Pricing Analytics for Rotable Spare Parts

Omar Besbes
Columbia Business School, New York, New York 10027, ob2105@gsb.columbia.edu

Adam N. Elmachtoub
Department of Industrial Engineering and Operations Research & Data Science Institute, Columbia University, New York, New York 10027, adam@ieor.columbia.edu

Yunjie Sun
Department of Industrial Engineering and Operations Research, Columbia University, New York, New York 10027, ys2888@columbia.edu

In this paper, we describe a comprehensive approach to pricing analytics for reusable resources in the context of rotable spare parts, which are parts that can be repeatedly repaired and resold. Working in collaboration with a major aircraft manufacturer, we aim to instill a new pricing culture and develop a scalable new pricing methodology. Pricing rotable spare parts presents unique challenges ranging from complex inventory dynamics and minimal demand information to limited data availability. We develop a novel pricing analytics approach that tackles all of these challenges and that can be applied across all rotable spare parts. We then describe a large-scale implementation of our approach with our industrial partner, which led to an improvement in profits of over 3.9% over a 10-month period.

Key words: price optimization, reusable resources, rotable spare parts, large-scale implementation

Introduction

In this work, we focus on the engineering and implementation of a systematic pricing approach for thousands of rotable spare parts at a major aircraft original equipment manufacturer (OEM). (For confidentiality reasons, we refer to our industrial partner as 'the OEM.') In addition to manufacturing, a significant part of the business is associated with ensuring a high quality of customer service, which includes the ability to provide spare parts for all operating aircraft. (Note that aircraft is both singular and plural.) The management of spare parts is increasingly critical as the age of the fleet increases and the number of aircraft no longer under warranty increases. Due to the increasing opportunities for spare part sales, the availability of spare parts by competitor service providers has increased. Thus, pricing spare parts optimally, while maintaining a high level of service, has become progressively more challenging and paramount over the years.

The OEM faces the challenge of pricing thousands of different kinds of spare parts. Such pricing decisions present unique challenges, especially for an important subclass of parts called *rotable spare parts*, for which the selling process is quite different than for regular spare parts. Rotable spare parts, which is the subject of this work, are parts that can be repeatedly repaired and resold. Rotable spare parts represent a majority of spare parts sales, with price tags of up to hundreds of thousands of dollars per unit. When purchasing a rotable spare part, a customer will give their broken unit to the OEM in exchange for a functional unit. This swap of their broken part for a functional part is known as an *exchange sale*. The OEM will then send the broken unit to a repair facility. When the broken part is repaired and sent back, it is then placed back into the OEM's inventory and can be sold again. We remark that rotable spare parts are an example of a reusable resource, which is the subject of increased attention due to the increasing popularity of reusable resources such as ride sharing vehicles and cloud computing systems.

One key characteristic in selling rotable spare parts is that the total number of units is fixed for the OEM. At any given time, these units are classified into two groups: the on hand units, which are available for sale and the in-repair units, which are not available. Moreover, because there are many competitors in the market, when a customer asks for a rotable spare part but the OEM has no available units, the sale will almost certainly be lost to a competitor. Note that substituting one part number for another is not physically possible. Although dynamic pricing policies (i.e., a policy where the price changes depending on the number of available units) are natural to consider and were our original intent, we actually utilize a static pricing policy because it is near optimal empirically and theoretically. This stems from the results we have established in Besbes et al. (2019), where we prove that static pricing guarantees at least 78.9% of the optimal dynamic pricing policy, and 95.5% in the special case where there are at most two units (over 30% of the parts we consider have at most two units). Empirically, we are not able to find an instance where static pricing loses more than 2.5% of the optimal dynamic pricing policy.

The objective of this project is to redesign and improve the pricing process while utilizing all available data, and ultimately provide a systematic and scalable approach to pricing in this context. In collaboration with experts from the aircraft OEM, we developed a novel and systematic approach that leverages existing data and captures the special features of rotable spare parts to derive the best pricing strategy to maximize profit. Prior to our work, the existing approach to pricing rotable spare parts was driven only by the repair cost, and did not factor in repair time, competition, and the special inventory dynamics. More broadly, the project is also one that aims at changing the conversation about pricing within the organization, uncovering potential for systematic approaches to pricing of other offerings by the firm.

Unique Challenges

Our engineering approach for the price optimization tool we developed dealt with the following challenges unique to rotable spare parts:

1. *Modeling the special inventory dynamics and market competition.*

The rotable selling process has its own unique inventory dynamics, where the inventory constraints, stochasticity in repair time, and market competition all play critical roles. Capturing this information and synthesizing how it should impact price decisions is key for a successful model.

2. *No knowledge of price-demand relationship.*

The OEM, in accordance with common practice in the aircraft industry, does not often change the prices of rotable spare parts. This, in turn, provides very little information on the price sensitivity of customers and more broadly the relationship between expected demand and price.

3. *Limited data for each rotable spare part.*

Although at the aggregated level, the amount of available data is reasonably large, the data associated with each rotable spare part are very limited due to the slow-moving nature of the system. Most rotable spare parts are expensive and have lifetimes of many years. A typical rotable spare part may be purchased only two or three times per year, and thus parameter estimates are inherently noisy.

Timeline and General Approach

This project started in the summer of 2016 and culminated with a large-scale controlled experiment that ended in spring 2019. Figure 1 depicts a timeline of the project, highlighting some key steps.

It is important to note that an initial pilot program was conducted to obtain buy-in within the organization on the potential for adjusting prices, and specifically to ensure that the market is indeed price sensitive. With this successful pilot and early buy-in on the potential, significant effort was put in developing a tailored mathematical model that addressed all the business needs. In turn, the challenge resided in estimating input for the model when possible or proxies when it was not directly possible. This involved a major effort in aggregating various disparate data sources residing on different systems, cleaning the data, and quantifying the residual uncertainty on the inputs needed. Finally, we developed a pricing analytics tool that allows users to obtain optimized price suggestions

Figure 1 The Timeline of the Project Ranged from June 2016 to March 2019

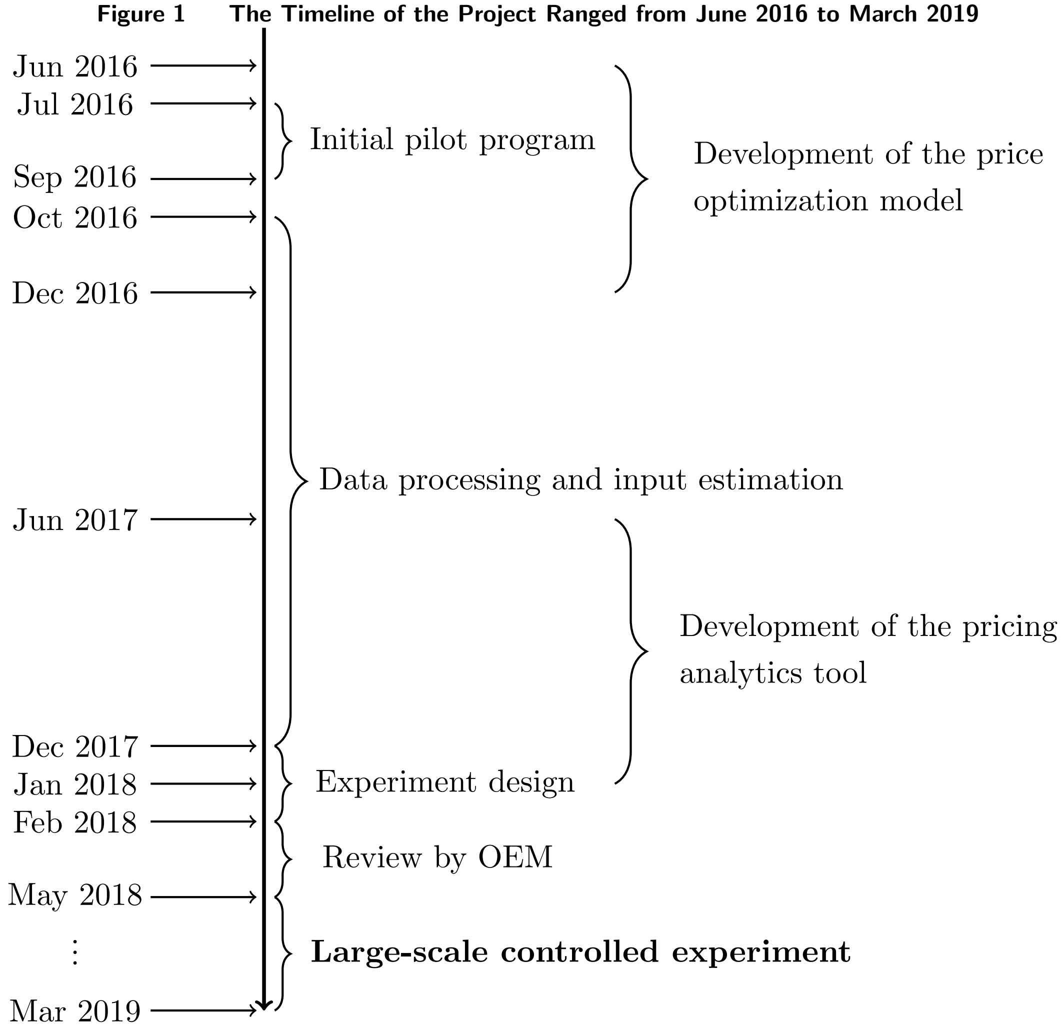

based on the available data and perform robustness checks. In Figure 2, we present an overview of the data-driven systematic approach we developed in the project.

We developed a model for rotable spare parts that can be used to understand the profit rate as a function of the total number of units, demand rate, repair time, and prices. To find the optimal prices, we estimate the inputs of the model using various data sets from the OEM and factor in other business constraints such as the maximum price change. Once the optimal prices are found on the estimated model, the prices are then robustified in the sense that we potentially adjust the price so that we ensure that it is robust to the uncertainty in all the parameter estimates. We highlight the importance of the robust approach to ensure that prices are not overly sensitive to our assumptions and the high level of uncertainty in

Figure 2 Our Data-driven Pricing Approach Includes Modeling, Optimization, and Implementation

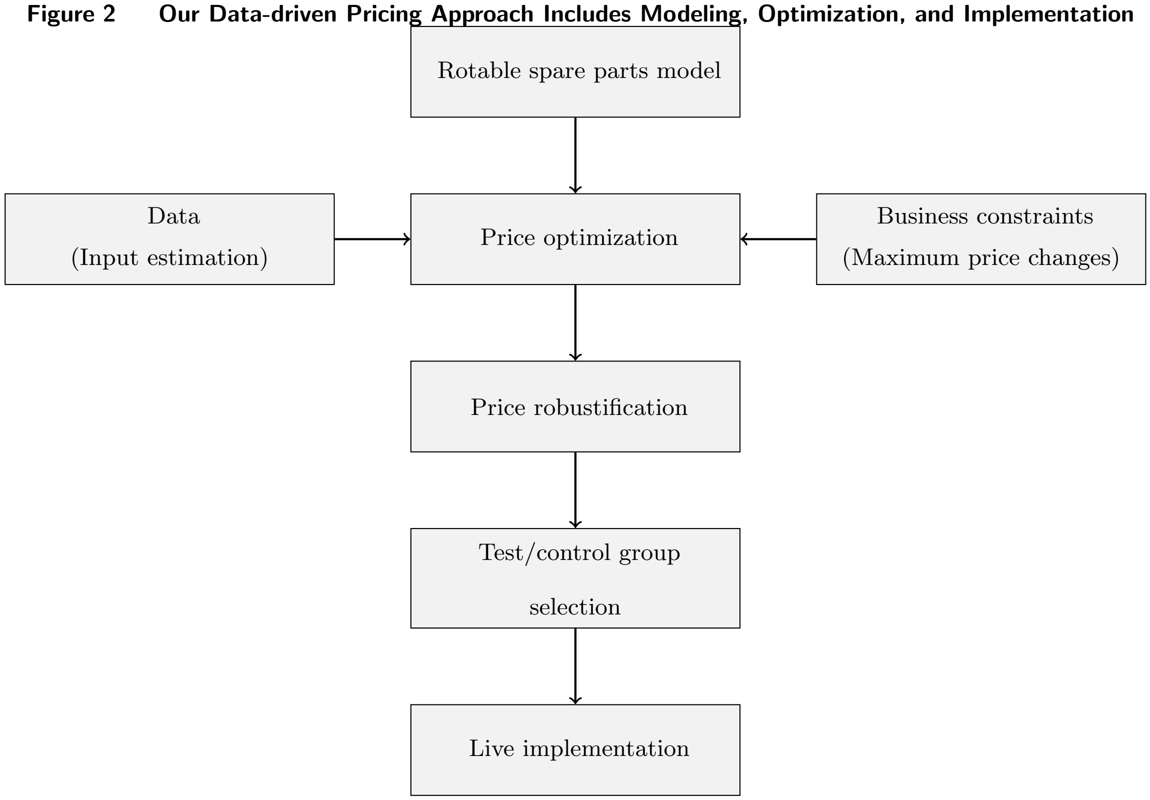

some inputs. A key feature of the decision support tool we developed is that the data and the uncertainty associated with some inputs (and their implications) can be overridden by the users, allowing users to challenge their own assumptions or potentially refine the inputs with knowledge not codified in the existing databases.

After review and testing of the tool, in conjunction with the OEM, we decided to launch a large-scale controlled experiment to test the suggested price changes. We use the robustified price to guide how we split the rotable spare parts into the control and test groups for the implementation. Based on 10 months of sales data, we found that the new system led to an increase of 3.9% in profits under a difference-in-difference (DiD) analysis. Equally importantly, this project has led to multiple other initiatives related to pricing analytics within the organization.

Related Literature

Research on inventory management of rotable spare parts has a long history dating back to Allen and D'Esopo (1968). Since then, many studies have been conducted on inventory management of repairable parts. Graves (1985) and Cohen et al. (1989) studied the problem of determining optimal inventory levels. Guide and Srivastava (1997) provided a review of models and applications of repairable inventory. More recently, several studies focused on aircraft spare parts. Simao and Powell (2009) used approximate dynamic programming to determine the inventory level of aircraft spare parts at each warehouse while Aisyati et al. (2013) studied the inventory policy using a continuous review model. Muckstadt (2004) contains a comprehensive overview of modeling approaches and solution methodologies for addressing service parts inventory problems.

At a high level, the system dynamics of rotable spare parts could be considered as a closed-loop supply chain. Many works, such as Fleischmann et al. (2003), Savaskan et al. (2004), Guide and Van Wassenhove (2009) and Calmon and Graves (2017), have been written in this broad domain. Another stream of works focuses on allocation and overhaul planning of rotable spare parts. These works include Tedone (1989), Arts and Flapper (2015) and Erkoc and Ertogral (2016).

Relatively fewer works focus on finding the best pricing strategies in selling rotable spare parts, typically in the context of reusable resources. Gans and Savin (2007) study dynamic pricing to maximize the expected profit for rentals. Their model uses a discounted reward structure with a discrete price ladder, although with multiple customer types. They show the near-optimality of static pricing in highly utilized rental systems where both the offered load and system capacity are large. In a similar model, we show in Besbes et al. (2019) that a static pricing policy is provably near-optimal in all parameter regimes. Our results hold even when the number of units is small, which is the situation for the rotable spare parts that we examine (see Figure 9). Lei and Jasin (2018) studied a related pricing problem of reusable resources where the service times are deterministic, which is not our situation (see Figure 7).

We remark that there has been a limited but steady stream of literature on successful pricing model implementations; these include Smith and Achabal (1998), Natter et al. (2007), Caro and Gallien (2012), Ferreira et al. (2015), Simchi-Levi and Wu (2018), and Xu et al. (2019). These implementations are typical in fast-moving industries such as fashion or online retail where there is a wealth of data. However, to the best of our knowledge, this is the first paper on the implementation of a pricing model in a slow-moving environment, and the first such paper in the context of reusable resources.

Rotable Spare Parts Pricing

Model and Assumptions

We begin by describing a model that captures the expected profit rate of a rotable spare part as a function of its price and inventory dynamics. We note that this model shall be applied to each (of the thousands) rotable spare part separately. For each rotable spare part, the total number of units is fixed. We let C denote the total number of units, which is also referred to as the pool size. Customer requests are assumed to arrive at the OEM according to a Poisson process with rate $\Lambda > 0$. Given the current price, p, of the rotable spare part, customers decide to purchase the rotable spare part if their willingness-to-pay exceeds p. We denote by $\lambda(p)$ the effective arrival rate at price p. When a customer decides to purchase a unit of the rotable spare part, that customer also simultaneously gives their broken unit to the OEM in the so-called exchange sale. The OEM will then send the broken unit for repair, and incur an associated expected repair cost c. The repair time (including travel time) of the rotable spare part is assumed to be exponentially distributed with mean

$1/\mu$ periods. We note that both interarrival times and repair times are each generated from independent and identically distributed processes.

Since both the interarrival and repair times are exponentially distributed, the rotable selling process can be modeled as a Markov decision process (MDP) with a one-dimensional state. Specifically, the state is one of $\{0, 1, \dots, C\}$, which represents the number of units that the OEM currently has on hand (available for sale). The transition rate of having i units on hand to $i+1$ units on hand thus is $(C-i)\mu$ because there are $(C-i)$ units in repair, each with an independent and identically distributed (i.i.d.) exponential repair time with mean $1/\mu$. In the other direction, the transition rate of having i units on hand to $i-1$ units on hand is simply $\lambda(p)$. Figure 3 illustrates the Markov chain embedding of our model. The numbers in the circles represent the number of units currently on hand in the system.

Figure 3 The Dynamics of Our Model is Described by a Markov Chain

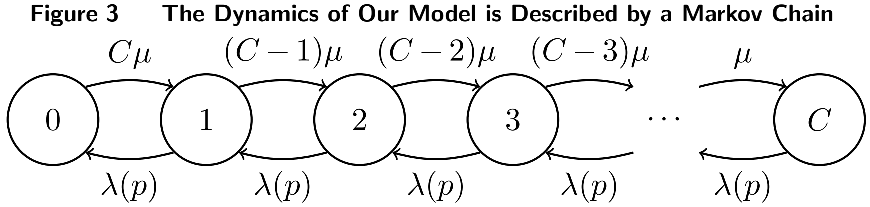

The goal is to find the optimal price of the rotable spare part to maximize the expected profit rate, that is, the product of the arrival rate $\lambda(p)$, profit per unit sold $p-c$, and availability $\mathbb{P}_0(p)$ (steady state probability of having at least one unit to sell). Since the demand rate $\lambda(p)$ is assumed to be decreasing in p, the profit of selling one unit, $p-c$, is increasing in p, and the stockout probability $\mathbb{P}_0(p)$ is decreasing in p, there is a nontrivial trade-off between making more profit per sale, the rate of selling, and the availability. See Appendix A for the mathematical formulation of the price optimization model.

In this model, we capture the competition in the market through the term $\lambda(\cdot)$, while the inventory dynamics, including the total number of units and repair time of a rotable spare part, are captured in the stockout probability $\mathbb{P}_0(\cdot)$. These key inputs were not captured by the 'cost+margin' method employed in the past.

Although in principle a dynamic pricing policy is optimal for this MDP, and was the original intent of the project, we instead relied on a static pricing policy for two reasons. First, from a performance perspective, numerical tests showed that the best static price loses at most 2.5% compared to dynamic pricing. This motivated our theoretical study in Besbes et al. (2019), which proved the near-optimality of a static price in such systems, in particular when C is small (which is often the case as we see in Figure 9). From the data, 2.5% is generally smaller than the average error of our estimated input paramaters. Second, from a practical perspective, a static pricing policy allows the OEM to keep its current practice of publishing a catalog of prices for rotable spare parts at the beginning of each year and maintaining the same price throughout the year. The OEM will not need to develop a new system for deploying the new pricing algorithm.

Justification of Assumptions

In deriving the price optimization model, we made two key assumptions that we seek to justify: (1) customers arrive according to a Poisson process, and (2) repair times are random, and specifically exponentially distributed. Throughout the rest of the paper, we will use two generic examples of rotable spare parts, a sensor and a jack, to illustrate our ideas.

Since selling rotable spare parts is a very slow-moving process, many rotable spare parts have limited sales during the year. Figure 4 depicts the average sales per year of each rotable spare part from 2010 to 2017. One can observe significant differences across spare

Figure 4 This Histogram of Sales Volume Reveals Slow-Moving Sales

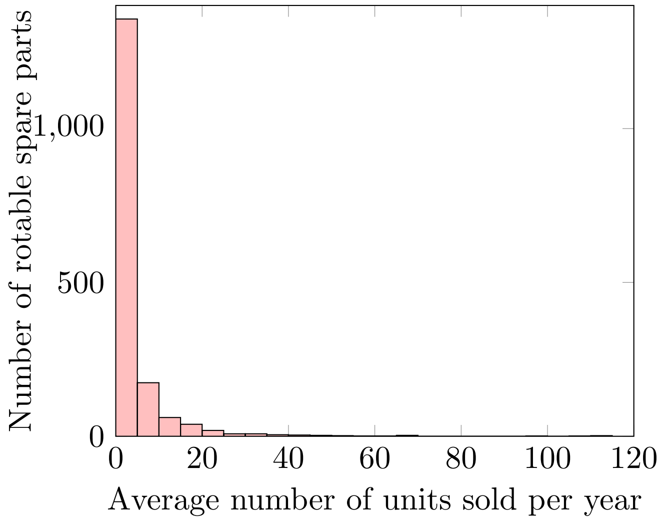

parts in terms of volume, but also that most of the parts are of a very slow-moving nature. The majority of rotable spare parts have less than three units sold per year, in which case estimating the repair and interarrival times is naturally very noisy.

To understand the distribution of interarrival and repair times, we focus on rotable spare parts with higher volumes. For example, Figures 5a and 5b shows the empirical interarrival time distributions for the sensor and jack. Similarly, we look at the empirical repair time distributions for the sensor and jack in Figures 6a and 6b. Although the OEM may change the price of the parts at the beginning of the year by at most 3% due to inflation, we ignore such changes when calculating the interarrival time.

In Figures 5 and 6, we also depict the best-fit exponential curves (dashed line) to the data. The data suggests that both interarrival and repair times are inherently quite random, and exponential fits appear to be reasonable. Admittedly, for the repair time a better fit might be a deterministic time plus an exponential time, although this type of random variable would impose significant technical challenges. The exponential fit, while not perfect, captures most of the shape of the empirical distribution.

Moreover, the behavior above is fairly typical across other rotable spare parts with volumes of at least three units per year. To demonstrate this, we compute the coefficient of

Figure 5 The Histograms of Interarrival Times for Two Parts Approximates an Exponential Distribution

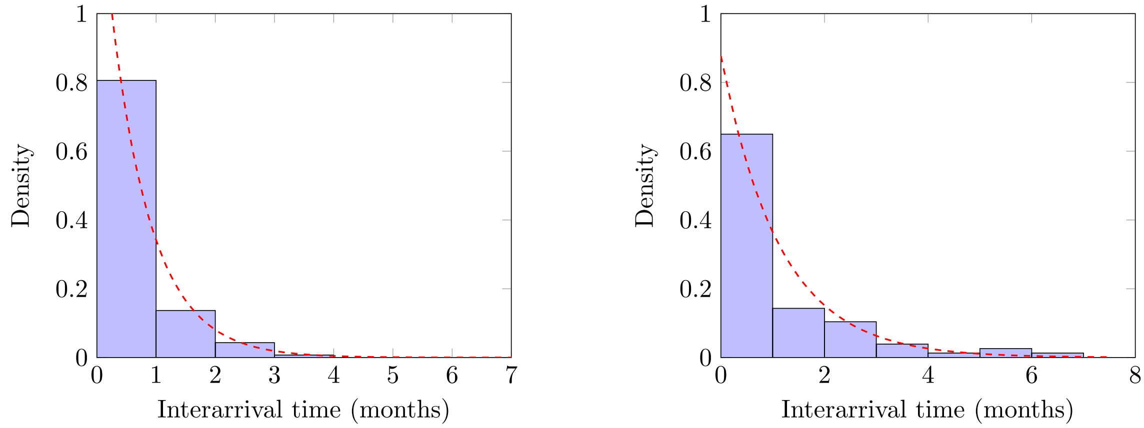

(a) Sensor (b) Jack

Figure 6 The Histograms of Repair Times for Two Parts Approximates an Exponential Distribution

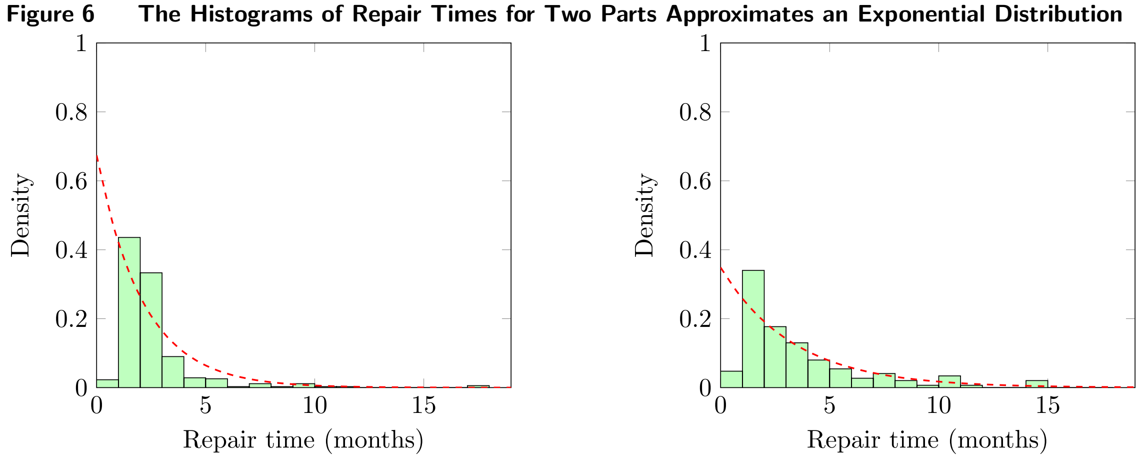

(a) Sensor (b) Jack

variation (CV) of the interarrival and repair times for each rotable spare part. In Figure 7, we report the empirical CVs of interarrival and repair times for all spare parts with more than three units sold per year.

A large proportion of the parts have coefficients of variation around one for both interarrival and repair times, in line with the assumed variation based on the exponential assumption. Moreover, the minimum CV we find is significantly far from zero, motivating the use of random times in our model. We shall later account for imperfections in the model

Figure 7 **The Coefficient of Variation of the Interarrival and Repair Times Reveal a Stochastic System**

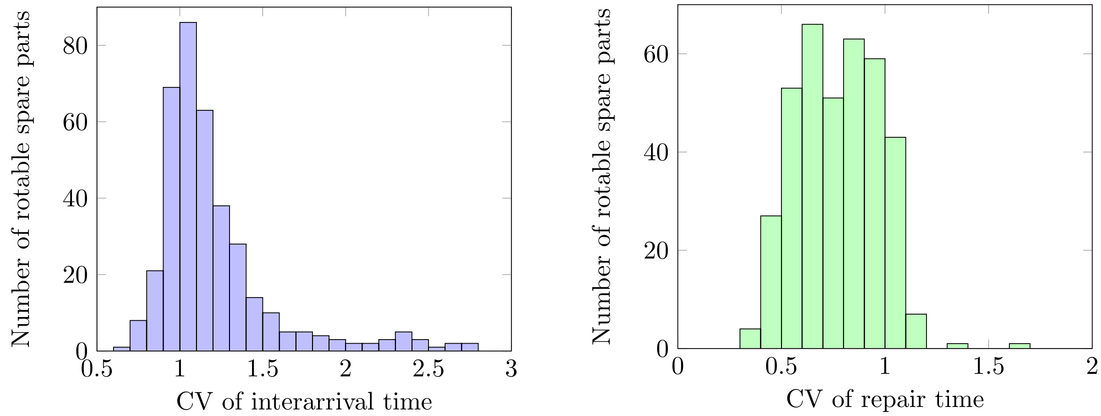

by adding a robust component to our pricing algorithm. While the data is not perfectly in line with the exponential assumption (which would lead to an exact CV value of one), the exponential assumption captures the randomness associated with repair times while allowing us to have a tractable function to optimize.

Price Sensitivity of Customers

Questions initially raised by our industry partner revolved around how price-sensitive customers were, and whether there was potential for significant price optimization. To test this, and before engaging fully in a revamp of the pricing process, we conducted a simple pilot program on a limited group of rotable spare parts in the summer of 2016. The approach in the pilot program was to discount parts with low sales and low risk of backorders so as to not disturb the existing supply chain. To do so, we selected the parts and discount levels according to a decision tree (see Figure 8).

We changed prices of rotable spare parts in three leaves of the decision tree. We started the pilot program on July 12, 2016 by reducing the price on 182 parts. An additional 95 parts belonging to the same leaves of the decision tree were used for the control group. We considered 77 days before and after the start date of the pilot program and performed a DiD analysis. Our analysis suggested an overall increase in profit of 17% and an estimated increase in sales volume of over 44%.

Figure 8 A Decision Tree Was Used for Selecting Discounts for Rotable Spare Parts in the Pilot

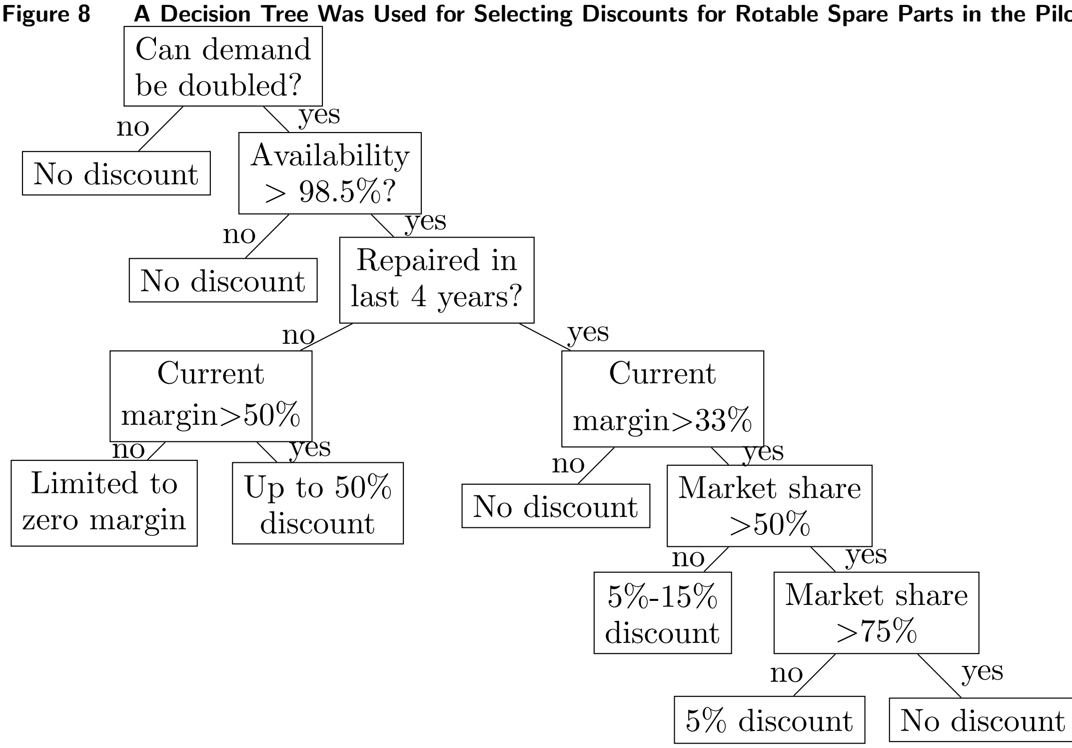

Although the pilot was relatively short and the parts included in the pilot program represent only a small portion of rotable spare parts, mainly the parts with very limited sales, the results of the pilot program indicated that there is significant potential for optimizing prices. This provided some evidence that customers are indeed conscious of price changes in the rotable selling market and that this is an environment where data-driven pricing can be leveraged to optimize prices. Furthermore, the pilot program convinced our industry partner of the need to systematize their pricing approach, and they decided to launch our pricing analytics approach across the entire rotable spare parts supply chain. Next, we detail how we estimated key inputs for our model.

Input Estimation

The estimation of inputs required extracting, aggregating, unifying, and cleaning data across multiple databases at the OEM. Moreover, estimating the demand function $\lambda(p)$ was a particularly challenging task because the OEM rarely changed the prices of parts beyond inflation adjustments.

Pool Size

The OEM identifies rotable parts at two levels, *individual* and *family*. An individual part has its unique part number while a family may contain several individual part numbers if they have different generations or are symmetric parts (one is applied to the left of the airplane and the other to the right). Given various business constraints at the family level, and the need to aggregate data for very slow-moving parts, we performed our analysis at the family level. Aggregating all data sources, we derived the mapping of individual part numbers to their family part number, applicability (models of airplanes for which the family of parts can be used for repair), and the description of each family. Aggregating data from on hand and in-repair inventory for each individual part, and leveraging the family-individual mapping above, we obtain the pool size (that is, the total number of units of inventory across the entire family). The pool size corresponds to C in our price optimization model. Figure 9 depicts the histogram of pool size of rotable spare parts at the OEM. For our two running examples, we have 10 units of the sensor and 11 units of the jack.

Figure 9 **The Pool Size for Rotable Spare Parts Is Primarily Under 10 Units**

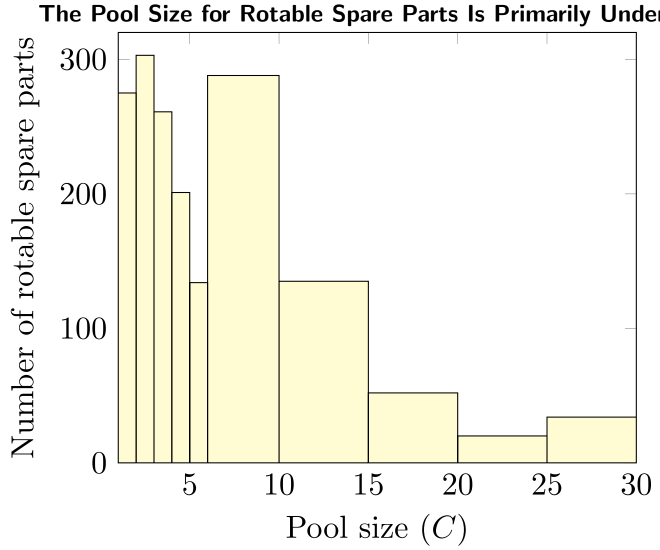

Repair Time and Cost

We collected the records of repair orders for a 10-year period from 2008 to 2018. Using the mapping of part number to family number, we calculated the average repair time (also called Total Turnaround Time (TTAT)) and repair cost for each rotable spare part family. These estimates correspond to the parameters $1/\mu$ and c in our price optimization model. In addition, we also calculated the standard deviation of these two quantities which we used later when we robustified our optimized prices. The average repair time of the sensor is 2.88 months with a standard deviation of 2.92 months, while the average repair time of the jack is 3.79 months with a standard deviation of about 3.34 months. The costs for these examples are not reported for confidentiality reasons.

Price-Demand Relationship

The estimation of inputs above is fairly straightforward and the main difficulties are dealing with missing data, outliers, and accounting for the remaining uncertainty (which we deal with by robustifying our prices later on). However, a key input for which no estimate is available in the data is the relationship between price and demand, that is, the function $\lambda(p)$. Estimating a demand curve is challenging in general, but even more so in a setting with almost no price experimentation. (A price experiment or dynamic learning approach would take years because of the slow-moving nature of rotable sales).

The given data provide only one point on the demand curve, which is the current price and the corresponding current demand rate. For example, for the sensor, the current demand rate is about 1.5 sales per month at its current price, and for the jack, the current price has a demand rate of about 0.8 units per month. Quite notably, the prices never change except for annual price inflation adjustments and thus a demand curve cannot be constructed explicitly from data. This presents a unique challenge. *Can one still approach* *pricing in a systematic fashion without a demand curve?* Is there a proxy for such a demand curve despite the structure above? Below, we discuss a systematic approach we developed for this environment. This should be seen as providing a starting point in estimating the demand curve and we will discuss how we deal with the remaining uncertainty later when we "robustify" the suggested prices.

The approach we take is one of assuming a linear structure of demand and attempting to obtain a proxy for demand at an alternate point, a hypothetical price where the OEM would be able to capture the full market share Λ (or close to it). Using the current price point and this alternate point, we simply fit a line between these two points to generate $\lambda(p)$. Figure 10 illustrates the proposed demand model. With such an approach, the question becomes one of (1) estimating a price at which one would capture full market share, and (2) estimating the total demand, Λ, in the market for any part. After discussions across the firm with experts, we made the assumption that the firm can get full market share if the price is set to the average repair cost or half of the current price, whichever is higher. Note that in the setting with no inventory constraints, Cohen et al. (2015) show that using the price from a linear demand model with the correct price-intercept can result in large profits for many nonlinear demand models. In our setting, we account for the inventory dynamics in the profit objective and estimate the intercept on the demand axis to generate a linear model.

Since we have already estimated the current demand rate of each rotable part, we next need to estimate the current market share so that we can calculate the full demand rate. Note that if $\tilde{p}$ is the current price and MS is the market share, then the estimate of the full demand rate Λ is $\lambda(\tilde{p})/MS$. Although the OEM may have knowledge about the market share of its rotable spares business at an aggregate level, the market share at each rotable part level is unknown.

Figure 10 Our Demand Model Is Assumed to be Linear by Connecting 2 Points

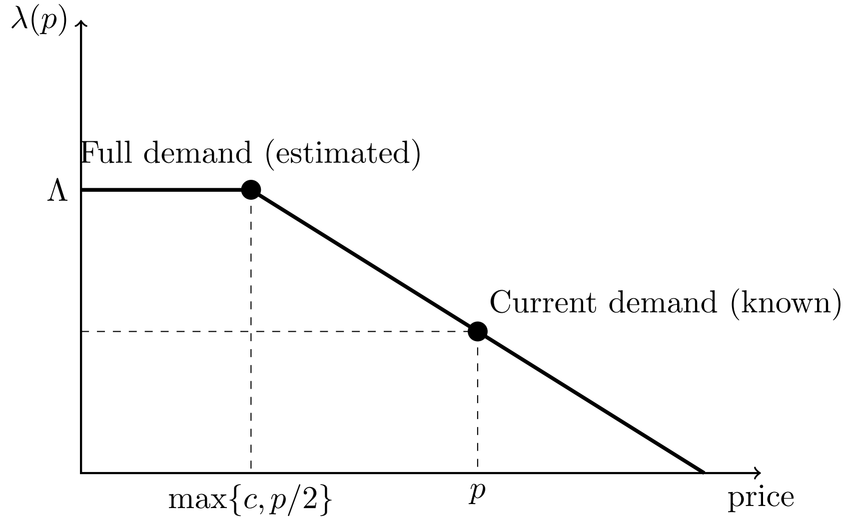

It is important to note that in theory, there are many ways to obtain the market share at the part level and we investigated many of those approaches. Repair logs of airplanes are one source; however, only around 70% of customers use a common software package and, in many cases, the removal of a rotable spare part is only for testing or is a precedent step in fixing another part. This data source, while effective in theory, proved to be highly incomplete in practice. We also investigated the data in user manuals to try to infer frequencies of replacements needed, but these led to inaccurate estimates of demand because the service manuals are more geared toward inspection than replacement. We also analyzed the rotable purchase history from warranty customers and the OEM-owned aircraft, of which all the purchases of rotable spare parts can be safely assumed to be from the OEM. However, the purchase frequency of warranty customers and the OEM-owned aircraft are not representative of the entire fleet because the models of the OEM's aircraft are limited and the aircraft are relatively new.

Given the limitations above with all available data sources, we decided to adopt a simple but robust way to estimate the current market share. Rather than focusing on parts, we focus on customers and airplanes. We assume, according to common practice in the aircraft industry, that each rotable spare part needs to be replaced at least once every 10 years. For any rotable spare part i, we can calculate the total number of customers who should have bought the part over 10 years, T_i, by simply counting the number of airplanes to which the rotable spare part i can be applied. On the other hand, from the sales data, we can extract the number of distinct customers who bought the rotable spare part i, denoted by B_i. In turn, the market share of rotable spare part i is estimated by $MS_i = \frac{B_i}{T_i}$. Notice that in this approach, we underweighted the customers who regularly shop from the OEM and overweighted the customers who seldom purchase from the OEM.

As an example, we estimated the market share for the sensor to be 30% and the market share of the jack to be 22%. We noticed that most of the estimates may have some error associated with them due to the limited amount of recorded data for each rotable spare part. We directly addressed this in the development of the pricing analytics tool by providing estimated bounds on the market share and generating prices that are robust to changes in the inputs.

Price Optimization and Robustification

A natural approach would be to simply compute the optimal price for each spare part under the estimated inputs by solving Equation (1) in Appendix A. However, as stated earlier, there are multiple sources of potential errors in the estimation of the inputs given the unique environment in which we operate. We develop a robust pricing approach to account for such potential errors in calculating the suggested price for each rotable spare part. The main idea in calculating the suggested price is to treat the estimated inputs as ranges of possible values rather than just a fixed value because the initial estimates may not be accurate due to lack of data. We select the price that works best on a variety of scenarios, which would make the suggested price be robust to estimation error. In Appendix B, we provide the algorithm to calculate the suggested price for a given rotable spare part.

In our running examples of the sensor and the jack, the optimal price for the sensor is to reduce the price by 11%, while price reductions associated to 95% optimality are 18% and 4%. After evaluating the three candidates, the final suggested price is chosen to be the optimal price from the model, which is about an 11% decrease of the original price. In the jack example, the price candidates are 24%, 16%, and 10%. After evaluating the expected profit on randomly generated inputs, the highest-price candidate is the selected suggested price and this corresponds to a 10% price decrease. The main reason the maximum price candidate is chosen as the final suggested price for the jack is due to the high variance in the repair cost, in which case the algorithm tends to be more conservative.

Implementation

In this section, we discuss the visualization and decision support tool we developed for helping the OEM in the implementation and in the controlled experiment conducted for 1,702 rotable spare parts.

Visualization and Decision Support Tool

In Figure 11, we show a screenshot of the visualization and decision support tool. We implemented this tool in Python using the Tkinter package for the GUI. The top left of the screenshot shows a panel that allows the users to search and select the rotable spare part (by family number) they want to analyze. After the user selects a specific rotable part, its basic information, such as part number, usage, applicability (i.e., models on which it can be applied), and past sales is displayed in the lower left corner. Next, the estimated inputs of the price optimization model are displayed in the top middle region of the tool. Instead of displaying the estimated market share, we choose to display the range of the market share. The minimum and maximum market share levels are calculated using the formula described in the algorithm in Appendix B.

Figure 11 **We Provided a Visualization and Decision Support Tool for the Implementation**

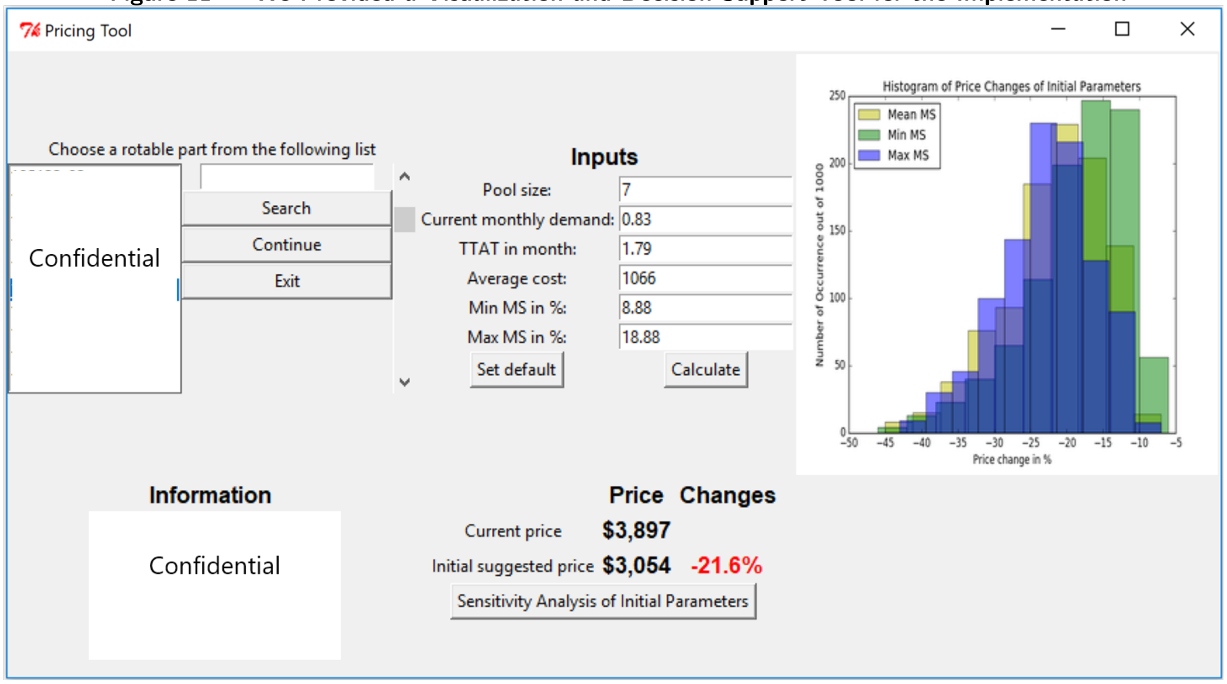

We designed the entries of inputs to be editable so that if the users do not agree with our estimation, they can manually overwrite the value. Note that if the minimum and maximum market shares are changed, then the estimated market share is updated to be the average of the minimum and maximum. Another advantage of making the entries editable is that it allows the users to do what-if analyses, which may help them in other areas of their operations. For example, they may want to understand the impact of an increase or decrease of the pool size, and how to negotiate with repair agents on repair time and cost.

When the users agree on the input values and click calculate, the suggested price as well as the percentage change will be displayed at the lower middle part of the tool. This price is the output returned by the algorithm described in Appendix B. In addition, since the estimation of market share contains the most uncertainty, we built a sensitivity analysis function in the tool to show how different market share levels will impact the suggestion. This corresponds to the histogram on the right side of Figure 11. To create the histogram, we generate $(\frac{1}{\mu'}, c')$ 1,000 times using the same procedure described in Step 2 of the algorithm in Appendix B. Then, we calculate the optimal prices under the different sets of inputs focused on the various market share estimates (MS, MS_{min}, MS_{max}). Finally, the histograms of percentage of changes associated with the optimal prices are plotted on the visualization and decision support tool. The sensitivity analysis aims to provide guidance to the users on the direction of the price change (whether to increase or decrease the price). As long as the users believe the range of the market share is correct, the sensitivity analysis will provide the frequency of optimal price changes. Even if the users choose not to follow the suggested price, the sensitivity analysis may inform them which direction the price should go to maximize the profit rate. The users can leverage this information as well as other business constraints to make final decisions.

Overall, the tool we developed serves as a decision support tool, which gives a suggested price that is robust to input estimation errors for a given rotable spare part. This tool helps the users in making pricing decisions and understanding the implications of the assumptions they have about market share, pool size, repair time, and repair cost.

Experiment Design

We now describe the controlled experiment we conducted to test the effectiveness of our pricing analytics tool. To ensure the control and test groups are comparable, for each rotable part selected in the test group, one needs a 'similar' rotable part in the control group. One natural way is to look at the estimated inputs of the price optimization model of each rotable spare part. If every estimated input is similar between two rotable spare parts, then one could claim these two parts are similar, and more importantly, the suggested changes of these two parts would be similar as well. However, this approach did not work because there are few pairs of rotable spare parts where all estimated inputs are similar. To overcome this difficulty, we propose a procedure in Appendix C for selecting parts into the control and test groups.

The key idea behind this selection procedure is that the usage of each rotable spare part is a natural classifier of different rotable spare parts. Instead of focusing on the similarity of all estimated inputs between two different rotable spare parts, which is a high-dimensional clustering problem, we just focus on the outcome–the suggested price changes–which is the result of running the price optimization algorithm described above. This leads to a one-dimensional pairing problem, resulting in a selection method that is easy to explain internally in the organization.

Using the above procedure, 852 out of 1,702 rotable spare parts are selected for the test group and the remaining 850 rotable spare parts are in the control group. Within the test group, 744 rotable spare parts were selected to receive price decreases while the prices of 108 rotable spare parts were increased. We handed the visualization and decision support tool and the lists of the control and test groups to the team at the OEM in the spring of 2018. The OEM used the tool to refine price changes based on scenario analyses and business constraints, and decided if the price should be adjusted from the default recommendation, and if so, to what level.

Implementation Results

The OEM changed the prices of the rotable spare parts in the test group on May 4, 2018 and sent a notification to its customers. The notification sent by the OEM did not include the part numbers of the rotable spare parts for which it changed prices so that we could isolate the effect of the price change from the marketing effort.

We collected the sales data of the 1,702 rotable spare parts from July 7, 2017 to March 5, 2019. These data represent the sales of rotable spare parts for 208 working days before and after the implementation of price changes. We perform a DiD analysis in the aggregate level to measure the effect of our price optimization model as well as the decision support tool. Since the objective of the price optimization model is to maximize the expected profit rate, we focus on the profit changes in the two groups before and after the implementation. From the DiD analysis, we see an estimated impact of 3.9% *increase* in profit from our pricing analytics tool. We omit the details of the analysis due to confidentiality reasons.

To ensure the increase in profit is not impacted because some days have very good (or poor) sales records, we provide a confidence interval of the DiD estimate of the profit increase. The interval is generated by generating 10,000 randomly generated bootstrapped data sets consisting of 208 working days with replacement in both the before and after periods. For each data set, we compute the DiD estimate and find the confidence interval of the estimate to be [3.61%,3.95%]. Conservatively, the new pricing algorithm adds millions of dollars of profit per year in the rotable spare parts business for the OEM.

Conclusion

In this collaborative project with a major aircraft OEM, we investigated the problem of setting appropriate prices for rotable spare parts. We adopted a data-driven price optimization approach to maximize the expected profit rate, which can also be used for a broader class of problems concerning reusable resources. This approach captures special system dynamics such as fixed pool size, random repair times, exchange sales, and market competition, most of which are not taken into account in the legacy approach. In addition, the algorithm proposed and the tools developed allow users to understand the implications of input errors and to be robust to such errors.

We conducted a large-scale controlled experiment and received encouraging implementation results, namely a 3.9% increase in profit according to a DiD analysis. The successful implementation demonstrates the power of a structured and systematic pricing analytics approach, even in a slow-moving environment. As of today, the OEM is engaged in expanding and building an internal version of the visualization and decision support tool that can seamlessly integrate with internal data flows. This methodology and its refinements will be a core part of pricing rotable spare parts. In addition, this project opened up the discussion of a systematic review of pricing processes across the organization, which has led to multiple new projects including rental tool pricing and service center quote estimation.

Appendix A: Mathematical Formulation of Price Optimization Model

Letting $\mathbb{P}_0(p)$ denote the steady state probability of having zero units available (i.e., stock-out probability), then our objective to maximize the expected profit rate can be written as

$$\max_p \; \lambda(p)(p-c)(1-\mathbb{P}_0(p)), \tag{1}$$

where the stockout probability $\mathbb{P}_0(p)$ can be expressed as

$$\mathbb{P}_0(p) = \frac{(\frac{\lambda(p)}{\mu})^C}{\sum_{i=0}^{C} \frac{C!}{(C-i)!}(\frac{\lambda(p)}{\mu})^{C-i}}.$$

Note that $\mathbb{P}_0(p)$ can be derived by solving the balance equations corresponding to the Markov chain described in Figure 3. The balance equations signify that the total flow into a state is equal to the total flow out of every state, and the total probability of being in any state is one. Letting $\mathbb{P}_i(p)$ denote the steady state probability of having i units in stock when the price is p, the balance equations can be written as

$$\mathbb{P}_1(p)\lambda(p) = \mathbb{P}_0(p)C\mu$$

$$\mathbb{P}_{i+1}(p)\lambda(p) + \mathbb{P}_{i-1}(p)(C-i-1)\mu = \mathbb{P}_i(p)(\lambda(p)+(C-i)\mu) \text{ for } i = 1, \ldots, C-1$$

$$\mathbb{P}_{C-1}(p)\mu = \mathbb{P}_C(p)\lambda(p)$$

$$\sum_{i=0}^{C} \mathbb{P}_i(p) = 1.$$

Appendix B: Rotable Spare Part Pricing Algorithm

Step 1. Find the candidate prices

—Estimate inputs (C, $\frac{1}{\mu}$, c, market share(MS), $\lambda(\cdot)$).

—Calculate the base optimal price, p_{opt} by solving Equation (1).

—Compute the minimum and maximum prices that can achieve at least 95% of optimality under estimated inputs, denoted by p_{min} and p_{max}.

Step 2. Select a robust price

—Generate 1,000 inputs $(\frac{1}{\mu'}, c', MS')$ according to

$$\frac{1}{\mu'} \sim \begin{cases} Normal(\frac{1}{\mu}, std_{\frac{1}{\mu}}), \text{if } std_{\frac{1}{\mu}} > 0 \text{ and \# records} \geq 5 \\ Uniform(0.8\frac{1}{\mu}, 1.2\frac{1}{\mu}), \text{o.w.} \end{cases}$$

$$c' \sim \begin{cases} Normal(c, std_c), \text{if } std_c > 0 \text{ and \# records} \geq 5 \\ Uniform(0.8c, 1.2c), \text{o.w.} \end{cases}$$

$$MS' \sim \begin{cases} Uniform(MS_{min}, MS), \text{w.p. } \frac{1}{2} \\ Uniform(MS, MS_{max}), \text{w.p. } \frac{1}{2} \end{cases}$$

where

$$MS_{min} = \max\{\min\{0.8MS, MS - 0.05\}, 0.01\}$$

$$MS_{max} = \min\{\max\{1.2MS, MS + 0.05\}, 0.99\}.$$

—Evaluate the average profit rate of p_{opt}, p_{min}, and p_{max} under each set of generated inputs $(C, \frac{1}{\mu'}, c', MS')$ on the objective function in Equation (1).

—Return the price with the highest average profit value across all scenarios.

Appendix C: Part Selection Procedure

Step 1. Group rotable spare parts by their usage category (e.g., valve, pump)

Step 2. In each usage group, sort rotable spare parts in ascending order based on their suggested price changes.

Step 3. In each usage group with size at least two, select the rotable spare parts according the following rule. For $i = 1, 3, 5, 7, \ldots$, randomly assign one of i and $i+1$ into the test group and the other into the control group. If the size of the usage group is odd, randomly assign the last part into the test or control group.

Step 4. Combine usage groups of size one, repeat Step 2 and Step 3 for the combined group.

Acknowledgement

We would like to thank our industrial partner for their valuable help throughout the project.

References

Aisyati, Azizah, Wakhid Ahmad Jauhari, Cucuk Nur Rosyidi. 2013. Determination inventory level for aircraft spare parts using continuous review model. *International Journal of Business Research and Management (IJBRM)* **4**(1) 1–12.

Allen, Stephen G, Donato A D'Esopo. 1968. An ordering policy for repairable stock items. *Operations Research* **16**(3) 669–674.

Arts, Joachim, Simme Douwe Flapper. 2015. Aggregate overhaul and supply chain planning for rotables. *Annals of Operations Research* **224**(1) 77–100.

Besbes, Omar, Adam N Elmachtoub, Yunjie Sun. 2019. Static pricing: Universal guarantees for reusable resources. *arXiv preprint arXiv:1905.00731* .

Calmon, Andre P, Stephen C Graves. 2017. Inventory management in a consumer electronics closed-loop supply chain. *Manufacturing & Service Operations Management* **19**(4) 568–585.

Caro, Felipe, Jérémie Gallien. 2012. Clearance pricing optimization for a fast-fashion retailer. *Operations Research* **60**(6) 1404–1422.

Cohen, Maxime C, Georgia Perakis, Robert S Pindyck. 2015. Pricing with limited knowledge of demand. Tech. rep., National Bureau of Economic Research.

Cohen, Morris A, Paul R Kleindorfer, Hau L Lee. 1989. Near-optimal service constrained stocking policies for spare parts. *Operations Research* **37**(1) 104–117.

Erkoc, Murat, Kadir Ertogral. 2016. Overhaul planning and exchange scheduling for maintenance services with rotable inventory and limited processing capacity. *Computers & Industrial Engineering* **98** 30–39.

Ferreira, Kris Johnson, Bin Hong Alex Lee, David Simchi-Levi. 2015. Analytics for an online retailer: Demand forecasting and price optimization. *Manufacturing & Service Operations Management* **18**(1) 69–88.

Fleischmann, Moritz, Jo AEE Van Nunen, Ben Gräve. 2003. Integrating closed-loop supply chains and spare-parts management at ibm. *Interfaces* **33**(6) 44–56.

Gans, Noah, Sergei Savin. 2007. Pricing and capacity rationing for rentals with uncertain durations. *Management Science* **53**(3) 390–407.

Graves, Stephen C. 1985. A multi-echelon inventory model for a repairable item with one-for-one replenishment. *Management science* **31**(10) 1247–1256.

Guide, V Daniel R, Rajesh Srivastava. 1997. Repairable inventory theory: Models and applications. *European Journal of Operational Research* **102**(1) 1–20.

Guide, V Daniel R, Luk N Van Wassenhove. 2009. Or forumthe evolution of closed-loop supply chain research. *Operations research* **57**(1) 10–18.

Lei, Yanzhe, Stefanus Jasin. 2018. Real-time dynamic pricing for revenue management with reusable resources, advance reservation, and deterministic service time requirements .

Muckstadt, John A. 2004. *Analysis and algorithms for service parts supply chains*. Springer Science & Business Media.

Natter, Martin, Thomas Reutterer, Andreas Mild, Alfred Taudes. 2007. Practice prize reportan assortmentwide decision-support system for dynamic pricing and promotion planning in diy retailing. *Marketing Science* **26**(4) 576–583.

Savaskan, R Canan, Shantanu Bhattacharya, Luk N Van Wassenhove. 2004. Closed-loop supply chain models with product remanufacturing. *Management science* **50**(2) 239–252.

Simao, Hugo, Warren Powell. 2009. Approximate dynamic programming for management of high-value spare parts. *Journal of Manufacturing Technology Management* **20**(2) 147–160.

Simchi-Levi, David, Michelle Xiao Wu. 2018. Powering retailers digitization through analytics and automation. *International Journal of Production Research* **56**(1-2) 809–816.

Smith, Stephen A, Dale D Achabal. 1998. Clearance pricing and inventory policies for retail chains. *Management Science* **44**(3) 285–300.

Tedone, Mark J. 1989. Repairable part management. *Interfaces* **19**(4) 61–68.

Xu, Joseph, Peter S Fader, Senthil Veeraraghavan. 2019. Designing and evaluating dynamic pricing policies for major league baseball tickets. *Manufacturing & Service Operations Management* **21**(1) 121–138.