

BETTER RETENTION WITH UPLIFT MODELLING: SAVE MORE CUSTOMERS WHILE SPENDING LESS

TRIGGERED CUSTOMER CHURN

Customer Service Rep: Hi, I'm calling from North-South Mobile. As one of our most valuable customers, we'd like to offer you a free handset upgrade if you just extend your contract with us.

Customer: You mean I'm out of contract? That's great, can I cancel right now?

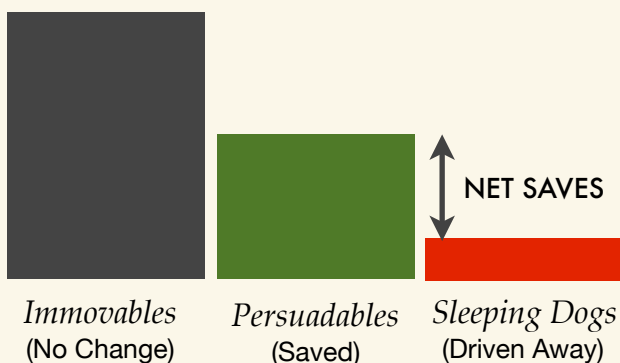
The core goal of retention activity is to persuade profitable customers who would otherwise have left to remain as customers.

The fact that a customer is high value, or is at high risk of defection, does **not** mean that they are likely to be savable. In fact, many retention initiatives trigger the very churn they are trying to avoid among some customers.

Most retention activity doesn't change the churn behaviour of the majority customers: those who would have stayed stay, and those who would have left leave anyway. We call customers in this group "*immovables*".

The interesting customers are the "*persuadables*"—those who would have churned but are persuaded to stay by the operator's retention activity—and the "*sleeping dogs*"—people who would in fact have stayed, but are triggered to leave by the intervention.

Successful campaigns save more customers than they drive away



Sophisticated operators measure the net effect of their retention activity by using rigorous control groups and can therefore work out a meaningful return on investment (ROI). But even if a campaign saves customers overall (and perhaps delivers a positive ROI), that does **not** mean it isn't triggering churn among some segments: it just means it's doing more good than harm overall.

The problem is that conventional, state-of-the-art targeting models are fundamentally designed to select customers on the basis of **churn risk** rather than **savability**. This is ironic, because often the customers at highest risk are some of the hardest to save, because the reason they are at high risk is often because they think they are being overcharged, receiving poor service or are otherwise dissatisfied.

Typical retention activity is counterproductive for about 30% of customers targeted in this way: by including them in retention activity, the operator is actually **spending money to drive away customers**.

*Uplift models
target on customer savability,
rather than on churn risk*

Uplift models are not just mathematically different from conventional response models: by modelling both treated and control customers, they predict how likely a given retention strategy is to work for each customer. This allows operators to concentrate activity where it is most effective—on the **Persuadables** segment rather than **Immovables** and **Sleeping Dogs**. This reduces costs and increases net saves, allowing higher ROI (or, in some cases, positive ROI where previously value was destroyed by counterproductive activity).

Uplift modelling is most effective for large consumer bases with relatively high churn rates, in operators with a willingness to use control groups to measure and optimize retention performance.

Banks, credit card companies, insurance firms and anyone else with contract-based relationships face very similar issues to telephone and internet providers, and can benefit from exactly the same techniques.

UPLIFT MODELLING

Uplift models seek to answer the question:

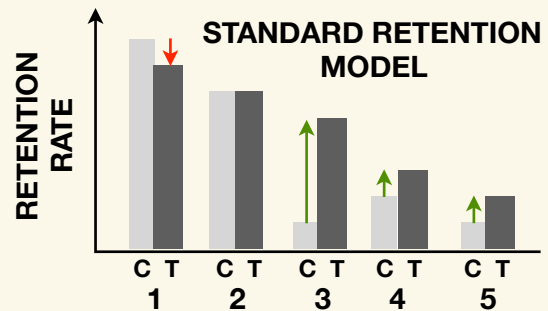
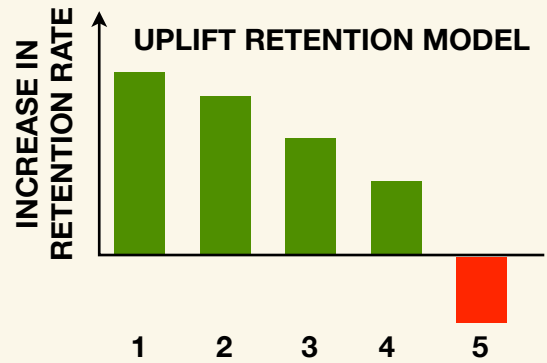
*What is the **decrease** in the probability this customer will leave if we take a given retention action?*

This prediction allows targeting on the basis of a model of the key quantity that directly drives return on marketing investment (ROI).

In contrast, traditional retention models seek to answer the question:

*What is the probability that a person will leave if we do **not** include them in the retention activity?*

This often leads to very different targeting, as shown. In this case, the action actually has a negative impact on the “top” (first) quintile from the model, and it is the third quintile that shows the biggest uplift in retention rate between the control group, **C**, and the targeted group, **T**.



GETTING STARTED WITH UPLIFT MODELS



Wherever you are on your uplift modelling journey, Stochastic Solutions can help. We were founded in 2007 by Nick Radcliffe, who has been developing uplift modelling ideas, algorithms and software since 1999.

Companies that already know they want to use uplift modelling in-house can license our Miró software and benefit from our training and support services.

Alternatively, if you already have data from a suitable campaign with valid, randomized controls, we can build a trial model for you and validate it in market to prove the ROI of the solution.

If you are earlier in the journey, and aren't yet using control groups to assess incrementality, we can help you to get started and measure the true impact of your current activities accurately. We can help you to design campaigns that will put you in a position to support uplift modelling.

And of course, if you aren't ready to bring uplift modelling in house, we can offer bureau modelling services on an ongoing basis.

Miró is our integrated solution for building uplift models. As well as supporting rich general-purpose modelling, interactive query and data manipulation, Miró features direct uplift model building capabilities and support tools for uplift modelling, analysis and operationalization.

These include:

- Significance-Based Uplift Tree Building
- Qini measures for assessing overall performance of uplift models
- Qini-based variable selection for choosing fields based on uplift relevance
- Uplift crosstabs for exploration & reporting
- Export and operationalisation of uplift models
- Support for noise reduction methods including stratified sampling and ensemble models

Miró is available for Linux, Mac & Windows for web-based, interactive and batch use.

Contact us at info@StochasticSolutions.com