Paper 2600-2016

Personal Lending: Customer Credit and Pricing Optimization

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

ABSTRACT

Financial institutions are working hard to optimize credit and pricing strategies for both adjudication and ongoing customer account management. Intense competitive pressures have generated increased revenue challenges for financial institutions. In response to these forces arising within the industry, there is a significant demand to improve the sophistication of methods to manage both the credit exposure and pricing. Numerous credit and pricing optimization applications are available on the market to satisfy these needs. We present a relatively new approach that applies an effect modeling technique on continuous target metrics. The effect modeling method (also referred to as uplift or net lift) can be applied to various continuous targets including revenue, cost, losses, or profit,

Examples of effect modeling to optimize the impact of marketing campaigns are known. See Radcliffe and Surry (2011) for a history and literature review.

We discuss essential steps on the credit and pricing optimization path: (1) setting up critical credit and pricing champion/challenger tests, (2) performance measurement of specific test campaigns, (3) effect modeling, (4) defining the best effect model, and (5) moving from the effect model to the optimal solution. These steps require specific applications that are not easily available in SAS®. Therefore, necessary tools have been developed in SAS/STAT® software.

We go through numerous examples to illustrate our credit and pricing optimization methodologies and solutions.

INTRODUCTION

An optimization approach relies heavily on creating an underlying ‘fact base’ of treatments and effects. This ‘fact base’ must be developed through a rigorous experimental design framework of treatment and control testing of various credit and pricing management strategies. Once treatment effects are measured and validated through repeated trials, the results can be fed forward into an optimization routine that can consider the pricing and management tactics in a broader strategic context. Through the application of optimization methods, pricing ‘resources’ can be more efficiently allocated across various behavior segments to maximize the goals of the enterprise. Resources can be represented through a variety of tactics including activity based pricing, risk based pricing, behavior pricing, limit management, and accept/reject approval cutoffs. As well, specific targets (or goals) can be set such as expected losses, revenues, net profit, or a combination of continuous variables across time.

A similar methodology can be applied to optimize fees, reward rates or other pricing components. The examples in this paper focus on credit and interest rate optimization. This particular optimization methodology is not that known for the financial industry. Therefore in this paper we have to start with a general overview of the methodology.

The paper is organized the following way. We begin with a case study that illustrates the importance of using proper measurements that consider treatment effects against a representative control group. Then we define and illustrate the notion of the treatment effect (uplift by Radcliffe and Surry (1999, 2011), true lift by Lo (2002), net lift by Larsen (2010)…). The role that effect modeling (uplift, true lift, net lift…) has on developing and implementing pricing and credit optimization strategies is explained. The mathematics behind the decision tree effect modeling SAS macro is explored, followed by discussions and examples of how price/credit sensitivity scores are applied to optimize the portfolio profitability.

The next section defines effect lift/gain charts. To compare effect models we use effect model power which is defined as the maximum value on the cumulative effect chart. Radcliffe (2007) proposed the Qini coefficient (an analogue of the Gini coefficient) to measure the effect model power. Larsen (2010) compares effect models using their performance in top deciles/quantiles. Our definition of effect model power is directly derived from business needs.

Effect modeling is an important but intermediate step in the whole optimization process. For a particular credit/pricing treatment, to make a decision on treatment's application ideally we need to know lifetime effect projections on profit and its major components such as revenue, cost, losses or other outcome factors. Given business constraints and objectives on the portfolio health such as loss-to-balance, revenue-to-loss, or profit-to-capital ratios, we can identify optimum population for this particular treatment. The simulation box is an application developed in SAS that generates projections up to 10 years, based on all input effect curves. The simulation box is the engine behind the final optimized solution. The concept of effect's fundamental shape is crucial for building the simulation box projections.

All examples in the paper are based on real data and real testing campaigns that were developed on credit cards and unsecured credit line portfolios at the Bank of Montreal. 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.

The decision tree effect modeling algorithm was developed and implemented in SAS back in 2001. The whole optimization methodology was shaped, finalized and implemented in SAS from 2001-2007. The methodology and SAS solutions were validated and tested on cards portfolio starting with 2002 and on unsecured credit lines starting with 2007.

In 2001 when the decision tree effect modeling algorithm and the SAS effect modeling macro were developed the author was not aware about Radcliffe and Surry (1999) results. At that time Radcliffe and Surry (1999) even used the term differential response modeling switching for the uplift modeling name later. It explains why we started using the effect modeling name in the Bank of Montreal since 2001. Giving the tribute to Radcliffe and Surry the author thinks that it'll be fair switching for the uplift name in the future.

The financial industry got comfortable using conventional modeling. There is a significant difference between conventional modeling and effect modeling. The author shares Radcliffe and Surry's (2011) concern that use of effect/uplift modeling has grown as slowly as it has. Effect modeling is about optimizing actions. However, conventional modeling scores can be helpful as input variables for effect models. Managing for example lending portfolios we have to adjust/change limits and pricing components according to new customer or external economic changes. All these pricing/credit adjustments are actions that require effect modeling optimization. We often can't identify even conventional responders for these actions. If the bank, for example, increases interest rates for the personal lending portfolio then everyone is a responder. The effect modeling helps to identify sub-segments of customers that reacted very negatively (vs the control) and sub-segments that practically did not change their behavior (vs the control).

The author has been using the effect modeling optimization approach for more than fourteen years. These ideas have been discussed with business and analytical leaders here in Canada and US. We have much unrealized potential in both industry and academia in the application of these techniques.

I. A CASE STUDY: XYZ-PRICING SOLUTION AT ACQUISITION.

As every bank the Bank of Montreal (BMO) has its proprietary interest rate pricing solution for customers that apply for unsecured credit lines. To maximize the portfolio profitability in 2008-2009 a challenger pricing solution – XYZ-pricing solution, was developed. For this specific example XYZ prices were about 1% higher than BMO standard prices. A preliminary analysis showed a positive incremental revenue due to the XYZ-solution. We still set up an experimental design to measure the incremental revenue of the XYZ-solution over the BMO standard one.

The objective of this experimental design is not just to measure the incremental revenue (the revenue effect) but also be able to optimize further the pricing solution.

XYZ-Pricing Solution Test Campaign.

- Portfolio: Unsecured Credit Lines.
- Launching Time: June, 2009
- 50% of applications are routed through BMO Standard Pricing Solution (Control).
- 50% of applications are routed through XYZ Pricing Solution (Treatment).
- XYZ prices are about 1% higher than BMO standard prices.

2009 was a very special year for the financial industry. Financial institutions observed a significant increase in losses due to 2008 Financial Crisis. Absolutely there was a need for smarter lending strategies. At the same time there is a significant competitive pressure in this particular sector of the lending business. Low risk customers can negotiate offers in several financial institutions at the same time. This part of the lending business is very appealing for development and testing analytical optimization theories and solutions.

The performance monitoring is set up by month/vintage since the offer is issued. We compare two groups of applications – control and treatment. Due to the random partition these groups are practically identical at the very beginning of the process. Some applications are rejected due to risk reasons. Risk rejected parts are about the same in every group. The rest applications get offers which include credit amounts and pricing rates. Some customers accept offers and some customers reject. In case if a customer accepts the offer we open or book a credit line account. All performance metrics are set to zero for non-booked customers.

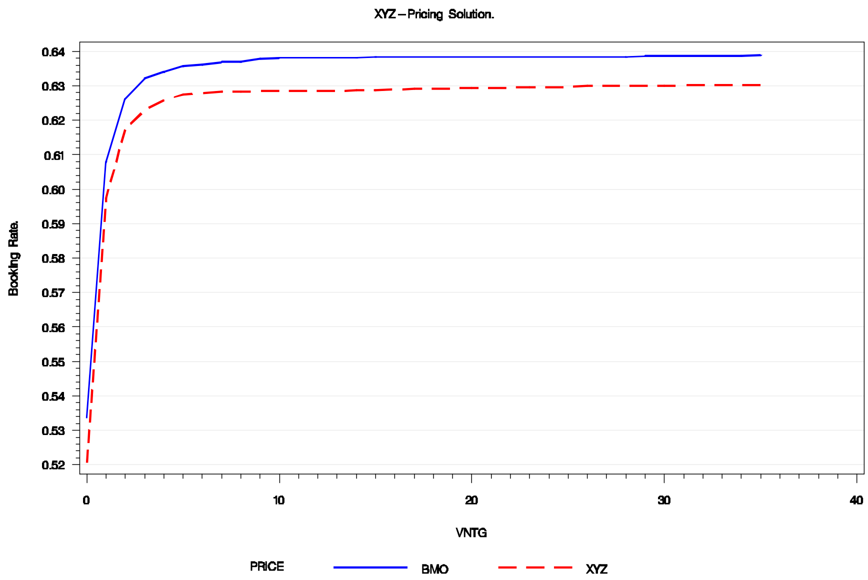

Figure 1. Booking Rates by Vintage: BMO vs XYZ-Pricing Solution.

Under higher XYZ-prices we booked by about 1% less accounts (a negative booking effect).

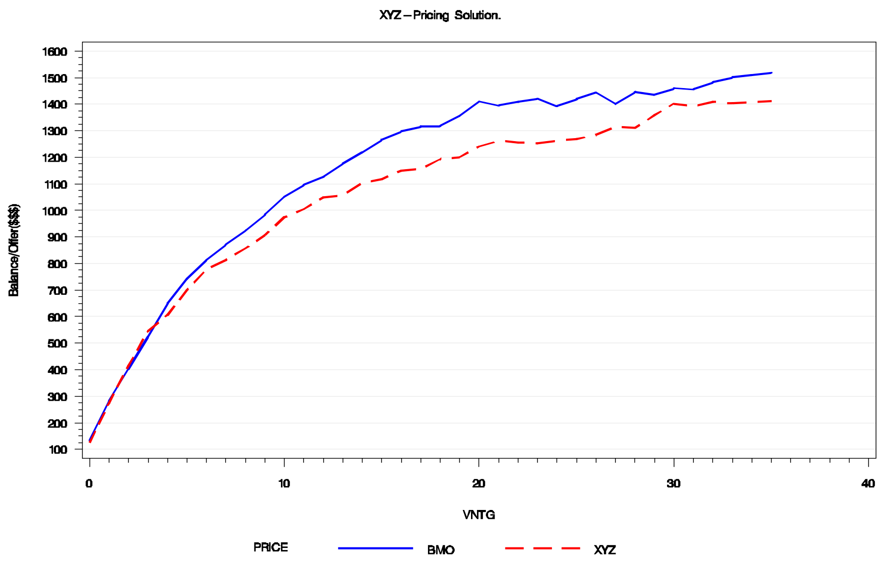

Figure 2. Balance/Offer: BMO vs XYZ-Pricing Solution.

Observe fewer balances for the XYZ treatment group (a negative balance effect).

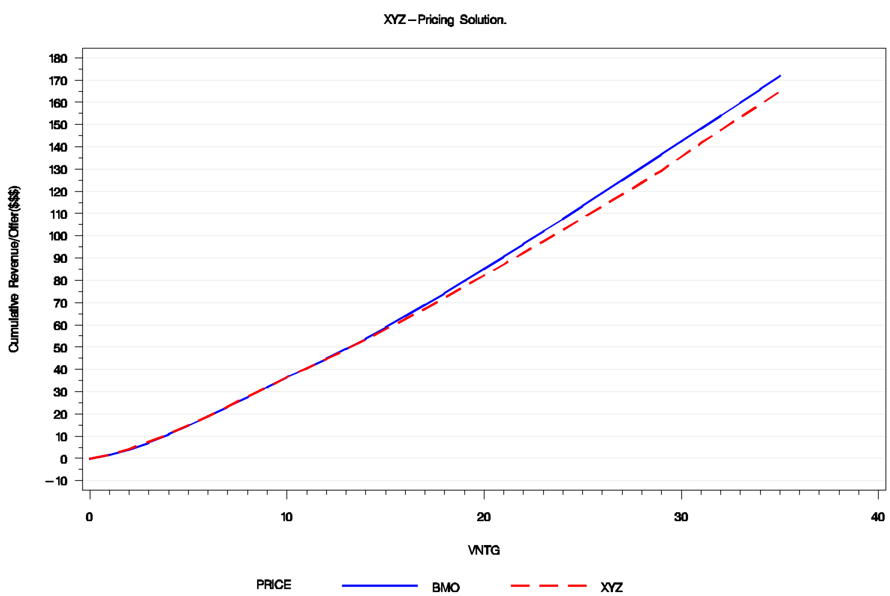

Figure 3. Cumulative Revenue/Offer: BMO vs XYZ-Pricing Solution.

Higher XYZ-prices did not compensate for fewer booked accounts and balances. The objective – to generate the additional revenue, is failed. Overall the higher in price XYZ-solution generates less in revenue on the segment of interest.

Drilling inside the test segment which we call the effect modeling, can help to identify sub-segments where higher or lower prices are more beneficial for revenue optimization objectives.

The effect modeling (uplift modelling by Radcliffe and Surry (1999, 2011), true lift by Lo (2002), net lift by Larsen (2010)…) is the main tool for our credit/pricing optimization methodology.

Wikipedia gives the following explanation: “Uplift modelling, also known as incremental modelling, true lift modelling, or net modelling is a predictive modelling technique that directly models the incremental impact of a treatment (such as a direct marketing action) on an individual's behavior. Uplift modelling has applications in customer relationship management for up-sell, cross-sell and retention modelling. It has also been applied to political election and personalized medicine."

Indeed numerous examples of effect modeling to optimize the impact of marketing campaigns with binary response targets can be found online or in publications. See Radcliffe and Surry (2011), Rzepakowski and Jaroszewicz (2012) for a history and literature review.

In this paper we explain applications of Effect (Uplift, Net…) modeling for credit limit and pricing optimization where target variables are mostly continuous – revenue, cost, losses and profit.

Eventually the developed methodology can be used to optimize any continuous parameter – fee, reward, or teaser rate....

II. TREATMENT EFFECTS, EFFECT MODELING, EFFECT/SENSITIVITY SCORES.

TREATMENT AND ITS EFFECTS.

Radcliffe and Surry (1999) defined the concept of incremental response for customers from a specially treated group (treatment group) vs customers from the control group. At the Bank of Montreal we traditionally refer to treatment's effects when talking about the treatment and how it affects customer's behavior. Also defining effects of a treatment over another treatment (vs over the control) is more comprehensive and beneficial for further discussions. The control is often associated with a "do-nothing-scenario" which is not always the case for us.

To define the effect we need at least two groups of customers associated with different treatments. Also we need a target function. The concept of effect is illustrated with the following example.

Credit Limit Increase Test Campaign.

- Portfolio: Unsecured Credit Lines (LC).
- Launching Time: April, 2010.
- Experimental Design: 60% customers were treated with a fixed credit limit increase and 40% customers were held as the control.

The performance monitoring is set up by month/vintage since the offer has been issued. We compare numerous performance/target metrics for two groups – treatment vs control. The following figure illustrates differences in performance for the balance target metric by month.

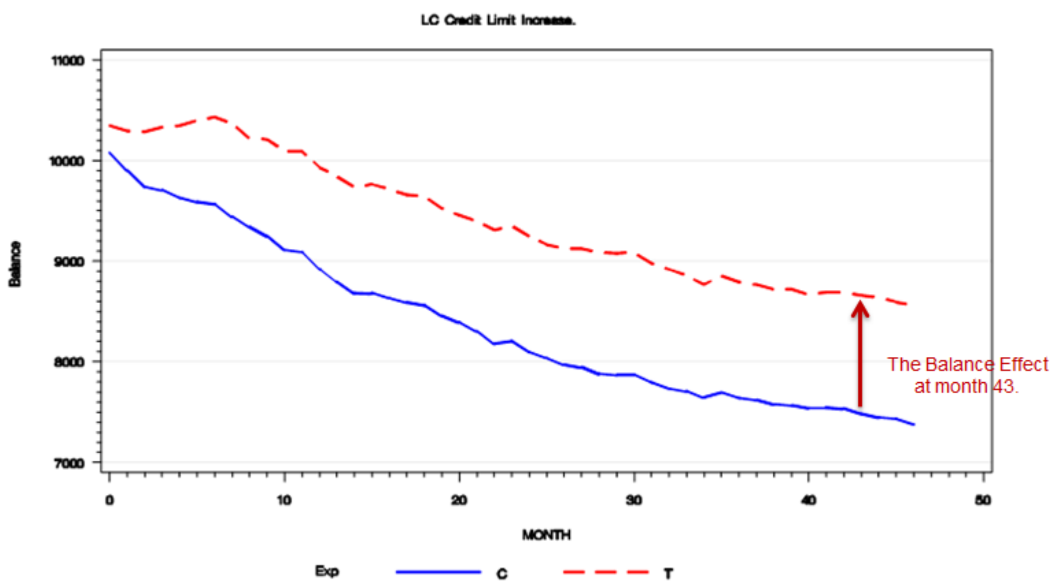

Figure 4. Balance/account by Treatment and Control Groups.

Both treatment and control customers are paying down their balances. Credit limit increase treated customers are paying their balances at a lower rate. Given a target metric, for example balance, we define the balance effect at a certain month, say month 43, as the difference between the average balance values on treatment and control groups.

The notion of effect is a segment level metric. It is not an account or customer level metric which makes the effect modeling very different from the conventional modeling. As mentioned before, the concept of the conventional responder is no help here. Everyone in the treatment group is a treatment responder.

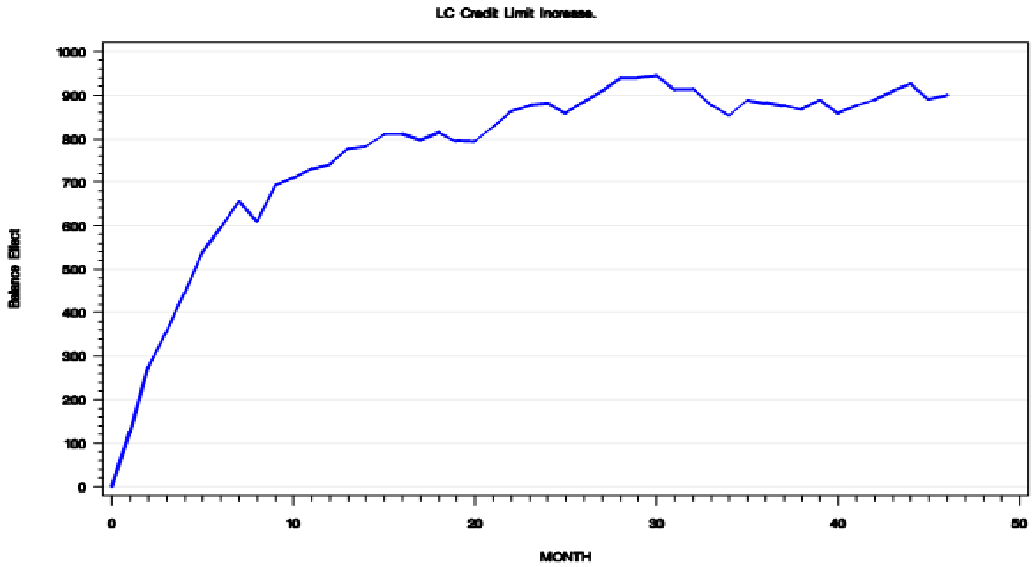

Figure 5. Credit Limit Increase Balance Effect by Month.

After 30 months treated customers revolve on average $900 more on their credit lines. The effect eventually levels off. We started testing effects practically since 2001 and observed that credit limit increase effects are very steady and can last for more than 10 years.

EFFECT MODELING.

Given a specific target function, for example balance, we want to drill down to the segment and find sub-segments that generate very significant effects and sub-segments that actually do not utilize the treatment (for example a credit limit increase) or even can be upset by the treatment. We do the drilling using the effect modeling.

In banking business the ultimate effect modeling target function is profit = revenue – cost - losses.

To reduce the overall noise we model sometimes profit components independently and sum up everything later.

Quite often cost effects for account management credit limit and interest rate treatments are insignificant and can be ignored. Bearing in mind the rare nature of loss events we would rather model loss effects separately.

Certainly the vital difference for the effect modeling from the conventional modeling is that the effect modeling is defined on the segment level. Nevertheless the conventional modeling suggests an instant and straight solution that we call sometimes the standard two-way effect model. Larsen (2010) calls it the difference score model (DSM).

Assume that we would like developing a balance effect model on the set of input predictive variables X. Indeed we can do it in two steps.

1. Develop a balance model on the treatment population: $B_T = B_T(X)$.
2. Develop the second balance model on the control population: $B_C = B_C(X)$.

The standard two-way balance effect model or the difference score model is just the difference of these two models:

$$Effect(X) = B_T(X) - B_C(X)$$

We wish having that straight and robust two-way effect modeling solution. Sadly the standard two-way effect modeling is not often the best way to go. With the effect modeling we run into significant challenges.

- High nonlinearity of effects.
- Quite often the effect represents only a small fraction of the target metric.
- At the same time since the effect is the difference between treatment and control values we run into a “double variance” problem which can cause unnecessary high model volatility.
- Business interpretation.

Our extensive study on effect modeling started back into 2001. Since that time we observe that the two-way effect models are particularly volatile. Volatility of two-way effect models is exceptionally high when we have to deal with continuous targets – profit, revenue, cost or losses. Business interpretation is another issue regarding two-way effect models. A thorough discussion on the two-way effect modeling failure can be found in Radcliffe and Surry (2011).

The effect is a segment concept. When we have the original experimental design that consists of treatment and control segments then it is evident how to measure and demonstrate a distinctive performance of the treatment group over control. Thus we came with the necessity of developing a decision tree effect modeling solution which has also a segment structure.

Decision Tree Effect Modeling SAS Macro.

In 2001 a decision tree based algorithm for the direct effect modeling was developed and implemented as a SAS macro.

Given a parent node, each time the direct effect modeling algorithm runs through the entire list of predictive variables and looks for a split value that maximizes the weighted power function which is the difference in effects between children nodes adjusted for the 95% confidence interval. A special weight function allocates higher values for splits that are closer to 50%. The confidence interval level can be adjusted as well as the weight function. The effect modeling SAS macro can be ran just with the option to maximize the difference in effects between children nodes. We found that in this case the effect modeling algorithm has a tendency chopping out nodes on edges. Edge nodes are becoming final nodes and the effect modeling tree is not “bushy” enough. To prevent it, we use weight functions that assign higher values for splits that are closer to 50%.

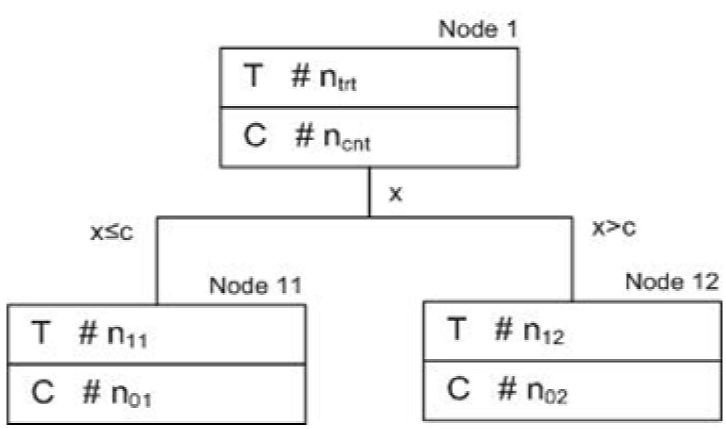

Figure 6. Splitting Node 1 into Node 11 and Node 12.

Here

- x is a splitting variable;
- n_{**} stands for the size of the corresponding treatment/control segment.

Node 11 and Node 12 effects can be calculated at the split point "*c*" in the following way.

$$Node_{11} : \Delta_1(c) = \mu_{trt}(x \le c) - \mu_{cnt}(x \le c)$$

$$Node_{12} : \Delta_2(c) = \mu_{trt}(x > c) - \mu_{cnt}(x > c)$$

The difference of effects at point "*c*" is

$$Diff(c) = \Delta_1(c) - \Delta_2(c)$$

The standard deviation for the difference can be estimated by

$$\sigma^2 = \frac{StdDev^2_{trt}(x \le c)}{n_{11}} + \frac{StdDev^2_{cnt}(x \le c)}{n_{01}} + \frac{StdDev^2_{trt}(x > c)}{n_{12}} + \frac{StdDev^2_{cnt}(x > c)}{n_{02}}$$

The 2σ-confidence interval (roughly 95%) for Diff(*c*) is

$$I_c = (Diff(c) - 2\sigma(c), Diff(c) + 2\sigma(c))$$

If 0 does not belong to the 2σ-confidence interval then we define

$$Power(c) = min(|Diff(c) - 2\sigma|, |Diff(c) + 2\sigma|)$$

Otherwise

$$Power(c) = 0$$

Given a weight value w > 0 we define the weight function W(x) as

$$W(x) = -\frac{(x - 0.5 + w)(x - 0.5 - w)}{w^2}$$

The weight is 1 when x=0.5.

The weighted power at the splitting point "*c*" is defined as

$$wPower(c) = W\left({}^{n_{11}}/_{(n_{11} + n_{12})}\right) \times Power(c)$$

Obviously the algorithm is very heavy computationally. In some our effect modeling solutions the number of input variables exceeds 10,000. In this case there is an option to run the effect modeling macro to select the list of important variables with respect to the weighted power first. The variable selection can reduce the number of input variable from 10,000 to several hundreds. After that we can train the detailed effect model on this list only.

The effect modeling macro can be used to train conventional decision tree models in SAS which we find very helpful. We just need to create a "fake" control group with the target function set to zero everywhere.

Starting with 2001 this methodology was used heavily for credit limit and pricing optimizations in the Bank of Montreal.

In their 1999 paper Radcliffe and Surry (1999) revealed a development of a direct effect/uplift modeling. We believe it was the first appearance of the direct effect modeling. Details on their approach can be found in their newest publication Radcliffe and Surry (2011). Lo (2002) is also one of the first authors on effect/true lift modeling. Larsen (2010) offers a SAS course on Net Lift Models: Optimizing the Impact of Your Marketing Efforts. Recently Rzepakowski and Jaroszewicz (2012) gave their version of decision tree uplift modeling which is based on generalization of classical tree split criteria.

So far all published examples on effect/uplift modeling applications in financial industry are on optimization of marketing response, sales and retention efforts. In all these cases target variables are binary. We use the effect modeling to optimize credit and pricing strategies where targets are continuous.

Effect or Sensitivity Scores.

Price or credit sensitivity scores are available on the market from different analytical vendors. We define always an effect or sensitivity score through the effect modeling lense. It means that there is a treatment and a specific tested population with an appropriate experimental design. As a result we can use a proper effect modeling technique and build an effect or sensitivity score.

It can be any treatment effect/sensitivity score with respect to any target function – balance, revenue, cost, or losses. For credit limit or pricing optimizations we develop and use effect/sensitivity scores with respect to revenue, cost, losses and profit targets. Credit limit increase and decrease effect/sensitivity scores as well as pricing sensitivity scores are critical for optimization credit strategies for lending portfolios. Bank of Montreal analytical culture requires a proper experimental design to measure an action and even more appropriate design to develop effect/sensitivity scores. There is a test or even a series of test campaigns behind every analytically developed sensitivity score.

Sophisticated sensitivity scores can be developed for combinations of treatments. To compete for low-mid risk customer balances we develop sensitivity scores for combinations of credit limit increase and interest rate decrease treatments. The lending competitive pressure is exceptionally strong for this segment. The experimental design is quite special if we want developing sensitivity scores for treatment combinations. If we work on optimizing higher risk customers combinations of credit decrease and interest rate increase can be more appropriate.

Lending business is about an optimal management of customer balances. While the effect score is a general term we prefer using sensitivity score name when the target is balance. In this case the effect score is described in terms of balance effects. With the balance as a target the sensitivity score gives us specific amounts in balance effects that the treatment stimulated or caused to balance attrition.

In the next section we give examples of effect/sensitivity scores.

EFFECT SCORE EXAMPLES.

Credit Limit Increase Effect Score.

Credit Limit Increase (CLI) Test Campaign.

- Portfolio: Unsecured Credit Lines (LC).
- Launching Time: January, 2011.
- Experimental Design: 50% customers were treated with a fixed credit limit increase and 50% customers were held as the control.

We train the direct effect model using the decision tree effect modeling SAS macro with the cumulative revenue as the target. The final credit limit increase effect score (after pruning and cross-validation) contains 17 effect nodes.

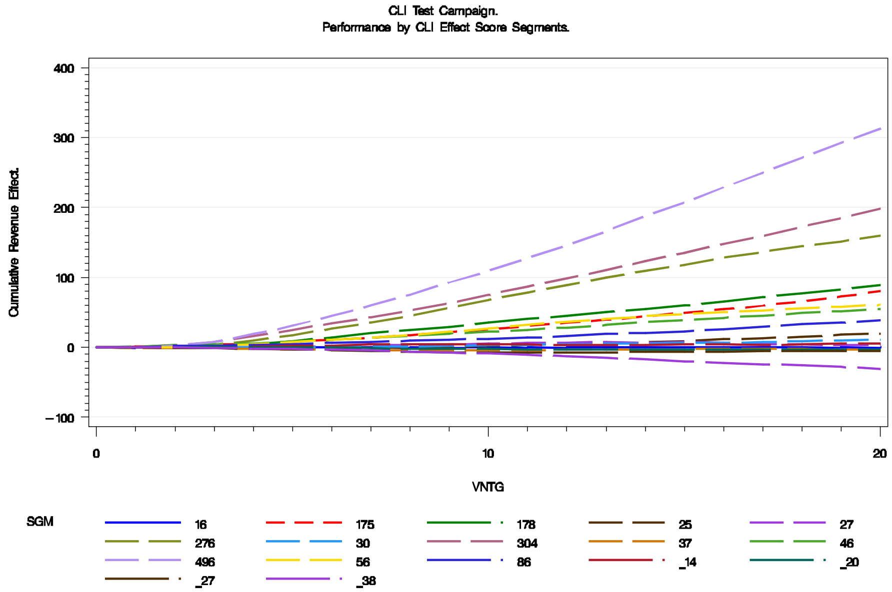

Figure 7. Cumulative Revenue Effect: Performance by CLI Effect Score Segments.

The effect modeling identifies heavy CLI users over the control as well as sub-segments/customers that hardly or do not use additional CLI at all. Customers from CLI sensitivity SGM=_38 are even upset with the CLI treatment and show a balance attrition with a small negative revenue effect. It does not mean that heavy CLI users must grow their total balances. They grow balances against the control. As we saw on Figure 4 both treatment and control customers could be paying down their balances. CLI users can be paying their balances at lower rates.

Brining losses and cost effect curves will complete the whole profitability picture. We'll get closer to optimization of the CLI treatment on this segment. The optimization can be done by revenue, losses and cost effect/sensitivity segments.

Interest Rate Increase Effect Score.

Interest Rate Increase (IRI) Test Campaign.

- Portfolio: Unsecured Credit Lines (LC).
- Treatment: 100 bps Interest Rate Increase (IRI).
- Experimental Design: 50% customers were treated with IRI and 50% customers were held as the control.

We train the direct effect model using the decision tree effect modeling SAS macro with the cumulative revenue as the target. The sample size of this test campaign was relatively small. Therefore the final interest rate increase effect/price sensitivity score has only four effect nodes.

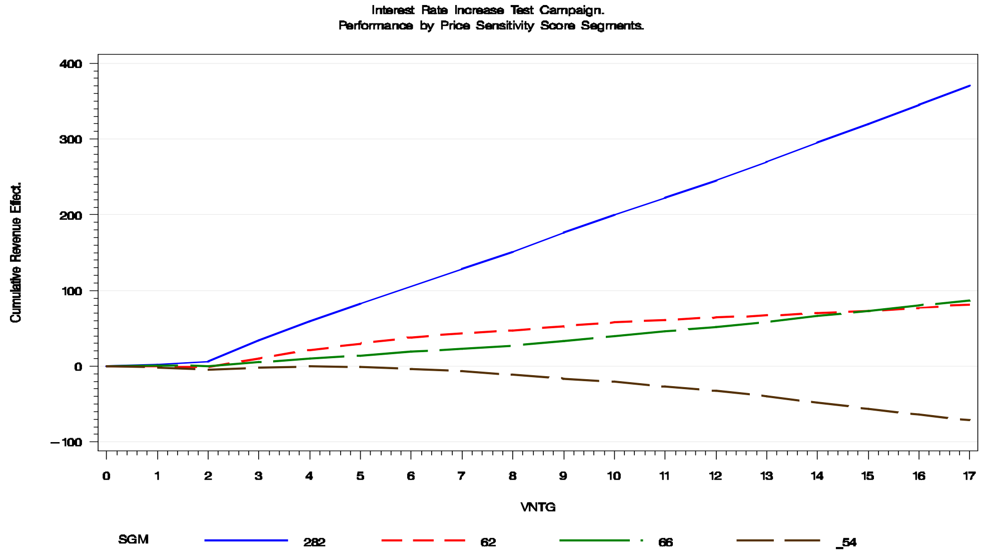

Figure 8. Cumulative Revenue Effect: Performance by IRI Effect Score Segments.

The effect modeling identifies a very price sensitive segment: SGM=_54. The segment revenue is affected negatively shortly after the interest rate increase treatment. On the next Figure 9 we see that SGM=_54 exhibits the highest balance attrition (vs the corresponding control). Increased rates can't offset negative effects on balances and generate incremental revenue. Definitely this kind of customer should be excluded from price increase campaigns.

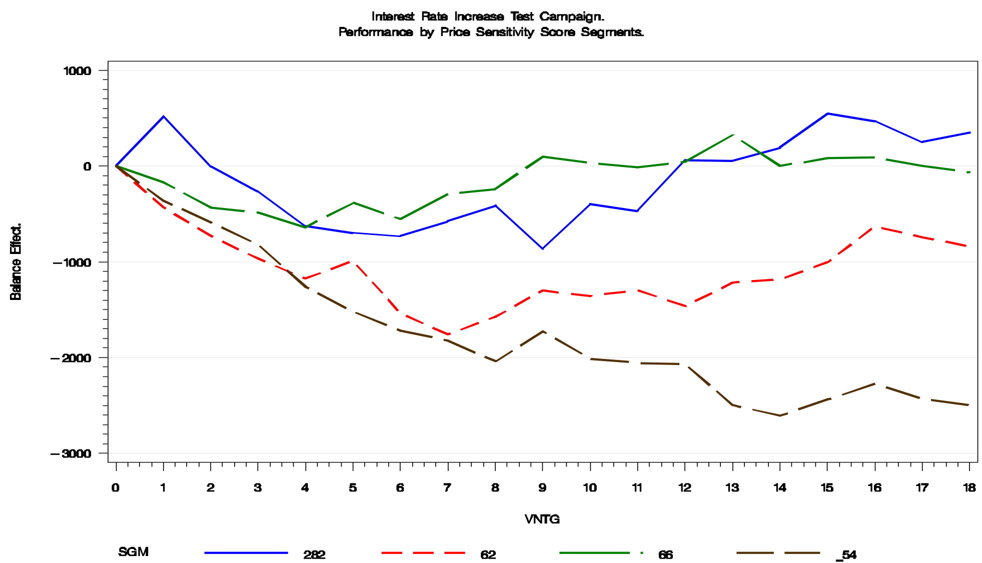

Figure 9. Balance Effect: Performance by IRI Effect Score Segments.

Customers from effect/price sensitivity segments 282 and 66 can take IRI stress. After 12 months balance attrition effects for these segments has practically vanished. These customers are suitable candidates for safe IRI treatments.

More analytics is needed to decision on customers from SGM=62. This will be addressed in following sections.

III. EFFECT LIFT/GAIN CHARTS, COMPARE/ASSESS EFFECT MODELS.

EFFECT LIFT/GAIN CHARTS.

We use the previous example on IRI effect score to illustrate the concept of effect lift or gain charts. IRI effect/price sensitivity score has four final effect nodes: 282, 66, 62 and -54. To build the cumulative lift chart we sort the population from highest-to-lowest scores. The x-axis shows the percentage of population with the highest scores. Given the percentage x, the y-axis shows the total effect on the x-percentage population. If instead of the total effect we put the average effect on the y-axis we generate the regular effect chart.

The same interest rate increase campaign was used also to validate a price sensitivity score that was developed by an external vendor. Let's call it the XYZ-price sensitivity score developed by the XYZ-team.

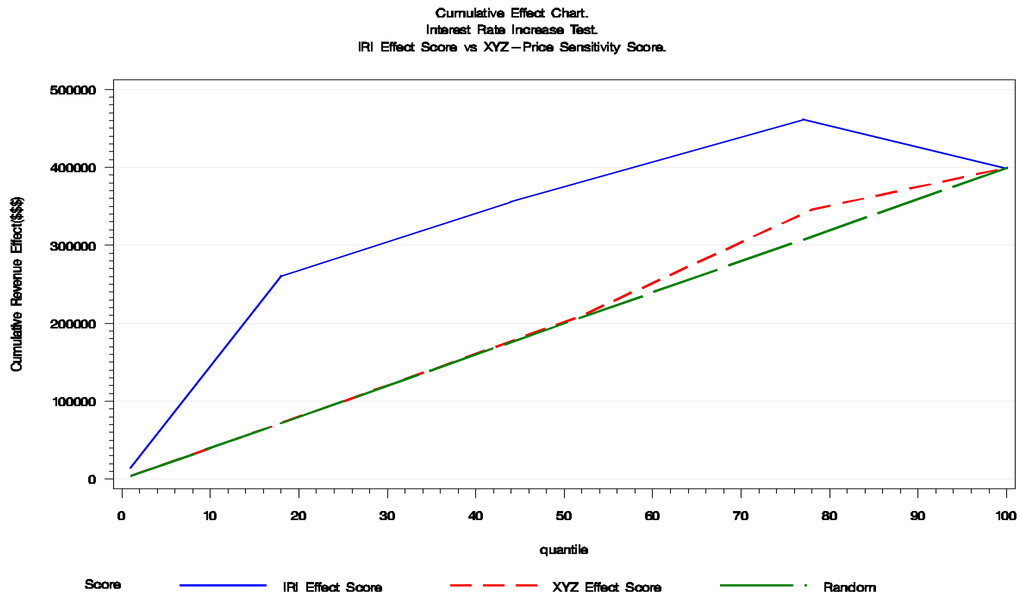

Figure 10. Cumulative Effect Chart: IRI vs XYZ scores.

IRI effect score was trained using the decision tree effect modeling SAS macro with the cumulative revenue as the target. We do not have details on XYZ-score development.

Overall for the observation period the treatment population generated $0.4MM in cumulative revenue effect (the incremental revenue over the control) due to the 100 bps rate increase. About 77% of the population, where IRI effect score>0, generated the maximal cumulative revenue effect of $0.47MM. Definitely by targeting population from 77% - 100% we have damaged the cumulative revenue effect and have provided a suboptimal decision.

COMPARE/ASSESS EFFECT MODELS.

To compare the two effect models we need to introduce a metric or even a series of metrics that can estimate a model's power. Since effect modeling is a segment level concept, assessing the effect model performance is more complex. See Radcliffe and Surry (2011) for an additional discussion on uplift/effect quality measures.

Radcliffe (2007) proposed the Qini coefficient (an analogue of the Gini coefficient) to measure the Effect Model Power. The Qini coefficient is a "normalized" area between the cumulative curve and the diagonal. Unlike for the Gini coefficient when we can think of the perfect response model a notion of the perfect effect model is not clear at all. We can come up with theoretical examples when the perfect effect model eventually does not exist.

Larsen (2010) compares effect models using their performance in the top decile/quantile.

Our definition of the effect model power comes from business needs. It is practical for business to associate the power of the effect model/score with the maximum value on the cumulative effect chart. Sometimes it is convenient to normalize cumulative effect values dividing by the number of treated customers/accounts. In this case we can compare cumulative lift charts for different samples.

We have to be cautious if the maximum cumulative effect value is attained at 100%. Quite often in these situations, the developed effect score does not create any business benefit.

The effect model power definition is directly connected with business objectives and therefore is preferable for our needs. Refer to Figure 10 for the effect model power definition. For the observation period the whole treatment population generated $0.4MM in cumulative revenue effect. Therefore $0.4MM can be thought as the effect power of the random model. The maximal cumulative revenue effect value for IRI-effect score is $0.47MM. This maximum value attains at 77%. Targeting 77% of the best population would increase the incremental revenue by 17%. One can suggest that the ratio of the developed effect model power over the random model power can be a good candidate for the effect model assessment. As we'll see in the next example that the random model cumulative effect (the total effect on the treated population) can be close to zero or even negative. Therefore we find that using the ratio of the model effect power over the random model effect power is not that helpful to measure the effect model success.

The XYZ-score performance is practically close to random up to about 50%. Even it is slightly better than the random from 50% - 100%, the XYZ-score attains its maximum at 100%. Given that using XYZ-score to cut a more profitable than random population is not feasible.

So far, the cumulative lift chart is the main tool for us to measure the effect model quality. The cumulative effect chart for the proper developed effect model is always above the random line. Higher the cumulative effect "bump" tells about a better effect model.

It is not uncommon when optimizing a pricing or credit management strategy that the very first test campaign shows about zero incremental results. The subsequent effect modeling can identify a subpopulation where the tested treatment generates the maximum incremental value which is typically positive. Also, the effect modeling identifies segments with about zero and very negative total cumulative effects. For these segments, treatment adjustments in combination with selection criteria adjustments need to be done. Based on optimization results of the first test campaign we design the second phase test which is likely to demonstrate significantly better overall incremental results. The process of approaching the perfect optimal solution is endless. After each optimization phase we improve the quality and competitiveness of the business significantly.

Creating an analytically driven business culture cannot happen overnight. It takes commitment, persistence, and resilience to succeed. Financial institutions that maintain a 'one size fits all' approach to customer management will find it increasingly challenging as we evolve to a more customer level approach to account management.

We will finish this section by exploring more cumulative effect charts based on the examples discussed in sections I and II.

XYZ-PRICING SOLUTION TEST CAMPAIGN.

This test campaign was discussed in Section I. Figure 3 shows that the higher in price XYZ-solution generates less in revenue on the segment of interest. The effect modeling can fix the negative cumulative revenue effect damage and help optimizing the pricing solution on the segment of interest. We train the direct effect model using the decision tree effect modeling SAS macro with the cumulative revenue as the target.

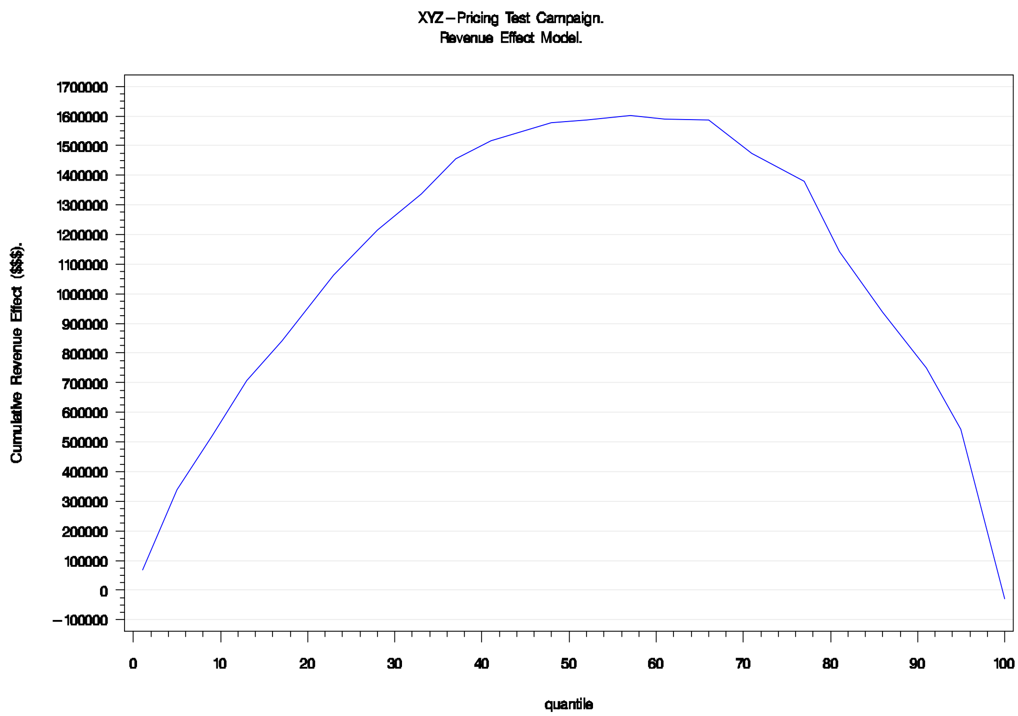

Figure 11. Cumulative Revenue Effect Chart: Revenue Effect Score for XYZ-Pricing Campaign.

Overall the cumulative revenue effect is slightly negative. Targeting 57% of the population with higher XYZ-Prices and keeping the rest 43% at lower standard prices would generate additional $1.6MM in revenue effect. In this case the optimization is very significant – turning a negative test into a positive business opportunity.

XYZ as well as other standard financial industry price optimization solutions are generally simplistic and based mostly on risk and limit dimensions. To win the competitive game we need justified complexity and sophistication in pricing and credit optimization solutions. The decision tree effect modeling algorithm provides such complexity. For this particular pricing effect model among strongest input variables are: income, requested amount, marital status, number of months in the last residency, customer's net worth, inquiry data, competitive data on revolving loans, a bureau score and more.

CREDIT LIMIT INCREASE (CLI) TEST CAMPAIGN.

This test campaign was discussed in Section II. CLI effect model was developed using the decision tree effect modeling SAS macro with the cumulative revenue as the target. CLI effect score contains 17 final effect nodes. Figure 7 presents cumulative revenue effect performance of all CLI effect nodes by month or vintage.

Figure 12 shows the cumulative revenue effect chart of the CLI effect score on the independent validation data. CLI is a revenue, balance as well as losses generating treatment. The shape of the CLI cumulative revenue effect chart presented on this figure is very typical.

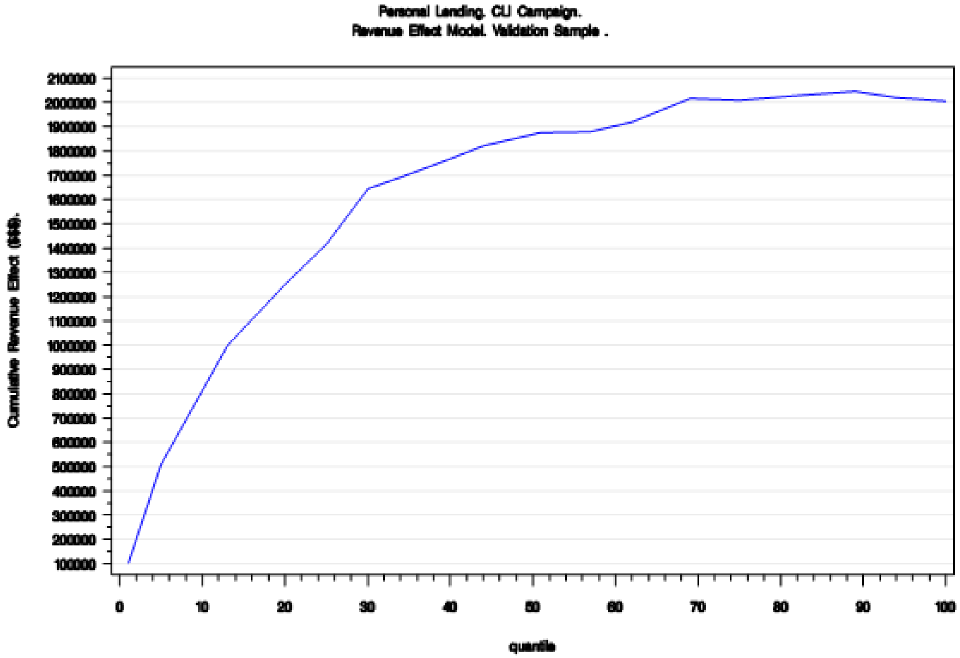

Figure 12. Cumulative Revenue Effect Chart: Revenue Effect Score for CLI Campaign.

The effect modeling identified customers that got even disappointed with CLI treatments. Customers from 89% - 100% show minor signs of negative revenue effects. Customers from 69% - 100% practically do not utilize the given CLI treatment. Customers from 51% - 69% generate low revenue effects. These low positive revenue effects won't offset negative effects in losses. Eventually the optimization step, which brings losses in to the game, removes 51% - 100% customers from CLI treatments as suboptimal.

IV. EFFECT MODELING AND CREDIT/PRICE OPTIMIZATION.

Effect modeling is an important step in the optimization process. However it is an intermediate step. In order to make a decision on a particular credit/pricing treatment we need to know lifetime effect projections on profit and its major components – revenue, cost, losses and more. Given business constraints and objectives on the portfolio health – loss-to-balance, revenue-to-loss, profit-to-capital ratios, we can identify an optimal population for this particular treatment and set up further testing campaigns. Quite often the performance time for the effect modeling is limited from 12 to 24 months. We can use the simulation box application to generate necessary extended projections.

The simulation box is a SAS based application/macro. Given a set of customers/accounts the simulation box generates all necessary projections (up to 10 years) based on customers' effect scores and input effect curves for main business metrics – revenue, cost and loss components.

THE SIMULATION BOX.

Building the simulation box is a systematic process and requires establishing various standards around treatment definitions and effect measurements. We just need to fill a SAS macro-template with specific input data. Definitely we need to input preliminary developed revenue, cost and losses effect scores. For every effect score segment we have to input corresponding effect curves – similar that we saw on Figure 7 or Figure 8. Effect curves for profit components are limited by the model observation period which is often from 12 to 24 months. To build up to 10-year projections we use findings on the long time performance of balance effect curves.

BALANCE EFFECT FUNDAMENTAL SHAPE.

Lending business is about managing balances. To price these balances healthy and competitively is the main optimization objective.

Three Balance Effect Phases.

Practical business sense suggests that after applying a credit/pricing treatment we can observe a balance effect building period for a certain time (phase 1), followed by leveling-off (phase 2) and a natural slow decline (phase 3). Since 2001 we accumulated hundreds examples on balance effect curve shapes.

Depending on a targeted population and a treatment, phase 1 (the balance effect building period) can last from 12 – 36 months. Phase 2 (the leveling-off period) can continue even beyond 10 years. Phase 3 (the natural slow decline) is less analyzed. Phase 3 requires the performance monitoring beyond 5 – 10 years. It's always a challenge to hold experimental cells untouched for longer than 5 – 10 years. Our current optimization solutions are mostly based on up to 10 years financial projections. Phase 1 and 2 trajectories are main components to build these projections. Phase 1 (the balance effect building period from 12 – 36 months) is the most critical to capture changes (effects) in customer's behavior. Following examples illustrate these ideas.

We start with Figure 5. Credit Limit Increase Balance Effect by Month. This CLI test was discussed in Section II. Observe the balance effect building phase 1 lasts up until 30 months followed by the leveling-off phase 2. After almost 4 years we do not observe signs of effect declining phase 3.

The next example is on a CLI test for cards portfolio.

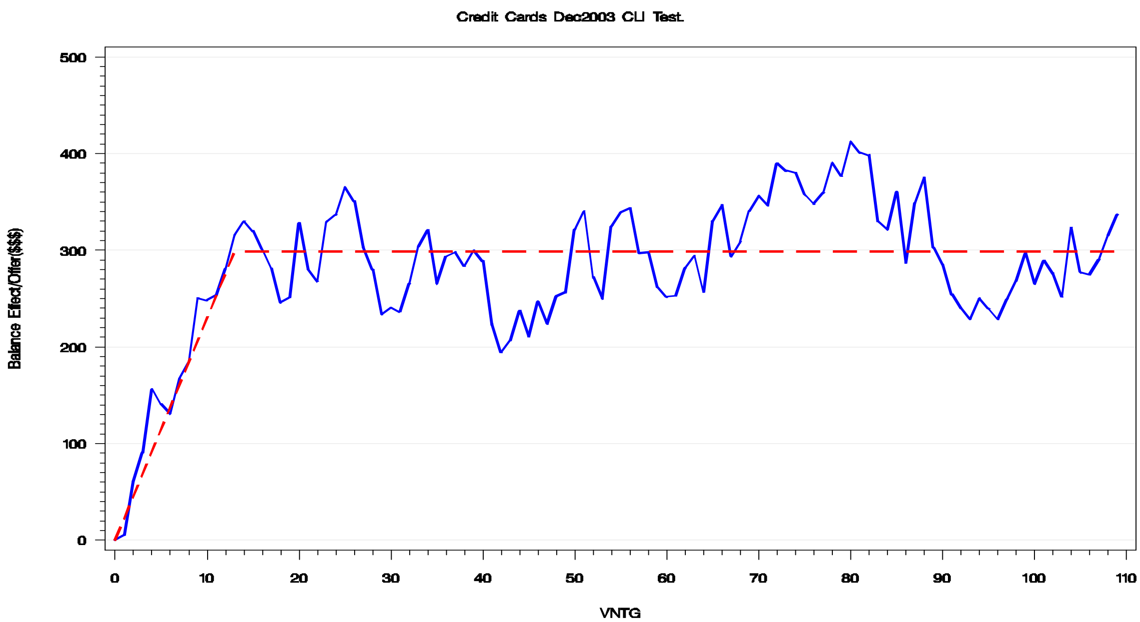

Figure 13. Credit Cards CLI Test: Balance Effect by Month.

The campaign was launched in Dec2013 on a significant part of cards portfolio. We stopped monitoring the campaign after almost ten years of performance. Balance effects are building during the first 12 months.

Even we observe some dips and upturns after 12 months overall it can be classified as a levelling-off behavior for about ten years. The first "dip" continued from Jul2007 – Dec2007. This was a pre-2008 Financial Crisis time when to get a card loan was relatively easy. Therefore the competitive pressure was significant. The "rise" is associated with 2009 – 2011 – after the Financial Crisis. Banks put conservative rules on loan's adjudication. CLI treated customers exhibited better responses and as a result higher effects from 2009-2011.

Figure 14. Loan Origination Pricing Test: Balance Effect/Booked Account.

Balance Effects are building until about 24 months and leveling off after that.

CONCLUSION

ESSENTIAL CREDIT/PRICE OPTIMIZATION STEPS USING THE EFFECT METHODOLOGY.

Essential Prerequisite: Superior Risk Segmentation.

This is how banks are competing. On average every adult Canadian has about three credit cards. Given a fitting credit history a customer can open a personal line of credit practically in any Canadian bank. The size of the credit and its rate depend heavily on customer's risk assessment which is done by the bank. About 50% of Canadian customers are of very low risk and can get a credit practically at any bank. Coincidentally these customers are not looking for additional credit. Banks are playing the competitive game on the other 50% of the population. The opportunity arises when the bank can identify low-mid risk customers that other banks classify into high risk. Obviously a superior risk prediction score is always important for building powerful loss effect scores.

Step 1. Define High Level Portfolio Segments of Interest and the List of Credit/Price Incrementals to test.

Many years ago the author used to believe that if we have all necessary data elements on a customer then we can decide on the best credit size and the price for the customer. It turned out that this is hard-to-unrealistic problem to solve. Instead author's choice is to go with the best guess first and use 2-4 corrections to attain the best solution for the customer. It can take a few years. Therefore it is important to set up the test road map.

Step 2. Set up and launch 3 – 5 year Testing Plan.

Step 3. Set up the Effect Performance Monitoring and Backing (Reversing) Process.

We need to monitor effects and also run a light effect modeling measurement to identify sub-segments that exhibit extreme negative changes (effects) due to the treatment. If such segments are detected we want to act immediately reversing the treatment. SGM=_54 from Figure 8 and Figure 9 is an example of such negatively performing segment. Practically after less than four months we could identify that this segment went to the negative territory and needs a reverse/remedy treatment.

Step 4. Effect Modeling and moving to the Optimization.

Depending on how much knowledge we've accumulated on behavior and fundamental shapes of balance effect curves we can start analyzing effects as earlier as after 6 months of performance. The standard time before credit/price effect model developments can be from 12 – 24 months. Profit is the ultimate effect modeling target. It can help partitioning profit into several components and model them separately. Play and transform targets to reduce their variance. Transforming targets and particularly loss targets is a quite interesting analytical area. Analysis of outliers helps reducing the impact of the "double variance" and improves the effect model stability. "High Nonlinearity" of effects makes the variable reduction into another analytical challenge. Our direct effect modeling solution is the decision tree algorithm and can help defining the most important effect variables efficiently.

When the effect modeling segments are developed we can act on some "low-hanging-fruits" immediately.

In Section II we discussed the interest rate increase test. After 18 months of performance some segments showed almost zero balance attrition. Our knowledge on shapes of balance effect curves suggests that these segments can take IRI stress. Customers of this type are suitable for similar IRI treatments. On the other hand the effect modeling identifies a segment that moves to the negative revenue effect territory almost immediately. Customers from this segment need reverse/remedy treatments as soon as we identify the negative effect. For the rest segments the following step is essential.

Step 5. The Simulation Box - Building the long time projections.

The business decision on applying the treatment has to be made on the long-time performance given business objectives and the risk appetite. As soon as we got the new and better optimized action plan we can move to step 1 and close the loop that supposed to bring us to the optimal solution eventually.

REFERENCES

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

I wish to personally thank the following people for the contributions in creating this discussion paper.

Allan Esser, Professor, George Brown College, and former colleague from Bank of Montreal.

Alexei Marcilio, Graduate Student, George Brown College, Analytics for Business Decision Making.

Alex Martkovich, Graduate Student, George Brown College, Analytics for Business Decision Making.

George Brown College is one of Canada's first colleges to offer a fully on-line program in business analytics that prepares students for turning information into knowledge, and converting intelligence into action. For more information about their unique program, go to:

http://www.georgebrown.ca/uploadedfiles/ProgramPDFs/2014-2015/Business/B412-2014-2015.pdf

For further information on joint industry research collaboration opportunities, please contact: aesser@georgebrown.ca

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