

NEXT BEST OFFER CAMPAIGNS

Lessons learned from a decade
of marketing optimization

by Robin Way



CORIOS REDPAPER — Papers that get read.™



CORIOS REDPAPER — Papers that get read.™

NEXT BEST OFFER CAMPAIGNS

Lessons learned from a decade of
marketing optimization

by Robin Way



Key findings

1. The fiscal contribution of cross-sell marketing campaigns can be significantly enhanced by adopting mathematical optimization approaches that augment traditional campaign strategies.
2. Like any complex analytical strategy, adoption of mathematical optimization for marketing campaigns requires skills, patience and diligence, which is usually well worth the return on investment.
3. Based on dozens of real-world projects, the experience of other marketing organizations in using mathematical optimization to build next-best-offer cross-sell strategies will help your organization reach similar objectives faster and with more organizational buy-in.

Introduction

There are many campaign designs that will produce strong results, which are produced by the ideal mix of offer strategy, timing, incentives, delivery, controls, fulfillment, brand, and of course, previously unmet customer needs. The pinnacle of cross-sell campaign performance is reached by having the ideal allocation of next best offers ("NBOs") to customers via a mathematical optimization-enabled campaign discipline, practices and tools.

In my career, I've been fortunate to support some of North America's leading companies in the financial services industry, including credit card issuers, retail banks, insurance carriers, payments networks and brokerage firms. What I've learned after a decade of intensive focus on customer strategy execution in this context is that the most exciting discoveries about ideal cross sell are made when using a practice that builds on a strong predictive scoring foundation and leverages mathematical optimization to allocate the best offers to customers. That statement has also proven accurate for customer strategies other than cross-sell marketing, but that one particular strategy is where I've built the greatest concentration of evidence. Few other approaches allow the practitioner to incorporate the policies and preferences of a diverse set of constituencies in the firm, and blend that point of view with the implied preferences of customers, as captured by predictive scores representing those preferences and needs. In fact it's the very presence of economic trade-offs, and the insights gained from inspecting them, that makes a mathematically optimal NBO strategy so exciting.

Earlier in my career I was employed at SAS Institute, who has created one of the very best mathematical optimization engines for enhancing the performance of campaign strategies. Back in the early 2000s, along with my colleague Jon Orme, who at the time was employed by SAS in Australia, we presented a business case to SAS' senior management to build a solution called SAS Marketing Optimization. Today that solution is recognized by technology industry analysts as a leader in the marketing discipline. The credit for that achievement goes to the R&D team who designed the solution. My own professional experience with optimal campaigns begins in the trenches, having implemented that solution in over 30 engagements with financial services companies in the US and Canada.

Practitioners from a diverse set of industries focus on cross-sell campaigns, but it is fair to claim that the financial services industry assumed a leading role. This is due to the relative maturity of decision sciences, offer targeting and performance measurement capabilities that have been developed in the financial services industry for decades now. That said, other industries bring unique perspectives to cross-sell, such as the focus on social marketing in

telecommunications, the focus on a vast breadth of products in the retailing industry, and the focus on customer delight in the entertainment industry. Mathematical optimization of NBO campaigns has the potential to work effectively in these settings and more.

This paper will review how mathematical optimization of cross-sell campaigns add value over and above other traditional offer allocation techniques; the value of information in balancing short term and medium term campaign performance; the challenges that campaign analysts will face when implementing a mathematical optimization capability for the first time; and some lessons learned from applying mathematical optimization to cross sell campaigns. Peppered throughout this paper are case studies that illustrate real-world scenarios of implementing optimized NBO campaigns and some of the unique aspects of those initiatives.

How mathematical optimization adds value to campaign performance

My colleague Brent Lever at SAS created a simple and very powerful numerical example of how a mathematical optimization approach to campaign selection out-performs traditional approaches.

Consider the situation depicted in Figure 1 with nine customers and three campaigns; each campaign is for a different product, and each customer-campaign assignment has an expected value to the firm, which is traditionally measured as a combination of an expected response score and an expected profit given response estimate. For instance, for customer 1 and campaign A, if the expected response is 2% and the expected profit given response is \$5000, then the joint expected value is $2\% * \$5000 = \100 . Also take into account for this simple example, some contact policies: each customer must receive one campaign offer, and no product can receive more than three campaign offers. The goal is to maximize overall expected profitability.

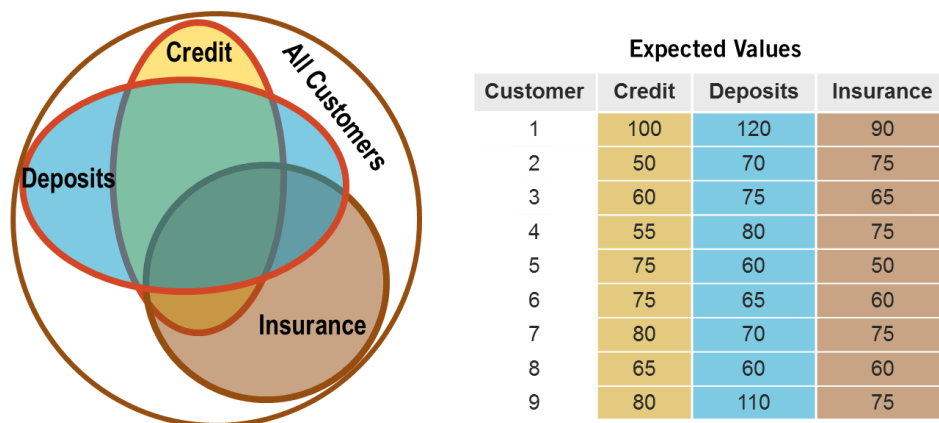


Figure 1: Setting up the mathematical optimization thought experiment

There are three traditional strategies of assigning campaigns to customers: campaign prioritization, customer prioritization, and mathematical optimization. With the campaign prioritization strategy, there is a hierarchy of offers whereby Campaign A receives top priority, then Campaign B, followed by Campaign C. Following this heuristic, Campaign A is assigned the three customers whose expected profit scores are the highest for that campaign, and the remaining six customers are then made eligible for selection by Campaign B. In turn, Campaign B is assigned the three customers whose expected profit scores are the highest from the six remaining, and Campaign C is assigned the remaining three customers. By examining Figure 2 below, you can see how this leads to some less-than-optimal assignments. It would have been smarter to assign Campaign B (rather than Campaign A) to Customer 1, in order to gain an additional \$20 of expected profit; but that campaign comes in second priority to Campaign A, so that option was not available to Campaign B for selection. The sum of expected profits for offers assigned to customers following the campaign prioritization strategy is \$655.

Customer	Credit	Deposits	Insurance
1	100	120	90
2	50	70	75
3	60	75	65
4	55	80	75
5	75	60	50
6	75	65	60
7	80	70	75
8	65	60	60
9	80	110	75

Figure 2: Campaign prioritization approach results in expected profit of \$655

The second strategy is customer prioritization, where the customers are prioritized into an expected profit order instead of campaigns. With this strategy, it's helpful as a thought experiment to imagine a flow of customers walking in a door, and then receiving their offers. Customer 1 walks in the door and is assigned the campaign with the highest expected profit for that customer, which is Campaign B. In turn, Customer 2 is assigned Campaign C, and Customers 3 and 4 are assigned Campaign B. At this point, all the offers for Campaign B have now been assigned, and all future customers to walk in the door will have to choose between Campaigns A and C. Again, this leads to less-than-optimal assignments, since it would have been smarter to assign Campaign B to Customer 9 relative to the previously-assigned Customers 3 and 4, but that choice was not available to the firm when those choices were rendered. In Figure 3, we observe that the sum of expected profits for the customer prioritization strategy is \$715.

Customer	Credit	Deposits	Insurance
1	100	120	90
2	50	70	75
3	60	75	65
4	55	80	75
5	75	60	50
6	75	65	60
7	80	70	75
8	65	60	60
9	80	110	75

Figure 3: Customer prioritization approach results in expected profit of \$715

The customer prioritization strategy is better than for the campaign prioritization strategy, and one reason is the granularity of available choices to the decision maker. When shifting from less granular to more granular choice strategies, the decision maker has more options from which to choose, especially if those choices can be made simultaneously and not sequentially. If a strategy to make a decision from among the most granular view of all choices simultaneously was available to the decision-maker, this would lead to the most profitable outcome. This is precisely the strategy provided by our third option, the mathematical optimization approach.

Using the mathematical optimization approach, the decision maker can observe all 27 choices of offers to customers simultaneously, and make the ideal allocation. Figure 3 shows these results with an expected profit of \$745, while also meeting all business rules or decision constraints. While this approach can be solved in a spreadsheet for only nine customers and three campaigns, consider the number-crunching required for solving this problem for millions of customers and tens of campaigns. Add to that challenge the introduction of more constraints, not only those that affect the top-down allocation of resources like budgets and capacity, but also bottom-up constraints, such as a contact policy invoked at the customer level over time (e.g., only contact each customer once every other month, and never repeat the same campaign in subsequent offers).

Customer	Credit	Deposits	Insurance
1	100	120	90
2	50	70	75
3	60	75	65
4	55	80	75
5	75	60	50
6	75	65	60
7	80	70	75
8	65	60	60
9	80	110	75

Figure 4: Mathematical optimization approach results in profit of \$745

Value of new information to balance short-term and long-term marketing decisions

In addition to maximizing expected profit from cross-sell offers, there are additional benefits from building mathematical optimization marketing campaign strategies. The decision-maker also benefits from understanding the contributions by resource that make cross-sell campaigns possible. These resources include outbound direct contacts via direct mail, call center, branches and digital channels, the corresponding incoming response handling, any customer incentives tied to an offer call-to-action, the budgets that fund these contacts, and the offer cells that differentiate specific offers from each other. .

Every one of these resources is finite, has different costs per unit, and get consumed by varying types of cross-sell offers at different rates. Therefore, the decision-maker has to make trade-offs in allocating resources to campaigns. The mathematical optimization engine takes all this information directly into account, and generates a series of metrics to characterize the importance of each resource and the point at which a resource gets fully consumed during the allocation of offers. At this point, there is a direct marginal cost to the campaign for consuming more of that resource, which is expressed in the same hard dollar terms as the goal to be maximized.

In the next chart, the percent change in expected profit is plotted against the corresponding percent change in two resources that are used to contact customers in a real-world cross sell campaign. Because these metrics are calculated as relative changes in profit for investment, they can be directly interpreted as elasticity metrics. The first resource is the sensitivity of

expected profit for an increase in the overall campaign budget (displayed as the blue line). The elasticity of profit with respect to overall budget is relatively flat, because most of the contacts used in this set of campaigns are resource channel-constrained, not budget-constrained. However, one channel that is very constrained by the business rules used in the optimization scenario is the interactive voice response (IVR) channel which is used to both receive inbound calls and also to make automated outbound dialing calls to selected customers. This channel resource is currently bounded, but if an increase in the number of IVR calls that could be allocated to this campaign were to grow by decision of the marketing strategists, the optimization scenario calculates that the increase in expected profit would grow sharply to a point, and then level off. This deceleration in the elasticity with respect to IVR calls is due to the penetration of customers that are likely to respond to an IVR campaign, which gets saturated at roughly the same point as the bend in the curve. After this point, there is not as much to be gained by making more IVR calls since the number of customers to respond has been captured and more investment will not result in greater profit, all other things being held equal.

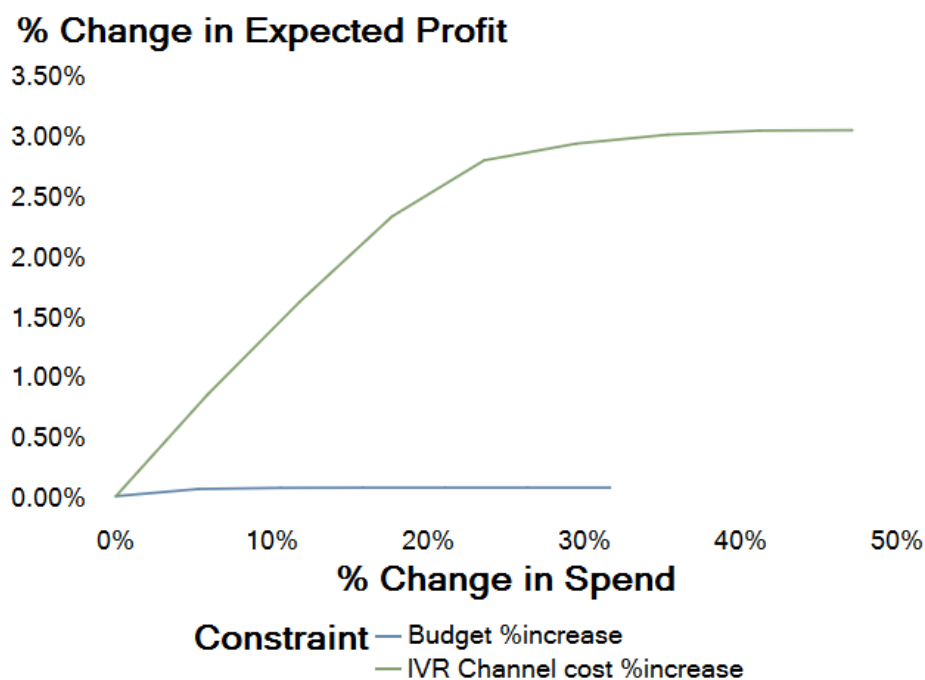


Figure 5: Comparison of marginal costs and returns for two constraint business rules

Consider for instance the capacity of the outbound telemarketing or call center to call customers for cross-sell offers. Compared to direct mail and email, as depicted in Figure 6, the call center is a relatively expensive channel through which to contact customers, and is also finite because it consumes minutes of call time and prep time that could be allocated to a different offer or a different customer. However, these costs are usually offset in part by the higher interactivity and hence higher expected success rate of a call to close a cross-sell opportunity. The challenge is to allocate the call center resource to the best offer-customer combinations, which is typically for those customers whose response rate for high margin offers will make the call center assignment an efficient one. Balance this with the choice of using a personally-assigned representative, such as a private banker, which is more finite and more costly to the firm, but through which the best customer-offer combinations could be allocated, and removed from the more general call center channel decisions, and this makes the selection among these trade-offs even more nuanced.

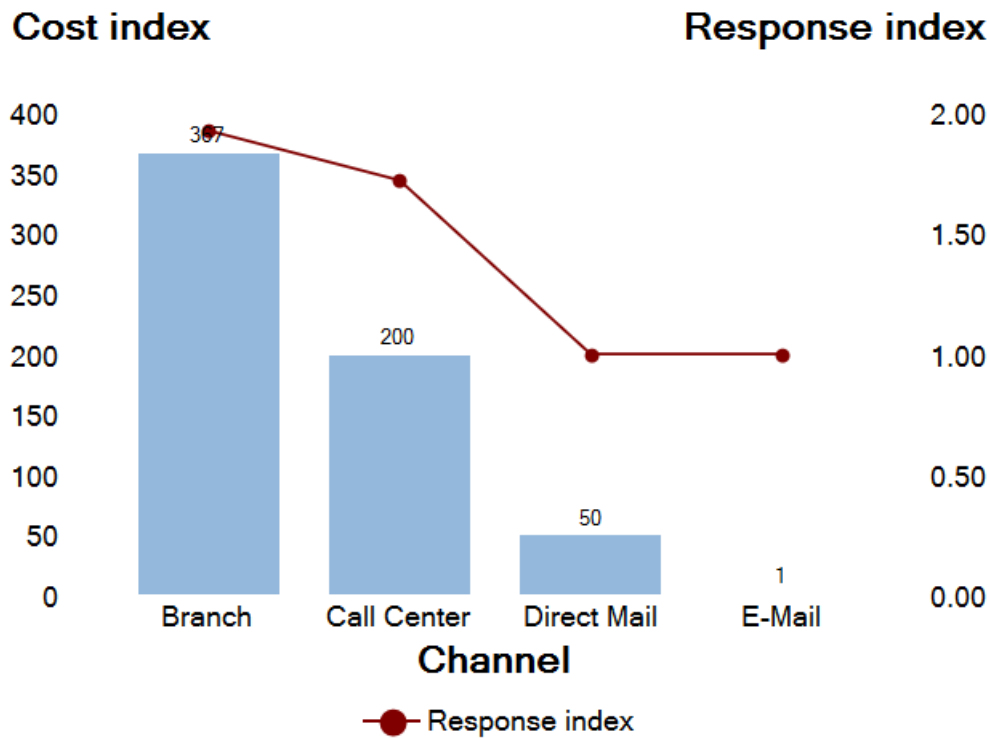


Figure 6: Comparison of channel-specific cost and response indices

Another trade-off that cross-sell marketers face is short-term response rate maximization for existing, well-known offers, compared with long-term maximization of innovative offers for which little or no campaign history is available. This is an issue we frequently encounter in the field, when management starts to realize that a catalog of well-developed predictive scores for campaigns and communications is an important asset to drive mathematically optimized campaigns. Reserving the necessary campaign contact volume to test new campaign strategies will be countered by the demand of competing campaigns to drive short term revenue. Campaign managers who recognize that next year’s revenue will be driven by introducing and properly testing new strategies today will place emphasis on making those short-term sacrifices. It is usually not yet clear to most campaign managers, however, that they need the proper metrics to gauge how much short term revenue they need to sacrifice, and a couple of important factors are at play. First, the discount rate of less certain future revenue compared with more certain immediate revenue needs to be applied, and the impact of the volume for those tests.

For instance, let’s say that a new offer called Offer B needs 75,000 contacts for the next month’s campaign in order to properly measure the response likelihood as well as the revenue per sale. Offer B is for the same product as the current champion, Offer A, and any customers eligible for that product can only receive either Offer A or B in the upcoming campaign. Product strategists believe Offer B, once perfected and rolled out in production campaigns, will improve the response rate by 75 basis points and improve sales revenue by 10%.

One factor the marketing team needs to consider is, if Offer B’s current-state test might not generate the same sales rate today as Offer A, will the opportunity cost to test Offer B be worth the sacrifice for-certain sales revenue by giving up the 75,000 contacts today. Even more important is, will the gain in new information produced by putting Offer B in market be worth

the 75,000 contacts, and could a similar gain in information be available by testing with only 50,000 contacts or fewer? This is a factor measured by the statistical power of the test itself, and by the experimental designs put in place by the marketing analytics team. This analysis might reveal that the certainty of new information would actually be stronger if the test was on 125,000 contacts. The marketer might also determine that the revenue to be produced on the most ideal Offer A customers won't really be sacrificed if they target more accurately.

The point is that the mathematical optimization discipline allows the marketer to measure these metrics directly, rather than just making the marketer guess blindly. The marginal cost of Offer B test can be measured directly as a function of the constraint forcing the 75,000 contacts into the Offer B cell, and this supposed trade-off in short term revenue can also be measured directly, before the campaign strategy is placed into the market. This underscores the importance of the test-and-learn discipline as a marketing analytics strategy that embraces techniques like experimental design, cell sizing, predictive scoring and mathematical optimization.

Benefits of mathematical optimization for cross-sell campaigns: a case study

The retail marketing group for a major bank who had developed a fairly mature cross-sell campaign strategy adopted a marketing optimization approach because they were struggling with numerous issues that limited their overall campaign profitability. This bank had developed a monthly cycle for multiple offers and channels, and a substantial depth in response scores for all these offers, but struggled to also balance these scores with their contact strategies and their boundary policies on contact channels.

The bank sought to maximize the profitability of their campaigns by optimally aligning offerings and customers. The group executes more than 30 campaigns and hundreds of offers every month and, unfortunately, faced ongoing roadblocks with their ability to manage, track, and optimize a myriad of campaign factors including:

- multiple marketing channels, including several call centers, email, direct mail, bank branch-based calls, and online banking
- cross-product halo effects, whereby the marketing analytics team had developed models to estimate response not only for Product A sales as a response to Product A campaigns, but also Product B sales as a response to Product A campaigns (and the like for the remainder of their product-offer campaigns)
- over-time contact strategies, including limits on the number of contacts per customer per 60 day period and rotation on offers to the same customer across campaign cycles
- offer cost across their channels, which could be as little as a few cents for email, up to \$15 for call center contacts and \$25 for bank branch contacts

- product sales goals for both high-margin offers (e.g., mortgage and home equity line of credit sales offers) and low-margin offers (e.g., deposit sales and service-oriented calls)
- margin contribution, by taking into account product margin estimates developed by their Finance function, as well as marketing-estimated margin contribution by channel costs and weighted by product retention likelihood
- customer response likelihood, for both direct sales and indirect sales

The bank's marketing team built a mathematical campaign optimization routine using a commercially-available marketing optimization solution. This solution allocated a margin-maximizing offer per client for each monthly campaign wave. After just four months of design and implementation, the bank implemented their first optimized campaigns in the field.

The optimized campaigns produced a net increase of \$22 in gross sales per customer for the first two months of optimized campaign execution. In the aggregate, the profitable sales growth for these two months averaged \$3.5 million per month in incremental campaign-driven income. The bank attributed this income growth directly to increased customer response and the profits from accepted offers. The bank now executes optimized campaigns on close to 80% of their campaign lead volume, several times per month.

Challenges faced by the marketing organization considering the adoption of mathematical optimization

Despite the very compelling value of adopting mathematical optimization for cross-sell campaigns, there are a number of cultural and process challenges that the organization will face as they develop a roadmap for introducing this approach in their marketing function.

The first and most fundamental shift is a cultural one. Many marketing functions are accustomed to a **blend of art and science in the construction and tracking of marketing campaigns**. Introducing a mathematical engine to augment the process of strategy design and development will make many traditional marketing strategy practitioners feel like they are having their role disintermediated by a black box. The flip side of technology adoption is that the organization can start to ask questions and have them addressed by a numerically-intensive approach in ways they've never seen possible before. Even the most analytically-savvy practitioners will face some challenges in upgrading their existing processes to ensure they integrate well with a campaign optimization discipline. For example, back-end performance measurement of optimized vs. non-optimized campaigns is one area that commonly faces some pressure to upgrade their measurement standards, to align with the test-and-learn methodology of which optimization is a major component.

Another challenge to organizations considering the adoption of marketing optimization approaches is the lack of a **sufficient history in test and learn-based marketing campaign design**

and measurement. Without this basis in measurement, they will struggle to develop the predictive response scores needed to allocate offers to customers. Often these organizations desire to adopt a marketing optimization approach for the right reason, but need to first develop their measurement and predictive scoring foundation. In this case, they have the opportunity to build a rapid model development and deployment process that will allow them to reach the point of beginning to optimize campaigns within a 6 to 12 month period, rather than following a more traditional analytics evolutionary path which might take 3 to 4 years to reach the same point of readiness.

The corporate adoption of a **customer-level contact strategy** is another factor with which many organizations will struggle. Contact strategies that place boundaries on how many times customers can be contacted proactively by the organization, by channel, offer and over time, represent an evolutionary stage of marketing that is itself difficult to reach without buy-in and coordination by product and channel teams. Marketing optimization approaches are unique in their ability to perform side-by-side testing and analytical evaluation of the fiscal and marketing impact of different contact strategies, to identify not only how many times the company may contact the customer, but also what is the likely change in marketing-generated income. The benefits of adopting a contact strategy is to avoid the state of customer fatigue where the customer will stop considering whether to review the offer at all. One analytic initiative that accompanies contact strategy development and testing is the long-run effects on customer retention and profitability, which takes years to measure properly, and good coordination on outbound channels to ensure contact rules are being followed properly.

Some organizations have adopted **net lift analytics** to measure the response of the customer to an offer in the presence of an incentive to respond. For instance, a marketing offer could contain one of three options for cash-back upon sign-up: \$25 cash-back, \$10 cash-back, and \$0 cash-back (i.e., the latter offer is really the control cell for the net lift modeling, in which case the firm doesn't provide an incentive to the customer). A fourth customer cell is the offer control cell for the baseline response modeling. A challenge that marketers will face, once they have developed the net lift model scores as well as the baseline response model scores, is which metric around which to optimize their campaigns. A marketing strategist with a strong affiliation for economic efficiency and a firm belief in the efficacy of their net lift models will argue for optimizing around net lift; whereas a more traditional marketing strategist will likely argue for optimizing on overall response volume and to drop those customers with low net lift to a smaller incentive. There is no one true global answer to this question, but our recommendation is that the marketing optimization approach should allow the organization to compare and contrast the two approaches directly for their fiscal and campaign impacts. Indeed some optimization approaches will allow the marketing strategy team to combine these two goal-seeking objectives together. We have found across a variety of projects that the correlation between baseline response and net lift response is on the order of 40% to 70%; this means that there is a substantial segment of incentive-driven customers who will respond mostly to the incentive in addition to responding to the offer, but there is separation between these segments. Because the marketing organization will pay for the incentive prior to benefitting from the net income that accrues from response, the ROI expectations on these campaigns is even more important to tune in advance for the optimal fiscal impacts.

A final important challenge faced by marketing organizations in the adoption of marketing optimization approaches for cross-sell campaigns is **insufficient spread in variable cost** between contact channels and offers. The only time we have ever observed a marketing optimization investment not offering a very compelling return for the marketing organization is when all offers are distributed across digital channels, where the variable cost per contact is low and essentially non-differentiated, other than for the fixed price investment up front for marketing program administration and creative development. This represents a case of the organization

facing no effective economic trade-off between offers to be assigned to customers; very often, these organizations also don't spend nearly as much time worrying about which one offer to deliver to the customer, because they are planning on the presentation of multiple offers at the same time. For purely digital strategies, the marketing optimization framework of selecting a limited number of offers for each customer, and optimizing expensive, finite-capacity channels for offer delivery, simply aren't aligned.

Lessons learned for cross-sell campaign optimization

We have implemented over 30 mathematical optimization applications for financial services companies, including cross-sell and marketing strategies as well as risk management, pricing and collections strategies. We've learned a lot about the steps needed to get the most out of mathematical optimization applied to these customer strategies and we've summarized some of the key lessons learned herein.

First, during early consideration to adopt a mathematical optimization approach for cross-sell campaign performance enhancement, take some time to set **clear business expectations**. Develop a sound appraisal of the business process for campaign design, campaign execution, the role of analytics in the current and future campaign process, and performance measurement. Then identify a clear position for the optimization functional capabilities, expected impacts and managerial communications as part of this planning and socialization process.

Next, identify **roles and responsibilities** for parties that will support a mathematical optimization function. While it might be obvious to assign this capability to an advanced analytics team, despite the advanced math running under the hood, this is more successful when placed in the domain of the marketing campaign strategy and execution team. The analytics team will provide some of the key inputs in the form of predictive response scores, but there are a wide range of scenario inputs and business rules that are more in the business domain. In order for this solution to be sustainable within the business long-term, it needs to be operated by the business analysts in the campaign function so they build the understanding and confidence to communicate to the rest of the marketing organization on the rules used for offer allocation. This team should establish a clear liaison to the marketing finance function to negotiate on business rules and support for marketing campaign funding, because one of the outputs of the marketing optimization scenario results will show the marginal return on the funding dollars for each campaign. Similarly, there is a clear role for each marketing channel liaison, such as allocating additional coverage for selected campaigns in the call center or the branches, where staff time is at a premium.

Set clear expectations for **early and ongoing results relative to business as usual ("BAU") marketing performance**. Organizations should consider setting up side by side testing for optimized and traditional campaigns so they can compare results before turning all campaigns over to the optimization engine. During the early stages of campaign optimization scenario construction, different stakeholders will react to the recommended campaign allocations with joy, with anger and with consternation. This is normal, and in every organization, compared to the traditional campaign approach, there will be some stakeholders that are winners and losers in terms of receiving campaign allocations. In order to avoid the reaction of "too good to be true", we recommend that practitioners build optimization scenarios initially with very simple

rules, and then add rules one step at a time. This will help the practitioner to understand the contribution of offer scores and business rules towards the optimal campaign allocation. Once a variety of scenarios have been constructed, the practitioner should develop interactive visualization charts and reports to view how the allocation of offers by products and customer segments makes intuitive sense given the inputs and rules.

Marketers who have relatively mature campaign processes already in place will often ask **which comes first, segmentation or optimization?** Option 1 is to run their campaign eligibility, selection and segmentation rules first, and then optimize the assignment of candidate offers to customers. Option 2 is to optimize all eligible offers first among the active customer base, and then use campaign strategies to select from among the optimal offers for each customer based on factors not used in the optimization scenario. Several commercial marketing optimization applications are designed to use Option 1, but it is our position that Option 2 provides a stronger and more economically efficient allocation of offers. This is because campaign segmentation rules are often based on gut feel and prior experience, but are rarely based on all the trade-offs that the optimizer function can take into account with regards to marginal costs and marginal revenues. By limiting the input set of choices the optimization engine can use via the campaign selection rules, the marketing organization is artificially putting limits on the return on investment they can earn. Instead, we advise running the eligibility rules as part of the optimization, generating the optimal set of (typically) three to five offers per customer, and then using the campaign rules to make the final assignment of offers to customers from within this subset of the most profitable offers. Furthermore, by viewing the stack-ranked list of profit per offer per customer, the well-advised campaign manager can place limits on the minimum expected profit that each offer allocated to each customer should earn the organization.

Each mathematical optimization is organized into scenarios, which are collections of rules and scores used to define the goals and boundaries that the optimization engine will pursue in allocating offers to customers. **The recommended initial scope of a scenario should be small** and focus on a single campaign for a single time period and the constituent offers in that campaign. This allows the team of practitioners to test their initial assumptions and evaluate whether they are mirrored in the mathematical optimization scenario results. As they gain experience, the optimization engine scenario can expand to running multiple campaigns and multiple time periods (e.g., the months within a campaign quarter) in order to evaluate the impact of rolling contact policies. Eventually, the organization will want to run longer-term optimization scenarios that evaluate the year-ahead marketing budget for budgetary planning purposes, to evaluate the potential impact of product introductions, and adjust for seasonal shifts in product adoption by customers and availability of channels.

Your marketing analytics team may believe they can invent a superior mathematical optimization engine or an offer allocation heuristic on their own, without needing to invest in a commercial class software application. They might be right; our recommendation is to define a **proof of concept** that pits the internal BAU with a commercial solution supplier or two. All parties should be able to complete this proof of concept in less than a month, with most of that time focused on sharing data with each competitor, defining and confirming that all parties will solve the same problem, but giving each competitor some leeway to bring their special sauce to the competition. The measurement objectives should be about more than the highest objective function, but should also focus on functional and process alignment for capabilities that are important for your marketing function and your team's skills. Make sure to consider a wide range of costs, both short-term and long-term, and any risks associated with each option. Whichever solution you choose, you should be able to justify a fairly rapid return on investment no matter whether you select a commercial solution or an in-house defined option. Finally, you should also evaluate each vendor's interactions with your team; the tenor and ease of

communication during the proof of concept are very likely to be representative of the working relationship during implementation.

In our experience, **no marketing organization possesses the perfect data** to adopt a mathematical optimization approach without some incremental effort. Instead, they should plan on identifying all their scenario objectives and the concomitant data requirements, and plan on collecting and structuring that data to fit the optimization scenario. Every organization's data will have some degree of cleansing required, and you should allocate team members with the sufficient data structuring skills to build all the input data, the staging to and from campaign databases and analytics tools. Examples of data sources that frequently need to be developed and refined during an engagement:

- **Response scores** for each offer are usually available, but factors to adjust response rates for the outbound channel often need to be built based on evidence or estimates, in order to offset the more-easily-estimated difference in cost per channel.
- Product-specific **expected revenue impacts** are usually provided to marketing by finance, but some further refinement is recommended. Most organizations are in a good position to estimate the net present value (NPV) of a future stream of income and costs. Some marketing organizations are content with revenue figures whereas others desire optimizing for margin contribution. Pursuit of the latter metric often results in an endless debate about the calculation of this margin metric and whether it takes enough of the indirect costs, fixed costs and transfer price structure into account.
- Financial services organizations should add the **risk-adjusted revenue** or margin into account when they are marketing lending products where a customer or prospect can default on the outstanding balance.
- Revenue per product sale need not be estimated down to the product-customer level, and it is unlikely that such a prediction model will be very accurate, nor will the marketing allocation decision change much when using **product-customer-level revenue estimates**, given our experience. It is highly desirable to stratify the revenue per sale by product-customer segment if possible. Finance is rarely in a good position to provide this breakdown to the customer segment level, so marketing will need to apply their best possible information to adjust product-level revenue metrics accordingly and prepare to defend those assumptions.
- Response rates at the customer level are often poorly calibrated, especially on the tails of the distribution. Most marketing analytics teams focus on the calibration of their response models for the predicted average response rate, but since a mathematical optimization engine will first allocate offers to customers on the **right-hand side of the response score distribution**, all other factors being equal, then some calibration across each decile of each response model's scores is a good idea.
- Some offers will require handling for inbound responses, such as products for which eligibility is determined at the point of the customer applying for the product and some underwriting activity then takes place. **Inbound response** scores can be defined using baseline outbound response rate as a starter, and then adding the incremental cost of that handling for each estimated responder.

- Similarly, the **cost of incentives** to drive response are also a function of outbound offer cost tied to the response rate, assuming that only eligible responders qualify for the incentive.

Few marketing organizations possess sufficient testing-based evidence to supply precise assumptions into the mathematical optimization process for the factors noted above, but our advice is not to despair. A helpful outcome of the optimization mathematics is that you can run multiple scenarios to test whether the allocation of offers to customers is really sensitive to your assumptions as you adjust them up and down in comparative scenarios. Some optimization of your marketing campaign results with pretty good data is better than never running any optimization scenarios for want of the perfect data.

Many marketers will have campaigns in place that tend to have produced steady results month-over-month, only to find that when they run their campaigns and scores through a marketing optimization engine, the recommended offer allocation can be quite different than the traditional mix. They might find that some offers traditionally under-represented in their mix rise to the surface as the top-allocated offer, and upon further review, the economics supporting that recommendation will make a lot of sense. In our experience, these offers are usually the ones in the middle range of the income generation stack with the middle range of the cost structure per offer. We're not quite sure why this is, but one explanation is that so much emphasis goes into product managers propping up low-margin offers and field sales demanding a strong mix of the sales compensation-generating big-margin offers, that the offers in the middle of the table get overlooked. Regardless, in order for the organization to believe that the optimized scenarios can be compared on an apples-to-apples basis, an appropriate optimization scenario should include business rules that constrain the allocated offer mix to +/- 10% of the offers for each campaign allocated via the BAU approach. The resulting offer allocation then takes full advantage of the differences in scores among customers and across products that are the best fit for each customer eligible for them. The marginal costs of the per-campaign constraints yield useful information about campaigns that should be increased or decreased relative to BAU in order to grow the overall profitability of the entire marketing campaign portfolio.

High-revenue products (e.g., mortgage loans) often dominate the offers being allocated by an optimization scenario, because the high revenue per sale dominates other offers by an order of magnitude when evaluated by the goal-seeking function, even if the response rates for those offers are relatively smaller than low-margin popular offers (e.g., deposit accounts). Our recommendation is to put ceiling business rules on cell sizes for such high-revenue offers to focus on those highest in the sales response likelihood, and route those offers through the highest-touch, highest-cost channels. We also recommend developing and using net response lift scores for those types of offers, as mortgage shoppers are frequently already in the market for such a product well before the campaign reaches them. This is often consistent with sales strategy for such products.

Response rates for offers are usually estimated as response for Offer A / contacts for Offer A, which is more specifically defined as a direct cross-sell response rate. Some marketing organizations with more mature performance tracking and model estimation capabilities also estimate indirect cross-sell response rates, which is defined as response for Offer B / contacts for Offer A (and all of the possible permutations across offers). The use of both direct and indirect cross-sell response rates produces a more accurate picture of the return on marketing investment and also allows the optimization engine to account for the benefits of lower-cost service offers that prompt a customer in latent sales mode to shift into an active sales mode. The investment needed to develop this wider range of response scores is substantial, and can

be dramatically accelerated when the marketing analytics organization develops a contemporary rapid model development and deployment platform.

Most optimization scenarios run in batch mode, usually in as little as 3-5 minutes for simple scenarios and as long as several hours for complex, larger scenarios. What's to be done when the marketing organization wants to enable the field with optimized offer allocations, but the customer reveals during a sales opportunity interaction that their situation has changed relative to the last time the marketing organization refreshed their profile? For instance, an outbound call to an existing customer who recently viewed several pages on the bank's web site and downloaded several information guides on investment strategies might reveal that the customer intends to get married and start a family; if this is new information to the bank, then the pre-staged optimized offers for a new car loan could be very much off the mark. Instead, the marketing organization should develop the ability to run a real-time optimization strategy for this single customer, by updating their profile for the information indicated during the early part of the call, and running an optimized offer allocation scenario intended for a single customer. This approach involves several technology components and analytic techniques that complement the marketing optimization engine, including real-time scoring and post-optimization offer allocation predictive techniques. The full description of such an approach is outside the scope of this paper, but more details are available from the author.

Additional case studies for cross-sell optimization, and other stages of the customer lifecycle

The applicability of mathematical optimization is not limited to traditional direct marketing cross-sell and next-best-offer campaigns. In this section of the paper, we present several non-traditional marketing campaigns that have been dramatically enhanced through the use of mathematical optimization, as well as a host of non-marketing customer strategies in the risk management, pricing and collections domains.

Recurring bill payment campaign optimization

We developed a series of predictive response scorecards and a mathematical optimization routine for a credit card issuer that wanted to maximize the dollar-weighted incremental response rate to a new marketing offer that offered targeted incentives to accounts that placed recurring bills on the credit account as the form of payment. Business rules were designed to control the outbound mailing and inbound incentive budgets and also provide sufficient capacity in all test cells to permit robust response rate performance testing. On an eligible universe of 8.6 million accounts, we targeted slightly more than 700 thousand offers, generating recurring monthly incremental revenue (for the life of the account participation in the program) of over \$725 thousand on a one-time marketing investment of \$695 thousand. Put another way, if the cardholders who signed up stayed on the recurring billing program for an average of six months, the incremental revenue to the card issuer organization would be roughly \$4.35 million on an investment of \$695 thousand, for an ROI of approximately 5.25.

Digital-only next best offer

We designed a simulation and optimization routine for a diversified retail bank to maximize the dollar-weighted response rate by assigning the ideal product offer from among multiple candidate offers to customers via a web and call center-based cross-sell engine. Business rules for the optimization routine included offer eligibility, response scoring, order presentment and cycling based on past contact and response history, and operational planning constraints. Simulation was used to project 30 days' forward site visits, impressions, click-thrus and responses over the web and contacts & responses via the call center. Expected incremental dollar-weighted ROI compared to business as usual were conservatively projected to exceed 15% over business as usual (equivalent to roughly \$2 average incremental response per customer) on a sample of 10 thousand accounts per day.

Brokerage sales lead distribution optimization

We developed a mathematical optimization routine for a retail investment brokerage to maximize weekly sales NPV by assigning the ideal sales lead per account and financial advisor across 2+ million accounts and thousands of advisors, by day over a multiple-period planning horizon, and by contact channel. Business rules used as policy constraints included advisor licensing tied to sales offer, advisor calling capacity, cost per contact and preservation of lead volume over days in the planning horizon. The incremental benefits compared to the prior approach for lead assignment was in the millions of dollars per multiple-week planning horizon.

Credit issuer initial line assignment optimization

We developed a mathematical optimization routine for a leading domestic credit card issuer to maximize risk-weighted net income by assigning the ideal initial credit line to each account at booking. Deliverables included behavioral segmentation tied to credit risk, utilization, balance and limit, and development of predictive scorecards for likely net income and gross credit loss. The optimization routine accounted for policies on total credit exposure, total expected loss, active rate and utilization, in order to calibrate the optimized assignment to current portfolio performance. The best-performing optimization scenarios projected an incremental net income contribution between 200% and 400%, and a margin-to-revenue ratio increase between 60% and 240%, compared to the business-as-usual performance on the historical accounts.

Credit issuer credit line increase optimization

In separate engagements, we developed a mathematical optimization routine for three credit card issuers to assign the risk-weighted net income-maximizing credit line increase amount to existing card accounts. The optimized assignment routine accounted for policies on total and at-risk balance exposure, percent and absolute increases in line per account and per credit risk

tier, and calibration of the optimized assignments to existing strategies. For these three engagements, the projected incremental net income ranged from over \$80 million annualized (on 1.5 million accounts), \$5 million annualized (on 350 thousand accounts) and \$60 million annualized (on 3.5 million accounts).

Auto portfolio collections optimization

We developed a mathematical optimization routine for the auto loan servicing department of a retail bank to maximize the risk-weighted returns on collections activities by assigning the ideal delinquency & recovery strategy to each of roughly 170,000 accounts in the loan portfolio. Key deliverables included behavioral segmentation and the construction of over 20 delinquency & collections scorecards and a collateral valuation model. The optimized strategy assignment routine accounted for variable costs and resource capacity per collection strategy, the eligibility of strategies per account given the maturity of the loan and the collateral book-to-value. Incremental benefits for the solution estimated at \$2 million in incremental net receivables per month.

Conclusion

As we have described in this paper, the benefits of the mathematical optimization approach to cross-sell offer allocation are simple to explain, and have brought many organizations substantial performance lift in their cross-sell and next-best-offer campaigns. The value of new information for testing new offers that will support longer-term marketing growth is usually offset by a trade-off in short-term response volumes for already-tested offers. A significant benefit of using a marketing optimization approach is to directly measure these trade-offs so that management can make the optimal investment.

Adopting a powerful approach like mathematical optimization brings its own risks and investment requirements, but a well-informed marketing organization can successfully adapt to these risks in order to deliver on the promise. There have been plenty of successes with other organizations, and the long-run benefits of these investments have been very strong in comparison with the investments. Few organizations have the perfect data platform in order to launch a marketing optimization approach with the desired precision, but again, the use of this approach yields direct feedback about the sensitivity of the marketing offer allocation to customers and the fiscal impact of perfect information, by allowing the marketer to ask themselves, “would I make a different decision if I had even better data”.

Interested in learning about how Corios can help you get more from your next best offer strategies? Give Robin Way a call at **503.295.1685** or email him at president@coriosgroup.com.

About the Author



Robin Way

Robin Way is the Lead Faculty Member for Banking at the International Institute of Analytics, and is the founder and CEO of the management analytics consultancy, Corios. He has over 25 years of experience in the design, development, execution, and improvement of applied analytics models for clients in the credit, payments, lending, brokerage, insurance and energy industries. Robin was previously employed with SAS® Institute's Financial Services Business Unit as a managing analytics consultant for 12 years, in addition to another 10+ years in analytic management roles for several client-side and consulting firms.

Robin's professional passion is devoted to democratizing and demystifying the science of applied analytics. His contributions to the field correspondingly emphasize statistical visualization, analytical data preparation, predictive modeling, time series forecasting, mathematical optimization applied to marketing, and risk management strategies. Robin's undergraduate degree from the University of California at Berkeley and his subsequent graduate-level coursework emphasized the analytical modeling of human and consumer behavior.

About Corios

Corios is a leader in the discipline of management analytics consulting focused on helping clients across industries to bridge the gap between their data and their business decisions. The company's custom offerings provide analytical solutions for risk management and compliance, marketing, pricing, and big data initiatives. Corios' solutions have identified business improvements worth hundreds of millions of dollars for corporations across North America.

Corios believes that data-driven techniques are the key to making profitable business decisions, and that data should be simple, approachable, and implementable for both analysts and decision makers. The Corios team is persistent in their efforts to promote transparency and common process across the analytical discipline.

To learn more about how Corios is bridging the gap between data and business decisions, please visit coriosgroup.com.

Next best offer campaigns

Copyright © 2014

Published by Corios
133 SW Second Avenue
Suite 202
Portland, Oregon 97204



All rights reserved. Except as permitted under U.S. Copyright Act of 1976, no part of this publication may be reproduced, distributed, or transmitted in any form or by any means, or stored in a database or retrieval system, without the prior written permission of the publisher.

Visit our website at: coriosgroup.com

