same style. We will refer to styles being sold for the first time as "first exposure styles"; a majority of Rue La La's revenue comes from first exposure styles, and hundreds of first exposure styles are offered on a daily basis.

One of Rue La La's main challenges is pricing and predicting demand for these first exposure styles. Figure 1 shows a histogram of the sell-through (% of inventory sold) distribution for first exposure items in Rue La La's top 5 departments (with respect to quantity sold). For example,

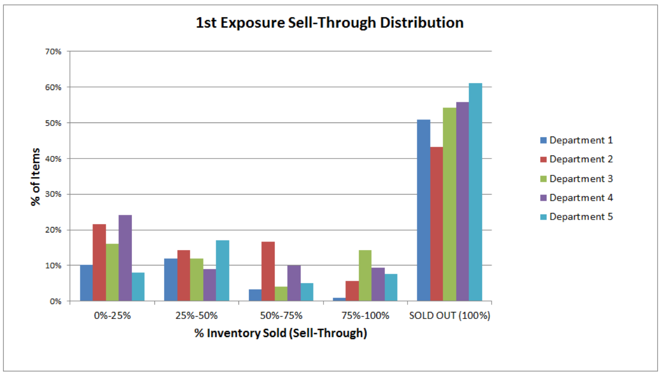

Figure 1 **First exposure sell-through distribution by department**

51% of first exposure items in Department 1 sell out before the end of the event, and 10% sell less than 25% of their inventory. Department names are hidden and data disguised in order to protect confidentiality. A large percent of first exposure items sell out before the sales period is over, suggesting that it may be possible to raise prices on these items while still achieving high sell-through; on the other hand, many first exposure items sell less than half of their inventory by the end of the sales period, suggesting that the price may have been too high. These observations motivate the development of a pricing decision support tool, allowing Rue La La to take advantage of available data in order to maximize revenue from first exposure sales.

Solution Approach

Our approach is two-fold and begins with using machine learning techniques to develop a demand prediction model for first exposure items; we then use this demand prediction data as input into a price optimization model to maximize revenue. Previously, Rue La La set initial prices based on some combination of the following criteria: percentage markup on cost, competitors' pricing, and the merchants' judgement/feel for the best price of the product. We show that applying machine learning and optimization techniques to these initial pricing decisions - while maintaining Rue La La's value proposition of giving their customers the best deal in the market - can have a huge financial impact on the company.

The first challenge we faced when building our demand prediction model was estimating lost sales due to stockouts. Although sales quantity is a natural choice for demand, it does not always represent true demand because of potential lost sales due to stockouts. As Figure 1 illustrates, a large percent of Rue La La's first exposure items sell out before the end of the event, thus the issue of lost sales is frequent in this setting. In fact, many items sell out within just the first few hours of an event. Our estimation method uses sales data from items that did not sell out (i.e. when sales = demand) to predict lost sales of items that did sell out. We used clustering techniques to identify groups of items with similar sales patterns.

The second challenge we faced was predicting demand for items that have no historical sales data. To do this, we developed a set of quantitative attributes that describe product, event, and assortment characteristics. We used these attributes and historical sales data (of other products) to build a variety of regression models to predict demand. We tested the performance of common models such as least squares regression, power/multiplicative regression, and semilogarithmic regression. One reason for the popularity of these demand functions as an input to price optimization is their set of properties, such as linearity, concavity and increasing differences, that leads to simpler, tractable optimization problems that can provide managerial insights. In addition to the common demand prediction models, we also tested several other regression models not typically used for demand prediction, such as regression trees.

Surprisingly, regression trees - an intuitive, yet non-parametric regression model - proved to be the best predictors of demand in terms of both predictability and interpretability. The obvious benefit of using regression trees is that they do not require specification of a certain functional, parametric form between regressors and demand; the model is more general than the common demand prediction models in this sense. In some respect, regression trees are able to determine - for each new style to be priced - the key characteristics of that style that will best predict demand, and they use the demand of styles sold in the past that also had those same key characteristics as an estimate of future demand. While this is effective in predicting demand, unfortunately this non-parametric structure leads to a more difficult price optimization problem.

We then formulate a price optimization model to maximize revenue from first exposure styles, using demand predictions from the regression trees as inputs. In this case, the biggest challenge we face is that each style's demand depends on the price of similar "competing" styles in the event, which restricts us from solving a price optimization problem individually for each style and leads to an exponential number of variables. Furthermore, the non-parametric structure of regression trees makes this problem particularly difficult to solve. We develop a novel integer programming reformulation of the price optimization problem and exploit its structure to prove a tight bound on its linear programming relaxation. Interestingly, this bound is independent of problem size and depends only on the expected revenue of a single style! We use this to create and implement an efficient algorithm that allows Rue La La to optimize prices on a daily basis.

Results & Impact

To implement our price optimization algorithm, we developed and implemented a fully-automated pricing decision support tool at Rue La La. It is run automatically every day, providing price recommendations to merchants for events starting the next day. The entire pricing decision support tool is depicted in the architecture diagram in Figure 2. On average, the entire tool takes

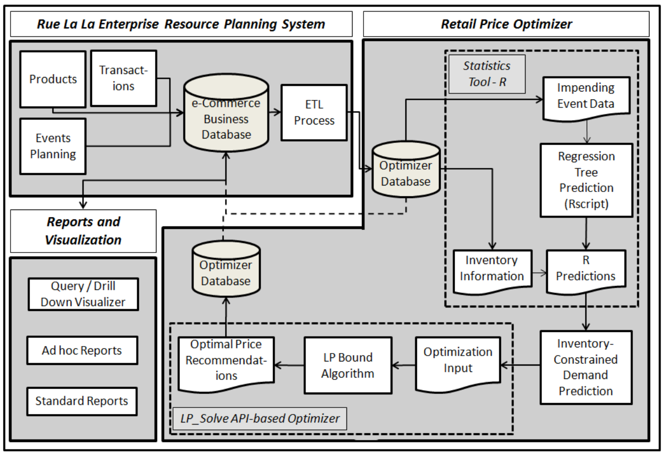

Figure 2 Architecture of pricing decision support tool

under one hour to price a day's worth of styles; the longest run-time we have encountered is 4.5 hours. These are reasonable run-times given Rue La La is running this tool daily.

Being able to estimate the tool's impact prior to implementation was key in gaining buy-in and approval from Rue La La executives to use the pricing decision support tool. We first share our historical analysis that shows an expected increase in first exposure styles' revenue of approximately 11%; then we present results from recent live tests that show a similar impact on revenue.

For our historical tests, we used data from 2011 and 2012 to build regression trees for our demand prediction model, which we then used to predict inventory-constrained demand for items sold in the first two quarters of 2013 *given the actual price of each style*. Inventory-constrained demand is simply the portion of demand that we are able to serve given the amount of the style's on-hand inventory. We compared the predicted inventory-constrained demand with the actual quantity sold, as well as the predicted revenue (predicted inventory-constrained demand $*$ price) with the actual revenue. This comparison is made to understand the accuracy of the demand prediction model and thus entire tool, since demand is the source of uncertainty. Figure 3 shows a summary of results for Rue La La's top 5 departments. Overall the demand prediction model is quite effective within each department, with all aggregate deviations within $+/-$ 8%. Although this chart only shows aggregate deviations, we also evaluated prediction accuracy for each style as part of our analysis.

Given the accuracy of our demand prediction model, the next step was to estimate the impact of our price optimization tool on revenue. For this purpose, we applied our predicted inventory-constrained demand and revenue corresponding to actual prices as our baseline. These two values (baseline inventory-constrained demand and revenue) were then compared to optimized results of our tool, that is, to inventory-constrained demand and revenue associated with the optimal prices suggested by the tool. See Figure 4 for a summary of these results for the same 5 departments.

This figure shows that the optimized prices found by the pricing decision support tool increase revenue by 10-13% across all 5 departments compared to the baseline; when extending to all

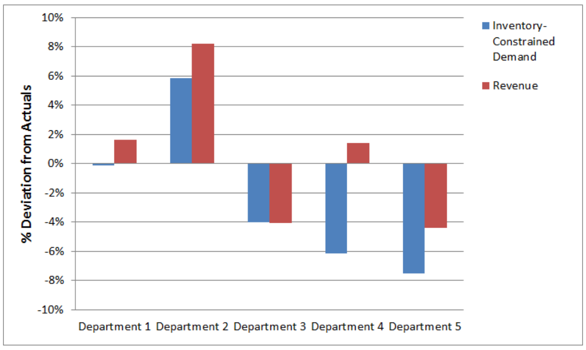

Figure 3 **Comparison of demand prediction model vs. actual sales**

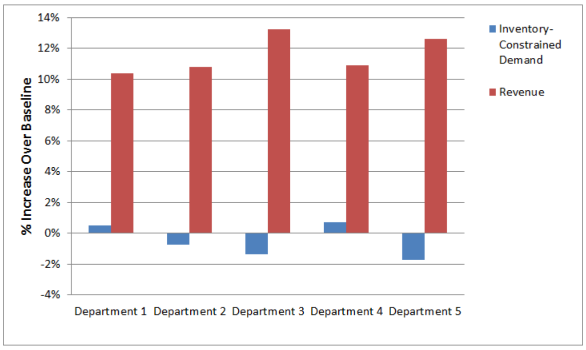

Figure 4 **Comparison of optimized results vs. baseline**

departments where the tool was implemented, the average revenue increase was 11%. Interestingly, the aggregate change in demand in order to achieve these results is very small. Upon further investigation of the data, it appears that there are two key contributing factors to this. First, in the optimized solution, prices are often set such that (*i*) demand increases for items with predicted low sell-through (percent of inventory sold) in the baseline, and (*ii*) demand decreases for items with predicted high sell-through in the baseline. Intuitively, the model tries to shift some of the demand from popular items that are likely to sell out of many sizes to less popular items.

The second key factor for the fairly stable aggregate demand is that in many cases, inventory-constrained demand does not fluctuate much with the price set considered for each style. One reason for this is that a change in true demand corresponds to a smaller change in inventory-constrained demand especially when limited sizes or inventory are available. Another potential reason is that since Rue La La offers very deep discounts, they may already be well below customers' reservation prices (i.e. customers' "willingness to pay"), such that small changes in price are still perceived as great deals. Thus the pricing model finds these opportunities to slightly increase prices without significantly affecting demand; in these cases, consumers still benefit from a great sale, and Rue La La is also able to maintain a healthy business.

Because of these positive results, Rue La La has recently implemented the tool, and we have concluded the first round of live tests that took place from mid-January through March 2014. Our

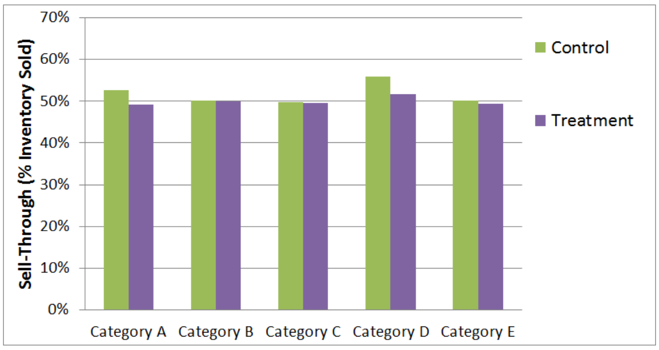

Figure 5 **Live test results: impact on sell-through**

goal for the live tests was to address two questions: (*i*) would the price increases dramatically affect demand, and (*ii*) what impact would the tool's recommended price increases have on revenue?

We identified approximately 2,500 styles where the pricing tool recommended price increases, and we raised prices on approximately 650 styles (our "treatment group") and kept prices the same for approximately 1,900 styles (our "control group"). Styles in both the treatment and control groups were divided into 5 categories based on the actual price of the style; specifically, all styles in the same category have a similar price. Furthermore, we chose styles to ensure that the set of styles in the treatment group vs. control group for each of the 5 categories had a similar product mix in terms of characteristics such as predicted sell-through (using original prices), event type, inventory, average price of competing styles, and department. By first dividing into categories based on actual price, many of these product mix characteristics (such as average price of competing styles) were naturally similar between the treatment and control groups.

To answer our first question regarding how the price increases affect demand, we calculated the average sell-through of styles in our treatment group and compared this to the average sell-through of styles in the control group for each of the 5 categories. Figure 5 illustrates these results. For each of the 5 categories, the average sell-through of the treatment group was within 4% of the average sell-through of the control group. Thus the results of these live tests suggest that increasing prices does not dramatically affect demand, which is important to Rue La La's business model of maintaining a certain level of scarcity of its styles.

With very little change in sell-through between the treatment and control groups, the treatment group has a big impact on revenue. Figure 6 shows a conservative estimate of the impact on revenue from the treatment group, where we assumed the price increases would decrease demand by 5% for each style. Note that the categories are ordered by their associated range of actual prices; Category A is for the least expensive products, whereas Category E is for the most expensive products. Thus, although the *percent* revenue impact is generally decreasing with higher priced categories, the *dollar* impact on revenue remains quite high. For confidentiality reasons, we chose not to present the dollar impact on revenue.

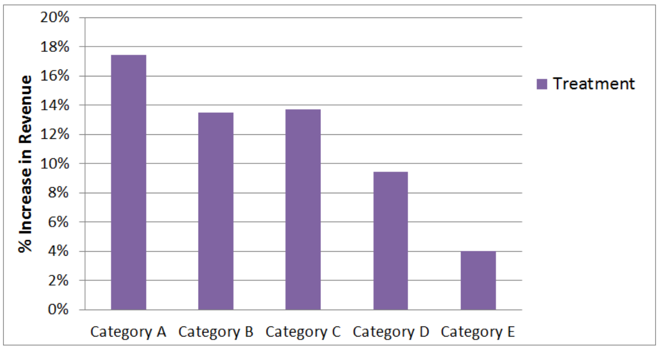

Figure 6 **Live test results: impact on revenue**

Overall, we conservatively estimate a 10% increase in revenue from using the price recommendations from our tool on the treatment group. If we assume that the price increases have no impact on demand (i.e. no impact on sell-through), then our results show a 16% increase in revenue from the treatment group.

These live test results are very promising and are similar to our historical test results; they suggest that the pricing decision support tool is able to identify styles where increasing the price will not significantly decrease demand, resulting in an increase in revenue dollars that drops straight to the bottom line without fundamentally changing Rue La La's business. Because of this, we are currently using the pricing decision support tool to make price recommendations on hundreds of new styles every day.

References

IBISWorld. 2012. Online Fashion Sample Sales in the US: Market Research Report. http://www.ibisworld.com/industry/online-fashion-sample-sales.html (accessed Jan 22, 2014).