Paper 1774-2018

Personal Lending: What is Customer Price/Credit Optimization? Is Experimental Design Inevitable to Optimize Price/Credit?

Yuri Medvedev, Ph.D., Chief Mathematician, Bank of Montreal, Toronto, Canada

ABSTRACT

What is customer price/credit optimization? The author has faced this challenging question numerous times with various analytical and business management professionals. On the surface, the question itself can appear straightforward to both analytics professionals and business management facing the complex decisions around customer value management. However, in reality, variations in understanding of customer price and credit optimization concept are quite vast and have led to a high degree of confusion around the optimization concept within the financial industry.

This paper will define the concept of customer pricing/credit optimization and the key role experimental design plays in the definition of the optimal solution. Eventually, any candidate for an optimal pricing/credit solution must be taken through a proper validation based on a proper experimental design. The paper will illustrate the concepts and use cases on real lending portfolios of leveraging proper experimental designs and examples of pricing and credit optimization solutions in detail.

It will be shown that by implementing a proper experimental design at the customer level, a pricing/credit optimization solution can be achieved and deliver 99% incremental profit over a non-optimized solution.

This is a follow-up of author’s previous presentation for SGF2016 (Medvedev, 2016) on customer credit and pricing optimization. The critical tool for the author’s optimization solutions is the effect/uplift modeling, developed using SAS/STAT® software.

INTRODUCTION

The author has been reviewing and developing customer level optimization solutions for lending portfolios since 2001, with the financial and analytical industries evolving significantly over this time. In the author’s view, there are significant opportunities to advance beyond conventional approaches to optimization in the financial industry. This is true of price/credit optimization, which is a critical component of the modern lending business. As financial institutions continually work towards optimizing their lending portfolios, increasing value to both the customer and the institution, such advances represent a significant opportunity for increasing profitability and driving financial outcomes.

Section I provides insight into common pitfalls in conventional approaches to pricing/credit optimization. The first is failing to properly leverage experimental design. The experimental design is a fundamental requirement to true customer level optimization, and that it is too often ignored in industry. Experimental design is the foundation for comparing two or more pricing/credit strategic solutions, and the only tool to validate an optimal pricing/credit solution. Without the proper use of the experimental design methodology, any resulting optimization remains only a hypothesis. The experimental design methodology is the industry standard for clinical trial studies in the health industry and must become a standard for the financial industry as well.

A second common pitfall impeding optimization lies in mistaking correlation for causality. The conventional approach to developing price/credit optimization is to leverage predictive modeling, which primarily captures correlations. However, if we want to model *causalities*, we have to approach solutions through an analysis of experiments where effect/uplift modeling plays the key role (see Medvedev, 2016 for an initial treatment of this topic).

The experimental design is the required component of any customer level pricing/credit or other optimization. Without an experimental design approach, no customer level pricing/credit, or other, optimization can be achieved. When reviewing a pricing/credit optimization solution product managers must ask developers if the solution was, at least, validated on a proper experimental design test. It is even better when both solution development and validation parts are supported by experimental design results.

To illustrate the experimental design based price optimization methodology, sections II-VI describe a specific instance of customer level price increase optimization for a portfolio of unsecured lines of credit. The price optimization solution is based on an experimental design test that the author launched as of April 2014. ***The effect/uplift modeling that ranks customers by their price sensitivity is the most critical component of author's optimization solutions*** (Medvedev, 2016).

Portfolio-wide price increases by 0.25% or 0.5% are common pricing adjustments for the lending business. In this paper, we illustrate how to optimize 0.5% price increase. The final step of our price optimization methodology is the Simulation Box development. ***The Simulation Box is the finest solution developed specifically for customer level optimization***. This solution allows business managers to run segments or customers of interest with given business constraints and objectives and effectively identify best customers to target.

In Section VI, various optimization scenarios are discussed, reviewing short-term vs long-term business objectives. Optimization solutions based on short-term objectives can be very sub-optimal and even hurt the business in the long-run. All our cases are supported with the actual experimental design data.

We demonstrate using a common portfolio price increase of 0.5% is actually very far from the optimal, as we show with the real data that the significant portion of the regular lending portfolio reacts very negatively to this common pricing change. ***Identifying these customers and removing them from pricing-up practically doubles the profitability effect.***

To optimize customer pricing, price elasticity concept traditionally plays the key role. However, the author finds that the traditional price elasticity approach is of a little help when we try to accomplish a customer level pricing optimization. Practicing with the real portfolios and data the author finds it's very difficult to develop realistic price elasticity and demand curves on the customer level. For the lending business, the demand curve that defines the price elasticity is actually the balance curve. Pricing portfolio managers want to know how a price change, for example of 0.5%, affects customer's behavior for all and every customer. It is shown that customer behavior varies across the portfolio, some being very price sensitive and others with no impact at all.

The author hopes that, over time, the financial industry will increasingly leverage experimental design methodologies when approaching pricing/credit and other optimization. Academia can assist significantly by developing and explaining these concepts.

The price optimization example and associated insights are obtained in the course of the on-going research project on pricing, credit and risk optimization conducted jointly by the Bank of Montreal and the Smith School of Business at Queen's University.

In the interest of protecting the bank's privacy, all data presented in this paper has been transformed. However, these transformations still preserve general shapes of curves, still allowing us to effectively demonstrate the overall ideas, methodologies, and concepts.

I. WHAT IS CUSTOMER PRICE/CREDIT OPTIMIZATION?

Our discussion is quite general and can be applied to any business and any customer optimization. In this particular paper, all examples are from lending portfolios. To make the discussion straightforward we explain all concepts and cases with pricing/credit optimization examples based on cards and unsecured lines of credit.

EXPERIMENTAL DESIGN TO COMPARE SOLUTIONS

Since we are going to talk about best or even the very best/optimal solutions we have to emphasize how we compare different solutions. This is how we come inevitably to the concept of experimental design which was introduced as long as in the mid of 18th century for clinical trial studies. Charles S. Peirce, in his publications as of 19th century, emphasized the importance of experimental design to make conclusions in statistics. Ronald Fisher, with his ground-breaking research with agricultural applications dated back to the first half of 20th century, is one of the founders of the modern theory on designing experiments.

To compare a strategy A to a strategy B on a portfolio or its segment, (Champion/Challenger strategy), we must design a proper experimental design test on the portfolio or the segment of interest.

This is the definition and the foundation of the comparison/assessment process. We can develop and justify a superior performance scenario A over B by different means. Eventually, we have to validate this hypothesis with the proper experimental design test. Otherwise, it is just an open hypothesis. Given a segment of interest and strategic scenarios A and B, we pick a randomized sub-segment to run under scenario A and the rest of the segment we run under scenario B. Only looking and comparing performance results for scenarios A vs B we can conclude and support statistically a superior performance of, say, scenario A over B. Therefore the critical role of the experimental design is proven to be required in order to develop the optimal solution.

OPTIMAL SOLUTION

Given

1. Target or objective function
2. Business dimensions
3. Business Constraints

We can search for the optimal, pricing or credit limit solution.

The ultimate target function for any business is certainly the profit which includes three main parts: i) revenue; ii) cost; iii) losses. Business dimensions or input variables define the space where we search/model for a better or the optimal solution. Business dimensions can include traditional customer lending metrics: balances at different time points, the history of credit limits, delinquencies, payment patterns, internal and external data, bank relationship data, demographic data and more. The author finds that derived variables can increase the list of business dimensions drastically and increase the power of the optimal solution. This optimal solution within the given space of input dimensions can be sub-optimal in the space extended with derived variables or other additional variables. So adding additional input/dimensions can turn the optimal solution into a sub-optimal. Business constraints can include descriptions of specific segments of interest, specific pricing or credit management constraints. For the pricing optimization, we can start with optimizing the price increase/decrease strategy only in the range from 0.1% - 2.5%. Metrics that define the business health can also be included in business constraints. For example, the traditional Loss/ Balance ratio or even better the Net-Revenue/Loss ratio can be parts of business constraints.

Given the target function, business dimensions and constraints, the optimal solution is the one that performs at least or even better than any other solution on every sub-segment defined by the given dimensions and constraints with respect to the target function.

In the event of conventional modeling, we have a target variable and aim to use the input variables to model the target. Building the perfect model is usually a great challenge. Similar, building the perfect optimal price/credit solution is usually hard to achieve. Therefore the role of a solution assessment becomes fundamental for the pricing/credit optimization process***, which can only be done using a proper experimental design test.***

The experimental design methodology is the industry standard for clinical trial studies, however, it is not common practice in the financial industry, particularly for pricing and credit limit optimizations. For more than fifteen years the author reviewed numerous solutions from internal and external analytical developers on pricing/credit optimizations. Proper assessments or validations of these optimization solutions using specific experimental designs were hardly ever completed by developers. In many cases, the author conducted his own assessments/validations based on properly designed experimental tests. Some of these assessment examples can be found in (Medvedev, 2016). Many of these optimization solutions were developed without the use of experimental design validations and therefore such solutions yielded neutral or even negative results. A business can implement a version of optimized pricing/credit solution, however, if this solution is not validated by an experimental design test then the actual performance of the solution can be very different from the predicted outcome and can even harm the business.

There are several reasons for this critical gap, and we will begin with a discussion of two major ones.

Reason #1: Lack of Proper Measurement Discipline for Actions and Treatments.

Through deep-dive analysis on a credit portfolio, recommended pricing changes can be identified with the target objective of improving profit. However, after launch, financial institutions routinely utilize pre-post or year/year analyses to assess the performance. Without implementing a proper randomized control group, the actual impact on the portfolio cannot be properly assessed and understood. The author had numerous discussions with business managers on pricing changes, some of these pricing changes were run with comprehensive experimental design based measurements and some were not. If an experimental design analysis was not a part of the solution development process then we observed often that actual pricing effect results were quite different from expected and far from optimal.

Reason #2: Confusion around Correlation vs Causality and Predictive vs Effect/Uplift Modeling.

We start with reminding the famous R. Fisher's discussion with his colleagues at the end of 1950th on smoking causing lung cancer. Fisher criticized his colleagues for making anti-smoking "propaganda". He certainly agreed that smoking and lung cancer are correlated. At the same time, he criticized that British medical institutions had committed "an error…of an old kind, in arguing from correlation to causation". Proving the correlation does not mean proving the causation. R. Fisher argued, a blind man is walking on a street and using the white stick. So these two events – the blind man and the white stick are highly correlated. Does the white stick cause the blindness? In the "smoking dispute" Fisher claimed, "to take the poor chap's cigarettes away from him would be rather like taking away the white stick from a blind man".

After almost 60 years from this discussion, we can still clearly see this confusion between correlation and causality in conventional approaches to optimization. Reviewing numerous optimization solutions on pricing/credit management the author observes over and over again the same "an error…of an old kind". Optimization developers make erroneous recommendations based on correlations but not on analysis of causalities. Analyzing a portfolio they try to build a series of predictive models – revenue, losses… Sometimes these models are quite inventive and often of a little help to find the true optimal solution. This flawed optimization process is quite typical for such developers.

The conventional predictive modeling is mostly about capturing correlations. If we want to model the causality we have to approach it through analysis of experiments where effect/uplift modeling plays the key role (see Medvedev, 2016 for the initial discussion).

In 2006-2009 the author disputed with some analytical vendors on optimizing pricing/credit at adjudication for cards portfolio. This was the time when the author set up the very first experimental design to optimize credit limit offers for cards. For some adjudication segments, the data showed a clear negative correlation between offered credit and the average segment balance after 12 months like in
Figure 1.

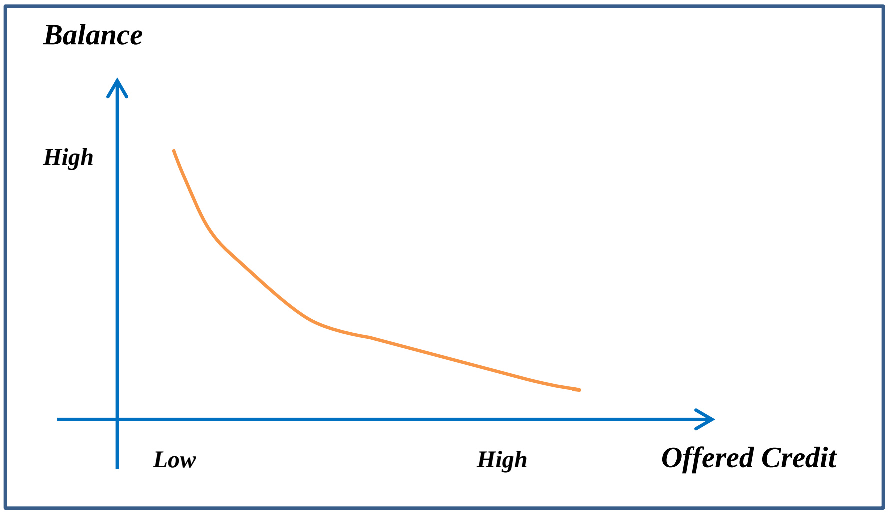

Figure 1. Cards Adjudication. Invite-to-Apply Segment. 12-month Balance by Offered Credit.

The business objective here is to grow balances for the segment. Building a conventional predictive model with "balance" as the target and predictive variables that include "offered credit" was not difficult. However, does the business want to be using this model that recommends decreasing offered limits to increase balances? Explanations behind this negative correlation are quite common. Higher credits were offered to low-risk customers where competitive pressure is very high. These low-risk customers have already several cards and do not rush to use the new one. On the other end, low credits were offered to mid-high risk customers. Activation and usage rates for these customers are high.

The experimental design analysis using the effect/uplift modeling optimization tool gave recommendations driving the objective of profit and more proper recommendations. Some low-risk customers got even higher and more competitive credit offers while special sub-segments in the mid-high risk region were recommended an increased or decreased offer. Similar negative correlation effects were observed between balances and offered rates for some line of credit adjudication segments. Here also higher rates are correlated with higher usage due to a higher customer risk.

The experimental design is the essential step for the development of pricing/credit optimization solutions and any candidate for an optimization solution must be validated on the proper experimental design test. The author's approach of effect/uplift based optimization methodology uses the experimental design both for development and validation steps.

The lack of understanding the difference between "correlation" and "causality" and using exclusively conventional predictive modeling to build optimization solutions represents the conventional approach in the financial industry. As stated, the conventional predictive modeling captures mostly correlations. If optimization via model causalities of actions/treatments, pricing/credit treatments in particular, is the desired outcome, then the effect/uplift modeling on experimental design tests is the way to go.

"No experimental design" means "no customer pricing/credit or other optimization".

In the following sections, we go in detail through our pricing optimization methodology with a specific example of interest rate increase optimization solution for a portfolio of unsecured lines of credit. To illustrate concepts and cases, in this paper we start with a simplified version of the price increase optimization. The optimization is explained only for the 0.5% interest rate increase which is a quite common price adjustment for lending portfolios. The next step when we want to optimize rate increases within a certain range, for example from 0.1% - 2.5%, will be a subject for subsequent papers.

From the author's direct experience, unfortunately not many industry pricing managers measure and analyze actual effects of this common rate increase on different customers and parts of the portfolio. As we see that a significant part of a regular lending portfolio can show from zero-to-very-negative profitability effects due to 0.5% rate increase. The author is not convinced that it is a common knowledge for financial industry pricing managers. Proper optimization of 0.5% rate increase treatment can nearly double positive revenue effects and practically eliminate the balance attrition.

Thus, the target or objective function for us is the profit on the portfolio of unsecured lines of credit. The pricing tool, we would like to optimize, is a rate increase of 0.5%.

In some cases when the portfolio pricing is very far from the optimal frontier, for example, if portfolio prices are significantly less than the market, there is no need in experimental design to justify small pricing adjustments up. This is a very rare case of portfolio's already poor pricing management. Portfolio price managers try to be in a reasonable distance from the optimal market pricing frontier which is very hard to identify clearly. In this case, the required foundation and the starting point for the pricing optimization is a proper experimental design test.

PRICE ELASTICITY APPROACH. HOW CAN IT HELP?

Since 1890 when Alfred Marshall introduced the notion of the price elasticity of demand, this concept is the key to develop price optimization solutions. In the lending business, we sell "*balances*" at a certain price "*p*" which must be special for every customer. This is how we can realize the customer level pricing optimization. If "*B*" is the size of customer's borrowing balance then the product $R = Bp$ is the point estimate for the revenue function. The revenue function attains its maximum at the price point where the revenue function derivative $\frac{dR}{dp} = 0$ or equivalently the price elasticity $E_p = -\frac{dB}{dp} \times \frac{p}{B}$ equals *1*. This is the microeconomics approach where the balance as the function of the price $B = B(p)$ is the critical input component.

The author finds that building an accurate customer level price-balance function is an unrealistic and very significant challenge for lending portfolios as well as for other retail portfolios.

The next challenge of applying the traditional microeconomic price optimization approach is related to the objective function. For the pricing optimization, it is commonly the revenue function. The traditional optimization approach typically builds revenue projections given the optimal price and it is another great challenge to assess the accuracy of these revenue projections. On the top of pricing changes, portfolio managers run numerous initiatives which are supposed to stimulate and grow the portfolio, however, with no proper experimental design, the actually observed portfolio revenue metric is not isolated and too noisy to assess accurately.

The author's effect/uplift pricing optimization approach runs several experimental design pricing tests, focusing on building customer level effect projections due to the tested pricing changes. We launch optimal pricing changes and monitor the effects with the new experimental design, the only way to accurately measure the optimization results.

II. THE CASE STUDY: PRICING TEST AND ITS EXPERIMENTAL DESIGN.

We now go into detail through author's price optimization methodology using the test campaign that was launched in April 2014 on a limited population of unsecured credit lines. The set of optimization steps turns out to be fairly standard. The author uses this approach to optimize customer pricing and credit offers. Evidently, the same methodology can be used to optimize practically any customer offer and in any customer based business.

INTEREST RATE INCREASE PRICING TEST CAMPAIGN

- Portfolio: Unsecured Credit Lines.
- Low-Mid Risk Active Accounts.
- Treatment: 0.5% Interest Rate Increase.
- Experimental Design: 50% random accounts were treated with 0.5% interest rate increase and 50% accounts were held as a control.

The experimental design for the whole campaign is quite complex and includes testing of pricing changes from 0.3% - 2%. To illustrate our optimization methodology we examine only the specific experimental sub-segment where 0.5% rate increase was tested.

The Bank of Montreal lending portfolio is quite generic for the Canadian lending business and therefore extrapolation of findings of this paper to other Canadian financial institution generic portfolios is feasible.

MEASURING EFFECTS

The performance monitoring is set up by month from the beginning of the offer/pricing change. The campaign population was selected in March 2014 (month=0) and treated with a 0.5% rate increase message in April 2014 (month=1) and we compare numerous performance metrics for treatment (T) and control (C). The following figure illustrates differences in performance for the balance metric by month for treatment and control experimental cells.

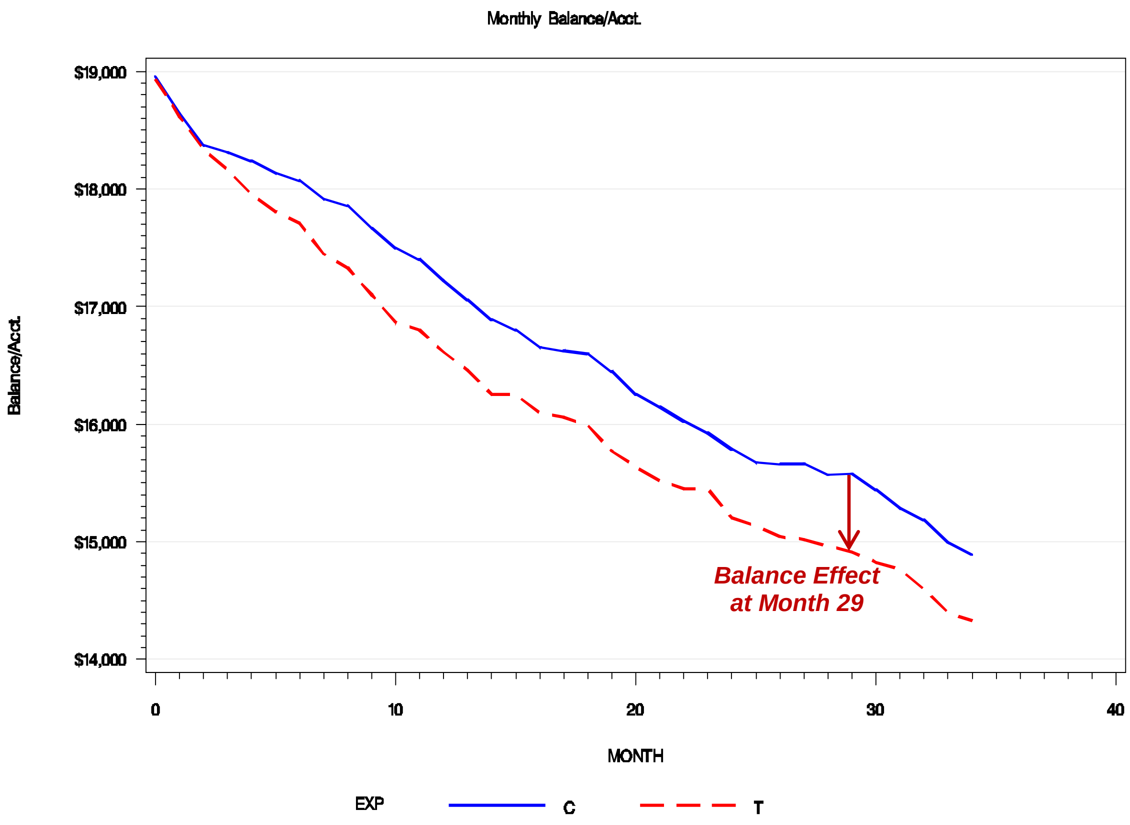

Figure 2. Monthly Balance/Account: Treatment vs Control.

The declining balance trend is very typical for active lending customers as these customers activate by borrowing a significant sum once, paying this balance off after that. The interest rate increase is a negative treatment for customers. Due to the randomization, at the starting month balance values for treatment and control cells are near identical. Going forward we observe negative balance effects due to the rate increase treatment. Given a metric such as balance, we define the balance effect at a certain month as the difference between the average balance values on treatment and control groups.

PROFIT EFFECT

The profit is the main target/objective function for our optimization example. The following finding is very critical on the way to develop an optimal solution that maximizes the profit effect.

Losses and cost effects due to 0.5% interest rate increase are insignificant. Therefore, revenue effects can be used to define profit effects for our pricing test.

For many years the author conducted plenty of pricing and credit change tests on lending customers. This critical finding is actually observed on all low-mid risk pricing tests conducted by the author. We have to affirm that our pricing/credit changes were within a reasonable range. We never erroneously impacted our customers with pricing or credit changes.

There exists an "industry belief" that increasing customer's pricing or load with additional credits we can escalate customer`s risk and may push the customer to a default sooner.

Analyzing lots of pricing/credit tests on low-mid risk customers the author could not find statistically significant and replicable evidence that tested price and credit changes escalate customer's risk. Default rate effects are insignificant due to numerous tested price/credit changes.

Let's elaborate more on insignificant losses effects. We observe that if a customer charges off then likely the customer charges off with the entire credit limit. Given customer's credit limit, his/her expected losses can be accurately estimated by customer's expected default rate. Pricing changes do not affect customers' credit limits, therefore if we do not observe effects on default rates then losses effects should be also insignificant.

Our pricing test is not an exception. On the following graph, we can see that charge-off rate curves for treatment and control groups are indeed very near.

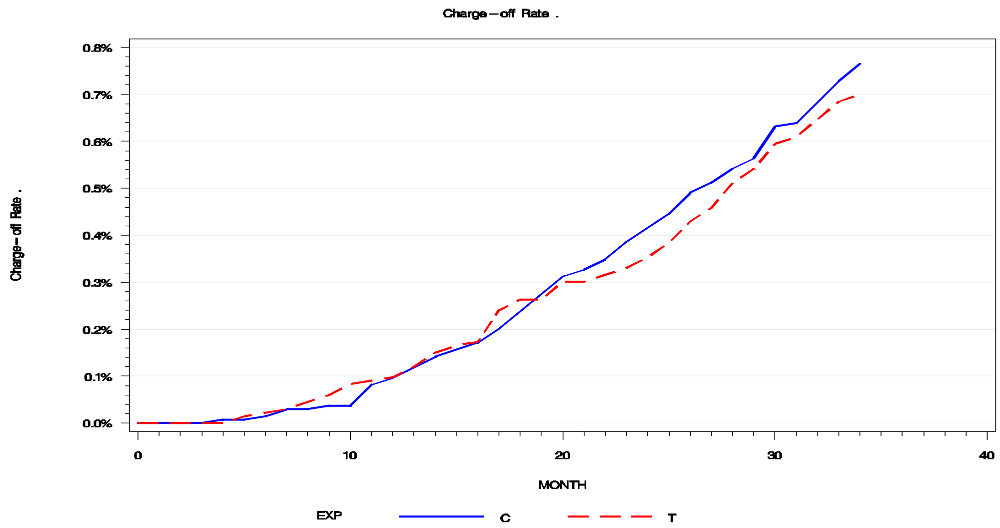

Figure 3. Charge-off Rate Curves for Treatment (T) vs Control (C) Cells.

The treatment group has even slightly less observed charge-off rates after almost three years of performance. Looking at 95% confidence interval around charge-off rate negative effect we can truly see that this negative effect is statistically insignificant.

Our main objective is to optimize the profit. This critical finding allows us to focus now exclusively on the revenue optimization.

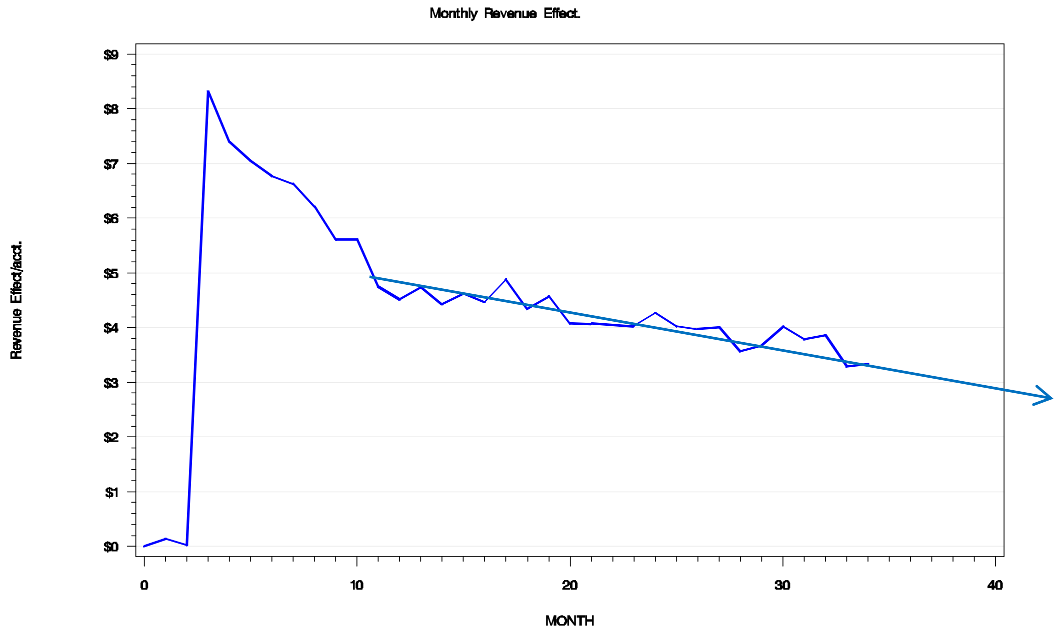

Figure 4. Monthly Revenue Effects/Account due to 0.5% Rate Increase Treatment.

The treatment group shows positive revenue effects over the control for almost three years of performance monitoring. However, we can observe a significant decline in monthly revenue effects of 60%, from $8.3 to $3.3 and effects are still declining. This suggests that there is a significant price sensitive population that is not happy with 0.5% interest rate increase treatment that reduces their balance or takes it somewhere else. Given the clear declining revenue effect trend, in one to two years, the revenue effect can move even further into negative territory and hurt the business in the long run. Therefore the quality of the treated population can be damaged significantly. Often the author asks pricing managers if they think and try to assess the impact of pricing initiatives beyond one year. It is unlikely to make such assessments if proper experimental designs were never set and analyzed.

Some customers can handle 0.5% rate increases and practically show no signs of changing behavior, and as part of the pricing optimization development in this paper, we identify such customers. Some targeted customers can handle even higher rate increases while others cannot take 0.5% rate increase stress but will tolerate a 0.3% increase. The author will address these cases in subsequent papers. The objective of this paper is to illustrate the price optimization methodology and its main development steps in detail based on the given test campaign. The customer level effect/uplift modeling is the critical tool for author's optimization solutions. We use the effect/uplift modeling technique to identify sub-segments that are very price sensitive, sub-segments that are less care about price changes showing minimal negative effects and sub-segments that can take the price stress easily.

III. PRICE SENSITIVITY EFFECT/UPLIFT MODEL.

The author refers to (Medvedev, 2016) paper for a more detailed introduction to the effect/uplift modeling concept. The effect is a segment concept. Therefore it is worth to remind that the vital difference for the effect modeling from the conventional modeling is that the effect modeling is defined on the segment level only. The effect modeling tool is specially designed to drill and analyze into experimental designs. Using the effect/uplift modeling tool we can identify sub-segments that generate significant positive, negative or no effects.

THE BEST EFFECT MODEL TARGET TO OPTIMIZE CUSTOMERS' PRICING.

The performance window for our test campaign covers 34 months. Optimizing the profit effect is equivalent to optimizing the revenue effect. We can choose several versions of the revenue based target for the effect modeling step. These versions can be based on measurement windows.

Single Period Window: The Customer/Account Revenue at the Last Performance Monitoring Month.

This is so-called a "snap-shot" target – the customer/account revenue at the last month 34 in our example. Given the long-term pricing optimization objective, the author finds this target not quite effective. Two customers can arrive at the same month-34-revenue value having very different trajectories.

Cumulative Period Window: The Cumulative Customer/Account Revenue for the Performance Monitoring Window.

The cumulative revenue target is significantly better for the long-term pricing optimization on the customer level. This is the target that we use to train the effect model presented in this paper. Given a customer/account, we summarize all revenue components for the performance window of 34 months. Customer's revenue includes interest and non-interest revenue parts, fees, reward/promotion costs and all other revenue components.

Projected Cumulative Period Window: The Projected Cumulative Target is the Best Option to Optimize Customers' Pricing.

The concept of the projected cumulative target is beyond the scope of his paper. We will address it in subsequent papers. The projected cumulative target was tested by the author on numerous optimization developments and is author's preferred optimization target. Here are rationales behind the projected cumulative target concept. Think of two customers/accounts having the following monthly revenue trails.

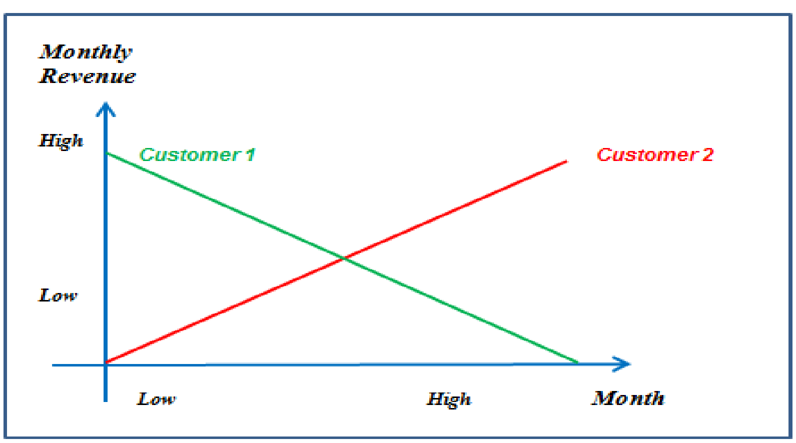

Figure 5. Monthly revenue graphs for two customers.

The cumulative revenue is the area under the graph. Thus these two customers have the same cumulative revenue values for the observation window. On the other hand customer 1 has already paid-off his balance and therefore can be less sensitive to price changes. Customer 2 is building his balance and may be very sensitive to price changes. Optimizing price for these two different customers will require different strategies.

The author developed a methodology that can effectively project customer's performance for the next several years. Using this projected cumulative target for 5 – 10 years window can really separate these two and similar customers.

INPUT PREDICTIVE VARIABLES.

The list of input or predictive variables includes all customer/account traditional metrics for the at least 12-month window. Both endogenous and exogenous variables (Equifax/TransUnion…) are on the list, as well as bank relationship data, demographic data and more. Many of these metrics have a time-series structure; for example, balances or payments on cards trades. In addition, the author finds that derived variables converting time-series metrics into new predictive variables can provide a significant lift to the effect model. The modeling begins with about 300 original input metrics that have a time-series nature and grows to a total list of over 10,000 with the inclusion of derived variables.

TRAINING THE EFFECT MODEL.

The decision tree effect modeling algorithm that was described in (Medvedev, 2016) is utilized in this project, with the cumulative revenue for 34-month observation window to train the model. The decision tree effect model splits the whole test population into 15 final nodes. Every node inherits the original experimental design with the randomly selected 50% treatment and 50% control cells. Revenue effects are measured for every node and quite different by nodes. As we've mentioned before, this effect model is the critical element in the final optimization solution.

IV. EFFECT MODEL VALIDATION. FINAL NODES AND THEIR PERFORMANCE.

The model was trained with 50% development population for the effect score/model, and 50% holdout sample for model validation. Cumulative revenue effect trajectories by effect score segments on the validation data are near to the trails on the development data.

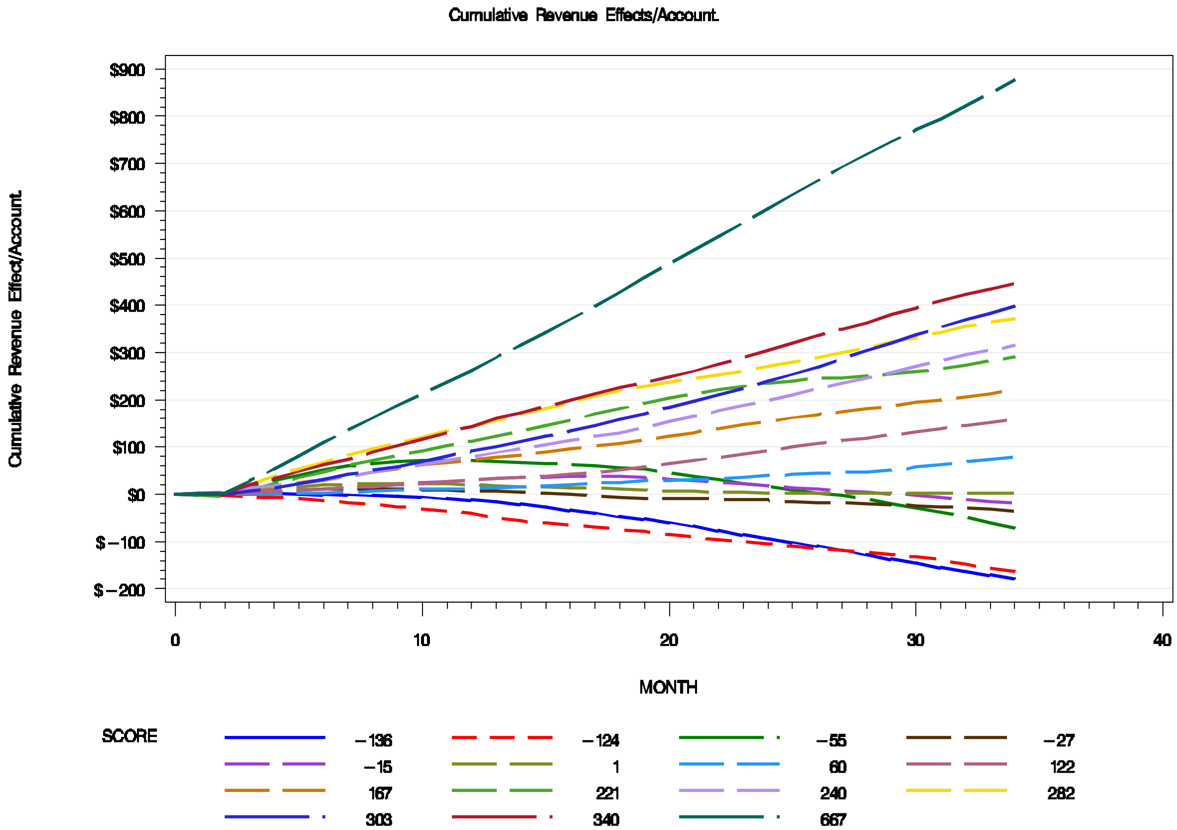

Figure 6. Cumulative Revenue Effect Trajectories by Effect Score Segments on the Validation Data.

From the model output, we can observe clearly there are five negative effect score segments. The effect score segment 1 is still in the positive territory but also trending down. If the business objective is to optimize the total revenue effect for 3+ years then we have to exclude all negative effect score segments and even the score 1 segment.

The effect modeling identifies 9 positive effect score segments where practically no balance attrition or other changes in customers' behavior are observed. These customers are more price elastic and therefore can absorb a price increase of 0.5%. The focus of the pricing optimization is on six price sensitive effect score segments - five negative effect score segments and the score 1 segment. Customers from these six segments are not accepting rate increases and causing attrition in balances (negative balance effects) vs the control group.

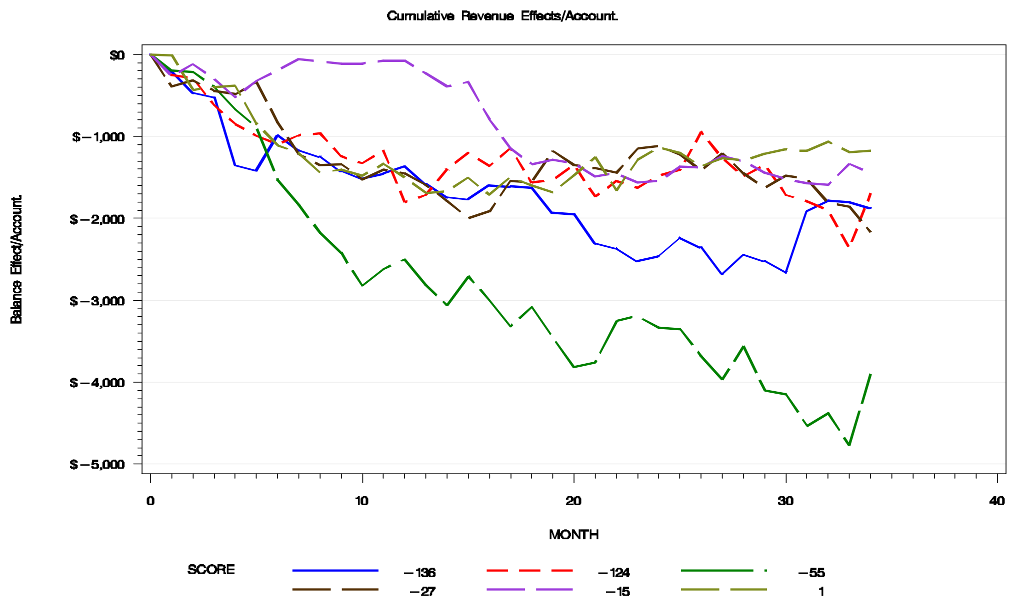

Figure 7. Balance Effect Trajectories by Price Sensitive Effect Score Segments.

V. SIMULATION BOX OPTIMIZATION TOOL.

The author has already described briefly the simulation box SAS macro concept in (Medvedev, 2016). Here we illustrate with more details how the simulation box works to optimize pricing decisions.

Business decisions on pricing/credit actions must be made based on long-term effects of said actions, preferably lifetime. As we discussed before the performance window for effect models is often limited to the 12 to 24 months, however, in our example, the performance window is extended to 34 months to yield a more optimized result. Given a set of customers/accounts, the simulation box generates all necessary effect projections (up to 10 years) based on customers' effect scores and input effect curves for business metrics of interest: revenue, cost, losses, and their components.

CUSTOMER PSYCHOLOGY: THREE YEARS TO ACCEPT OR REJECT THE CHANGE.

The following finding on shapes of effect curves was discussed in (Medvedev, 2016), based on almost 20 years of tests and analyses. The finding is fundamental to extend projections of effects beyond the limited effect model performance window. As the lending business is about managing balances, after applying a credit/pricing treatment we can observe three distinct phases: Phase 1 – Balance effect building; Phase 2 – Stable balance effect; Phase 3 – Declining balance effect. Phase 1 is the most critical phase and depending on a target population and a treatment can last from 12 – 36 months. Based on our pricing/credit effect measurement experience, phase 2 can last over 10 years, while with phase 3 there is a limited real knowledge. Given the length of these phases, the simulation box application is designed to extend effect projections up to 10 years.

The fact that the balance effect building period (phase 1) is shaped within approximately the first three years is very critical and characterizes customer's psychology.

This finding is indeed very critical and characterizes customer's psychology. Effects are building up when new customers decide on actions triggered by pricing/credit changes. It appears that if a customer did not react to a specific pricing/credit change within the first three years then the customer accepts it.

Customer psychology finding: three years to accept or reject a pricing/credit change.

As stated earlier, for our specific pricing example, we do not observe any effects on cost and losses metrics and therefore our focus is on revenue effects. The author designed a similar pricing test back in August-September 2011. On the following graph, the revenue effect trajectories for four effect modeling segments are observed.

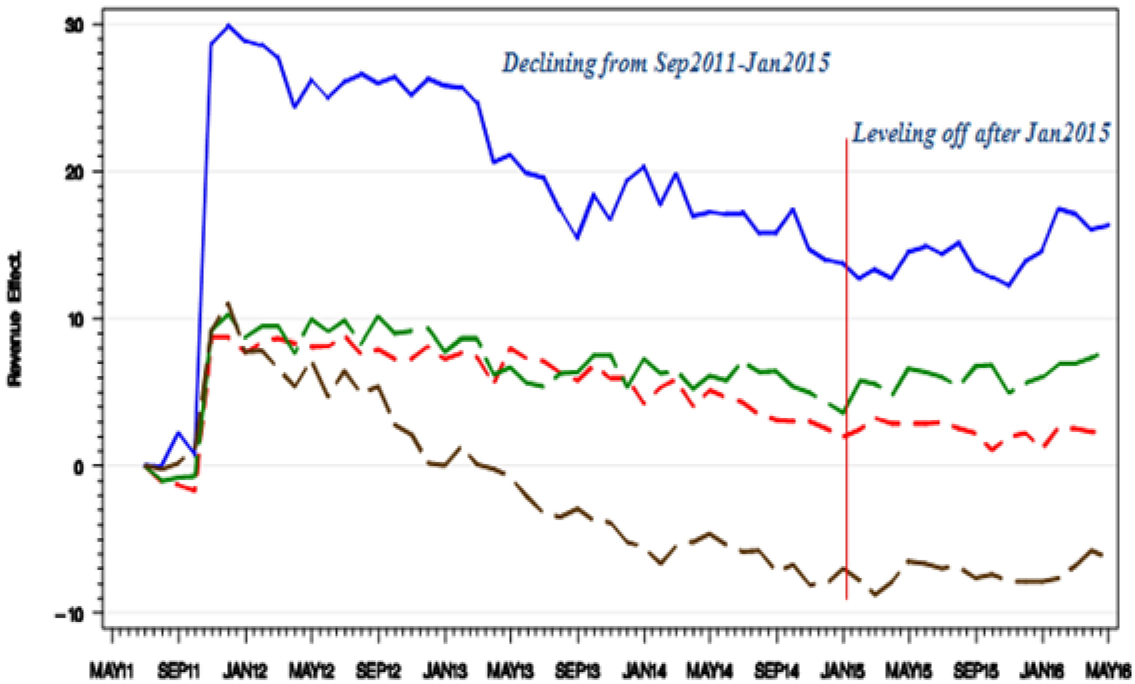

Figure 8. August-September 2011 Pricing Test. Revenue Effect/Account from Aug2011 – May2016.

All effect modeling sub-segments clearly demonstrate the declining trends for just over three years (40 months), and then stabilize, showing that customers are set up their minds on price/credit changes within these first three years. We use this critical finding to tune up the simulation box for our specific example and build revenue effect projections beyond the observed performance window of 34 months.

Customer's revenue effect score defines revenue effect trajectories for the observation window of 34 months. As per our finding, we can use the regression and extend effect trajectories for 40 months. We assume that effect trajectories are leveling off after that. Here is an example of genuine and projected revenue effect trajectories for the effect score segment -55 generated by the simulation box.

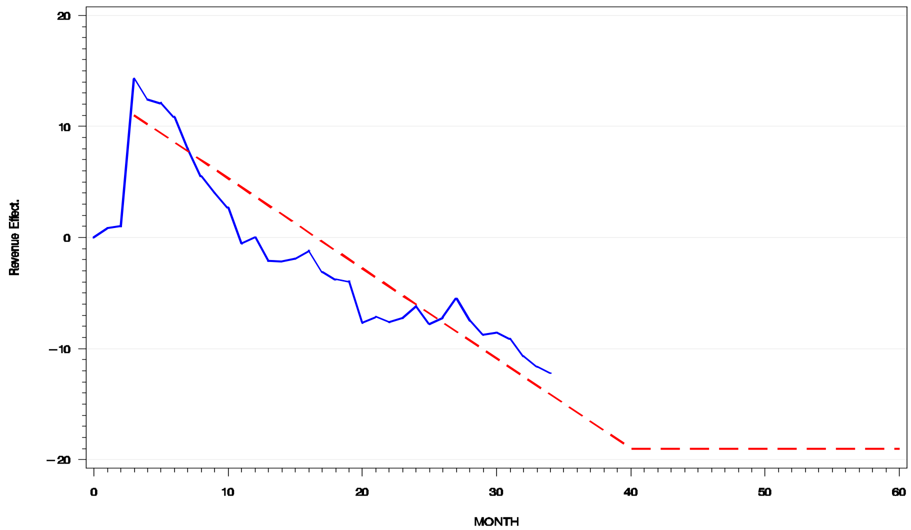

Figure 9. Effect Score Segment -55. Observed and Projected Revenue Effect Trajectories.

The similar method is used to generate other effects, for example, balance effects. As we observed that there are no cost and loss effects due to the pricing change. We load actual or observed effect curves for 34 months for key business metrics of interest into the simulation box which generates the projected effect curves up to 10 years for every effect score. Cumulative revenue effect trajectories for 10 years are critical to finalizing the pricing optimization given business objectives and constraints. On the next graph, one can see cumulative revenue effect trajectories generated by the simulation box for every effect score.

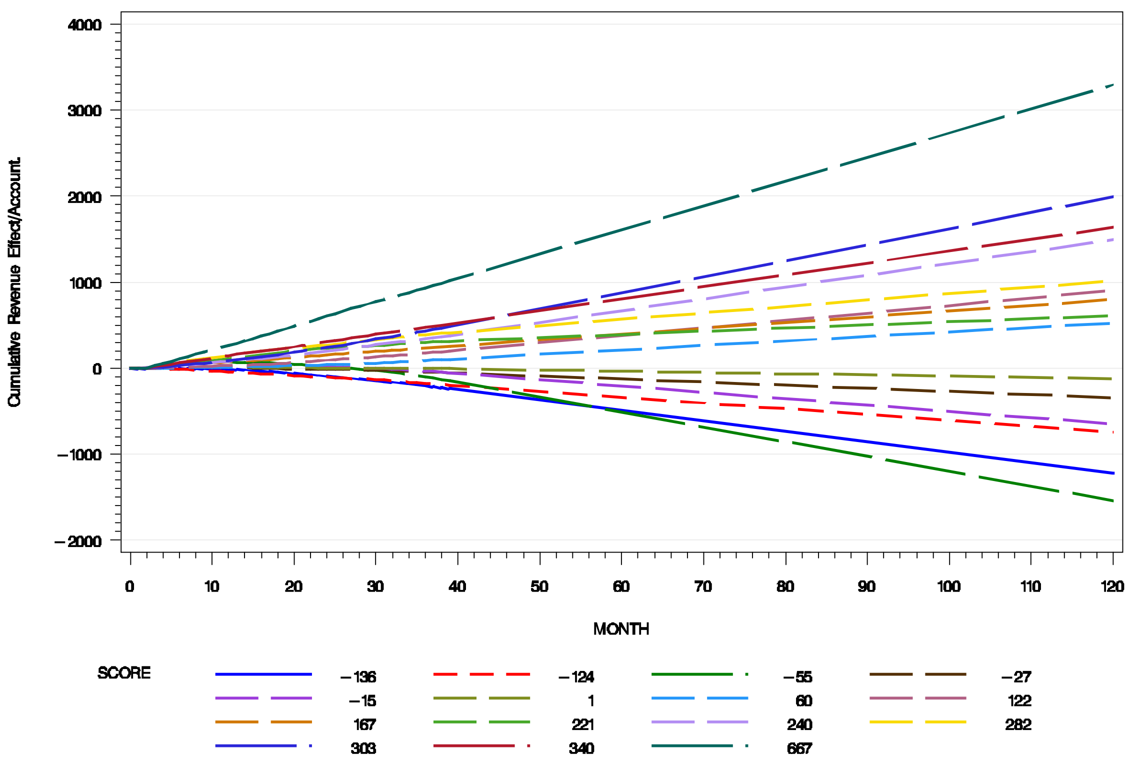

Figure 10. The Simulation Box Output: Cumulative Revenue Effect Trajectories by Effect Scores.

VI. IDENTIFYING THE OPTIMAL PRICING SOLUTION.

The simulation box solution allows for the business to specify its pricing objective; for example, it can be maximizing the profit effect over 12, 24 or any number of months from 0 -120 or a lifetime period. If we want to maximize the total profitability effect over the ten-year period then the simulation box suggests looking at the cumulative effect lift chart where the ten-year cumulative revenue effect is the target. To build the cumulative lift chart on the segment of interest the simulation box sorts the segment population from highest-to-lowest effect scores. The x-axis shows the percentage of the population with the highest scores. Given the percentage x, the y-axis shows the total ten-year cumulative revenue effect on the x-percentage population. On the next graph, we can see the revenue effect cumulative lift chart developed on the validation population.

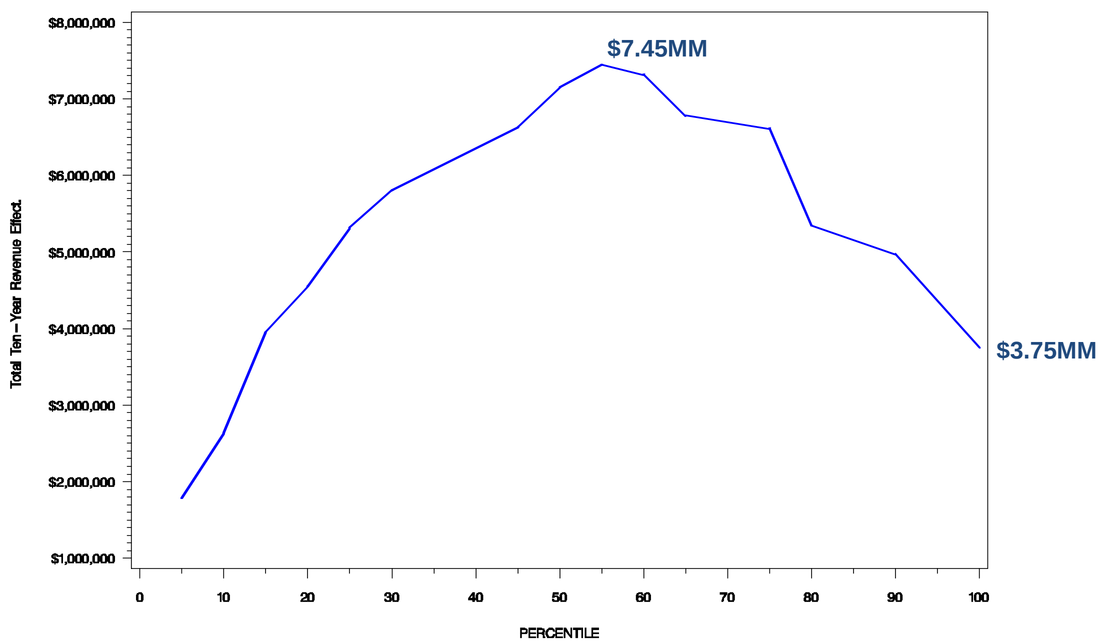

Figure 11. The Total Ten-Year Revenue Effect Cumulative Lift Chart

If we target the whole (100%) population under the simulation box evaluation then the total estimated ten-year revenue/profit effect is $3.75 million. If the business objective is to maximize the ten-year total profit/revenue effect then the cumulative lift chart recommends targeting only 55% of the population or all customers with effect scores greater than 1. The cumulative lift chart attains its maximum of $7.45 million at 55% point. We call this maximum ***the power*** of the effect model (see Medvedev, 2016). One can assess different effect models by comparing their powers.

The pricing optimization on this particular validation population brings substantial improvements. Targeting only customers with an effect score>1 almost ***doubles*** the total profit/revenue effect boosting it from $3.75MM to $7.45MM. Excluding the negative effect score customers practically eliminates the balance attrition due to the pricing increase.

The cumulative revenue trajectories for strongly positive effect score segments stay in the positive territory since the launch of pricing changes. These are customers that are not price-sensitive and can take pricing changes trouble-free. The optimization effort is all around handling properly negative and low positive effect score customers.

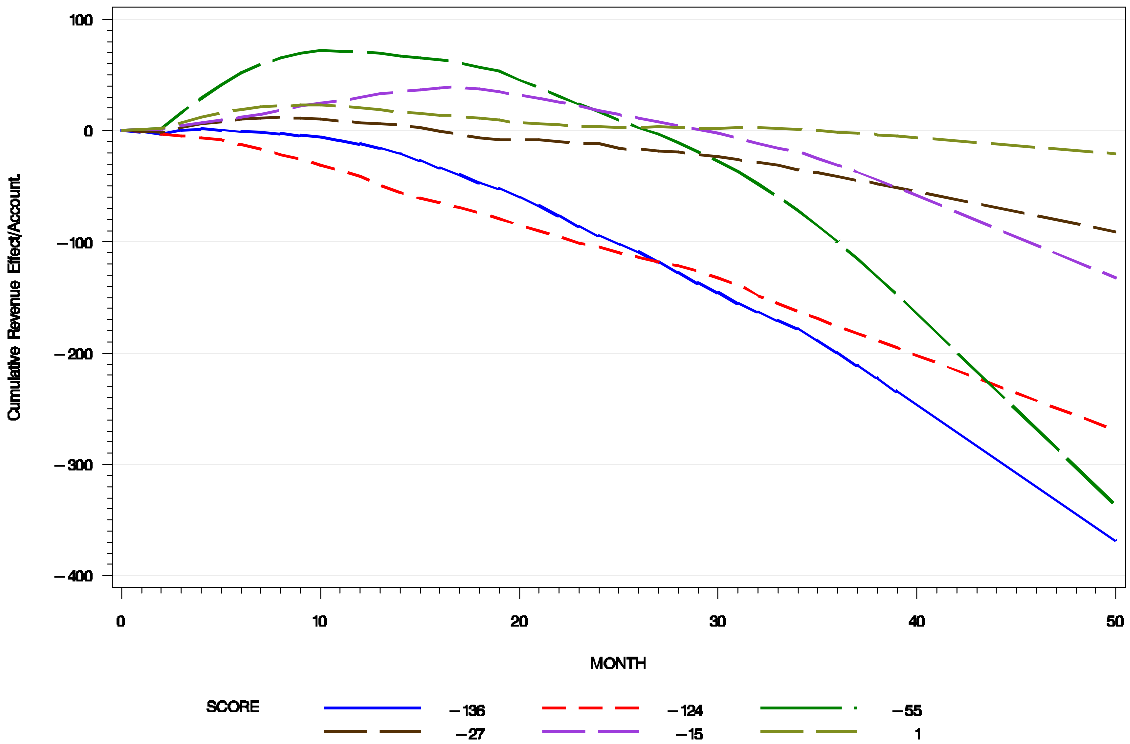

Figure 12. Negative and Low Positive Effect Score Customers: Cumulative Revenue Effect Trajectories.

Two effect segments -124 and -136 are obviously “negative” since the very launch of the pricing change. Therefore in any optimization scenario, these segments have to be excluded from pricing up. The rest effect score segments move in the negative territory within first 20 - 30 months.

If the business goal is to maximize the revenue effect for the next 12 months only, one would have to exclude only super price-sensitive customers with effect scores -124 and -136. Such short-term objectives are not rare events for some financial managers. In the long-run, optimization solutions based on short-term objectives can be very sub-optimal and even hurt the business.

Scenario	Optimization Objective	Optimal Target Population	10 Years Cumulative Revenue Effect (MM millions)	Growth over No-Optimization Scenario #1
1.	No Optimization Objective	100% customers.	$3.75MM	0%
2.	Maximize 12-month Revenue Effect.	Target 80% of customers. Exclude effect segments -124 and -136.	$5.33MM	42%
3.	Maximize 24-month Revenue Effect.	Target 70% of customers. Exclude effect segments -27, -124 and -136.	$6.78MM	81%
4.	Maximize 36-120-month Revenue Effects.	Target 55% of customers. Exclude effect segments 1, -15, -27, -55, -124 and -136.	$7.45MM	99% !

Table 1. The Simulation Box Scenarios on the Validation Population.

Maximizing 12-month revenue effect can increase the total long-run revenue effect by noticeable 42%. However, the long-term optimization, targeting only customers that can take pricing changes, can increase the revenue effect by remarkable 99%.

CONCLUSION

EXPERIMENTAL DESIGN IS CRITICAL

The critical conclusion of this paper is that without an experimental design, there is no customer level pricing/credit optimization. The experimental design is the foundation to compare two or more pricing/credit strategic solutions and is the critical tool for an optimal pricing/credit solution. This represents a significant opportunity in the financial industry today.

The author did numerous projects implementing an experimental design based optimization solution for pricing/credit strategies on lending portfolios and campaigns even demonstrated negative profitability effects. This paper goes on showing how optimization solutions can turn low positive or even negative effect tests into quite positive portfolio strategies as it’s illustrated with the specific pricing example of 0.5% interest rate increase on the low-mid risk population of unsecured credit lines.

CAUSALITY NOT CORRELATION

Implementing the conventional predictive and microeconomic approach of price elasticities to build pricing/credit optimization solutions leverages mainly correlation and not causality. This approach is widely utilized by the financial industry does not concur with customer level optimized outcomes. However, as building an accurate customer level price elasticity function is unrealistic and presents a significant challenge, the effect/uplift modeling methodology is the direct way to model causalities (see Medvedev, 2016 for the initial discussion).

OPTIMIZATION LEADS TO 99% INCREASE IN PROFIT

By implementing the optimization solution and its validation on actual experimental design test campaigns, we can measure and see profitability effects by targeting the optimal population vs the whole population. The effect/uplift modeling identifies price-sensitive customers and customers that can handle pricing changes. By targeting the optimal population increases the profit effect by 99% and practically eliminates the balance attrition.

As demonstrated with the optimization example that maximizing 12-month revenue effect can increase the long-run revenue effect by noticeable 42%. However, the long-term optimization, targeting only customers that can take pricing changes, can increase the revenue effect by remarkable 99%. This is the power of the customer level validated optimization. There are significant benefits for the financial and other customer based industries to adopt it and follow.

The effect/uplift modeling that ranks customers by their price sensitivity is the most critical component of the customer level optimization methodology. While portfolio-wide price increases by 0.25% or 0.5% are common pricing adjustments for the lending business, optimizing pricing changes at the customer level can increase profitability effects at such remarkable rates.

PRICE INCREASES DO NOT LEAD TO INCREASED CUSTOMER RISK

Analyzing numerous pricing/credit tests on low-mid risk customers the author finds that there is no statistically significant and replicable evidence that tested price and credit changes can escalate customer’s risk. In fact, default rate effects are insignificant due to tested price/credit changes. Therefore, low-mid risk customers find better ways to react to pricing changes rather than to default. Low-mid risk customers usually have a choice to take their balances to competitors if they are not happy with new prices.

'THE SIMULATION BOX' – IMPLEMENTING OPTIMIZATION THROUGH SOFTWARE

The Simulation Box enables the developed effect model to identify price sensitivity sub-segments where we can observe actual effect curves for business metrics of interest – revenue, cost, loss components, balance, attrition, et al. It also enables the ability to adjust the observation window for optimization to extend out beyond the typical 12-24 month and build long term observations up to 10 years to enable a much-improved optimization recommendation. Depending on the business objectives and constraints the simulation box identifies the best customers to target.

REFERENCES

Medvedev, Yu. (2016). "Personal Lending: Customer Credit and Pricing Optimization", SAS Institute Inc. 2016. Proceedings of the SAS® Global Forum 2016 Conference. Cary, NC: SAS Institute Inc. Paper 2600-2016. http://support.sas.com/resources/papers/proceedings16/2600-2016.pdf

RECOMMENDED READING

Larsen, K. (2010)." Net Lift Models: Optimizing the Impact of Your Marketing Efforts." SAS Course Notes, Cary, NC, SAS Institute Inc.

Lo, V. (2002)." The True Lift Model: A Novel Data Mining Approach to Response Modeling in Database Marketing". ACM SIGKDD Explorations Newsletter, Vol. 4, No. 2, 78–86.

Radcliffe, N. J. and Surry, P. D. (1999). "Differential Response Analysis: Modelling True Response by Isolating the Effect of a Single Action". Proceedings of Credit Scoring and Credit Control VI, Credit Research Centre, University of Edinburgh Management School.

Radcliffe, N. J. and Surry, P. D. (2011). "Real-World Uplift Modeling with Significance-Based Uplift Trees". Portrait Technical Report TR-2011-1, Stochastic Solutions.

Radcliffe, N. J. (2007). "Using Control Groups to Target on Predicted Lift: Building and Assessing Uplift Models". Direct Marketing Analytics Journal. An Annual Publication from the Direct Marketing Association Analytics Council, pages 14-21.

Rzepakowski, P. and Jaroszewicz, S. (2012). "Decision Trees for Uplift Modeling with Single and Multiple Treatments". Knowl Inf Syst (2012) 32:303-327.

ACKNOWLEDGMENTS

Writing this paper was an enjoyable experience. All ideas and concepts presented in the paper were discussed in detail with my analytical partners at Smith School of Business at Queen's University. I wish to personally thank the following people.

Andrei Bazhanov, Research Associate, Smith School of Business at Queen's University.

Yuri Levin, Professor & Stephen J.R. Smith Chair of Analytics, Smith School of Business at Queen's University.

Mikhail Nediak, Distinguished Professor (Associate) of Operations Management, Smith School of Business at Queen's University.

I would like to thank **Chris Osborne**, leading analytics executive and my former colleague at Bank of Montreal for his significant contribution, ideas and editorial support in this paper. I was fortunate to have **Chris** as well as my another former colleague from Bank of Montreal - **Allan Esser**, Professor George Brown College, with their critical suggestions.

As noted in the conclusion, this paper serves as a foundation and stimulus for a new solution in customer pricing/credit and risk optimization. These ideas are being leveraged by **Matt Fabian**, Director, Leading Global Information Services and Analytics firm, and by **Chris Osborne**. The approach described in this paper is being leveraged to guide the potential design of the very first true customer pricing/credit and risk optimization **AI Platform**. This AI Platform will provide financial institutions with optimization solutions at the customer level for any lending product. The **Simulation Box** will have the capabilities to run customer or segment level of analyses for interest and measure pricing effects, identify best customers for pricing/credit and other features changes. The experimental design concept developed in this paper will be the critical part of the Simulation Box AI Platform that will guarantee market performance matches recommended actions.

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the author at:

Yuri Medvedev

Bank of Montreal, Toronto, Canada

E-mail: yuri.medvedev@bmo.com