



"turning data into dollars"

Tom's Ten Data Tips – November 2007

Campaign Optimization

Campaign optimization can take place at three levels:

- Finding the best possible targeting model for a campaign
- Finding the optimal target group(s) *within* a campaign, possibly contacting the customer by one or more channels
- Determining the best possible offer *across* multiple simultaneous campaigns

To elaborate on these three levels:

- Arriving at the best possible model may be considered 'merely' a statistical challenge, but it's an intricate one. Much has been written about this, mainly in the Machine Learning literature.
- Finding the optimal target groups should be driven by explicit cost/yield considerations. As straightforward as this may sound, the entire concept of "optimization" deals with maximizing yield. As such, this merits special attention.
- Determining the best possible offer across multiple campaigns means choosing one offer in favor of another, and possibly planning the *order* in which offers will be made.

1. Campaign Optimization Is About Finding A Balance Between Contact Strategy, Short Term-, And Long Term Goals

There are three fundamental constraints that need to be balanced when optimizing campaigns:

- Business rules/contact constraints with regards to how a company chooses to engage in a dialogue with their customers. Possibly, some customers can only be approached via account managers, in other cases customers will have a maximum contact frequency, etc.
- Sales/performance targets, short term objectives. These are the quarterly or yearly (commercial) results that are expected from middle managers.
- Strategic corporate imperatives, long term objectives. Now we are talking about strategic goals that follow from mission and driving force of the corporation. This can be like GE's goal to be #1 or #2 in every market they operate in, or a critical mass that is necessary to compete effectively, etc.

The business world isn't perfect, and these three constraints are never perfectly aligned.

2. Insight In Cost/Yield Drives More Rational Targets

Challenges in campaign optimization are typically the result of misalignment between internal corporate targets (see tip #1). When the same customer is eligible for more than one campaign, and not all will be offered simultaneously, there needs to be some procedure for arbitration.

The discussions about which campaign should get priority can sometimes appear "politically tainted". However, the better one can empirically demonstrate cost/yield considerations (short *and* long term!), the more such discussions will converge to a rational optimum. That is because the argumentation will be fact-based so that all associates can make an objective assessment of pros and cons of alternatives under consideration.

3. There Are Three Kinds Of Optimization Across Campaigns

Optimization across campaigns, when a customer is eligible for more than one offer at the same time (and not all offers will be made simultaneously), can occur in three flavors (see also tip# 4-5-6):

- Customer centric; in this case one makes the offer that has the highest probability of being accepted (see also tip# 9)
- Company centric; response probability is multiplied with the Net Present Value of the product, minus cost per contact
- Marketing centric; based on customer lifecycle considerations, it's possible here to consider "loss leaders" if they offer sufficient opportunity for profitable cross- and up-sell, or to emphasize "strategic" products based on corporate directives

4. Customer Centric Marketing Hampers Revenue Maximization

The grand idea behind customer centric marketing is to offer the right product, at the right time, to the right customer. In customer centric optimization, one "simply" offers the product that customers are most likely to accept. The downside of this approach is that product profitability is entirely neglected in this equation. Few organizations are willing to walk the customer-centric talk this far, and probably rightfully so.

5. Aiming For The Highest Yield Is *Not* Customer-Centric

A company centric optimization strategy maximizes the yield. The optimum is $\text{MAX}(\text{response} \times \text{NPV})$. Do not be misled by the consideration of response probability. Profit from a customer need not bear *any* relation with the value he perceives. If anything, it is a priori more likely to be inversely related, since this is pretty much a zero sum game. More competitive pricing, economical for the *customer*, that is, equals lower profit margins for the company that they would need to make up for in volume.

It is good practice to take the contact (and response) history into consideration to avoid repeatedly offering products the customer has demonstrated no interest in.

6. Marketing Optimization Requires Deep Insight In Drivers Of Growth And Retention

Marketing centric approaches are neither about offering the product with the highest response probability, nor about offering the product with the highest yield. Instead, they are based on a profound insight in acquisition patterns that allow companies to not only consider immediate purchases, but instead regard the customer lifecycle. They require some ulterior goal like Life Time Value, presumably also in line with corporate strategy.

An example could be offering a so-called "loss leader", products that serve as a jumping board for future customer development. Alternatively, when acquisition of particular products are associated with a low propensity to churn, these may well contribute significantly to LTV, despite *themselves* being only moderately profitable. Such profound insights in the workings of the market place never come easy, and are typically the result of persistent data exploration.

7. Mind Your Success Rates In The Front Office

Campaign optimization should definitely take into account *how* the offer will be presented to customers. It matters whether this is in a human dialogue or via some "automated" channel like the web, Voice Response (VR), or direct mail. Humans are prone to "rejection fatigue", and this can occur when the most profitable offer (for the company) is consistently recommended, but this happens to be a product with very high yield and rather low response probability. A tactic like this can really burn out front-line sales staff in an "intelligent" analytical CRM environment.

Some may argue that an empirical optimization procedure would “discover” this relation by itself. This may be true, but what you *know* need not be learned by an algorithm. Moreover, one need not burden an optimization engine with discovering *the right* signal from the noise.

8. Finding The Optimal Targeting Model Is An Optimization Task

There are several considerations that need to be optimized when choosing the best possible model. How hard is it to implement a new model or replace an existing one? How long does it take to build the model, soup to nuts? Once in place, what is the expected lifetime of the model, or, what is the (expected) degradation of predictive accuracy over time? At any point in time one needs to consider how improvement from a new model should be weighed in light of replacement costs (and risks).

There is a tradeoff between a model with the highest possible accuracy in the short term, versus a model with slightly poorer accuracy initially, but a longer half time (model performance decays slower). Other considerations are how well (in terms of statistical power) one can monitor model performance, or how transparent the model is (is the outcome easily explicable). The “best” model is perforce a compromise, taking all these criteria (and possibly more) into consideration.

9. The Relation Between Propensity Scores And Response Probability Is Non-Obvious

In particular when multiple tools have been used, and the oversampling rate in the training data are different, calculating which offer has the highest probability of eliciting a response becomes a complicated matter to determine. The analytical relation between model scores and the empirical response probability is unknown to date and may never be analytically derived. All there is to do is integrate this function by brute force.

10. Campaign Optimization Is *Not* About Linear Programming

We hold the firm belief that *by far* the biggest gains in campaign optimization are the result of more rational corporate target setting. This is people work, essentially a negotiation outcome across multiple

layers of management. In our practice, the largest strides have been made when those discussing “the best” business targets, alternated between optimization levels #2 and #3, determining optimal target groups (on the basis of cost/yield), and finding the optimal offer for each customer.

The better this discussion is informed by a sound understanding of prevailing market conditions (see tip #6), and when cost/yield considerations (see tip #5) are part of the organizational DNA, the better off the corporation as a whole will be. This kind of fruitful strategic discussion will lead to targets that pull in unison towards a mutual corporate “sweet spot”.