

WHITE PAPER

Using AI/ML for Customer Churn Reduction in Private Banking/Wealth Management

A Point of View by TCG Digital





Contents

Case study - 1	
Customer Churn in Private Banking	2
Case study - 2	
Customer Churn at Wealth Management	6
Initial State of the Client	6
Client churn & downgrading specifically visible in the advisory service landscape	6
TCG's approach to churn management	7
A churn prediction model enabled risk classification of clients based on CRM data	7
A bespoke Banker Tool allowed for a more effective management of client relationships	8

Case study - 1:

Customer Churn in Private Banking

First of all, customer churn in private banking is an infrequent event. Typical rates are around 1-7% with some dependencies on geography and client segments. In addition, there is the challenge of addressable and non-addressable churn. Death, for example is a non-addressable churn issue (even if the assets remain in the bank).

The problem has two primary components. The first is predicting churn in an accurate fashion. Second, we must understand the drivers of that churn. The understanding is as important as the prediction, because the bank needs to develop strategies to address the potential causes – before a customer leaves the bank.

From a data science perspective, the first component comes down to dealing with an imbalance class data set. The second component is about surfacing the nuances in client groups and their behaviour patterns.

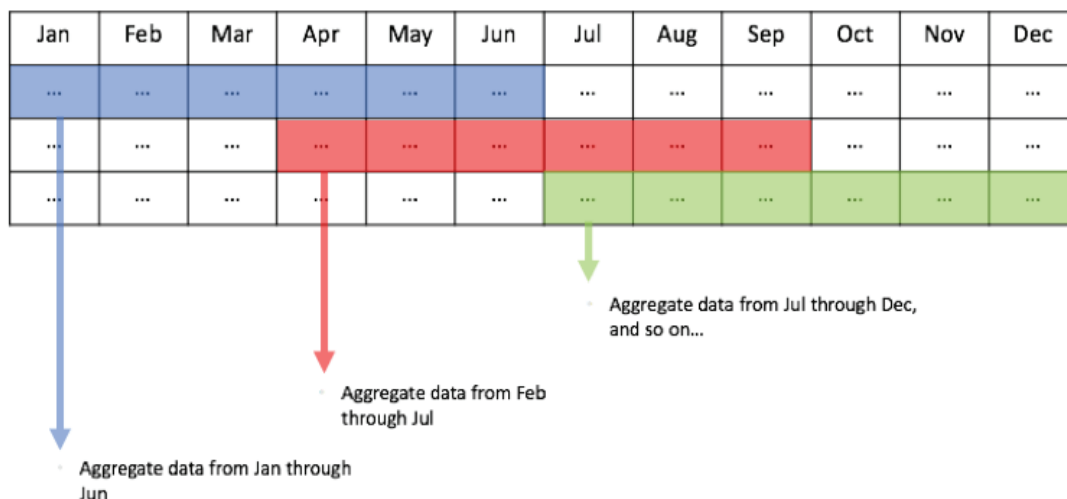
At TCG, we solve these problems by topological segmentation of data sets, building local models for groups surfaced by topology and providing statistical explains for those groups.

To elaborate how TCG’s AI platform tackles these challenges, we’ll use a recent customer project as an example.

Our customer is both innovative and data-savvy. They are committed to leveraging every piece of data at their disposal to create better customer experiences, predict and ultimately intercede prior a customer leaving.

The data they possess on their customer includes demographics, asset summaries, accounts, equities and holdings, trading history, transfer history, contacts and meeting logs. In general, banks capture clients’ data in a cross-sectional fashion, i.e. for every constant period they take a snapshot of every client. Concretely in the example here, our customer took a snapshot of all their clients at the end of every month.

In order to capture the clients’ temporal pattern, we started off the data transformation by aggregating client data statistics for a time window, say, 6 months, to predict client churn in the next 6 months. We moved the time window forward with an interval of 3 months and kept doing this until our time window hit the end of the analysis period. We then concatenated all these window slides except the last window into one training data set. We reserved the last window slide as an out-sample validation set for final model evaluation.



Once the training data set was ready, we kicked off our analysis workflow.

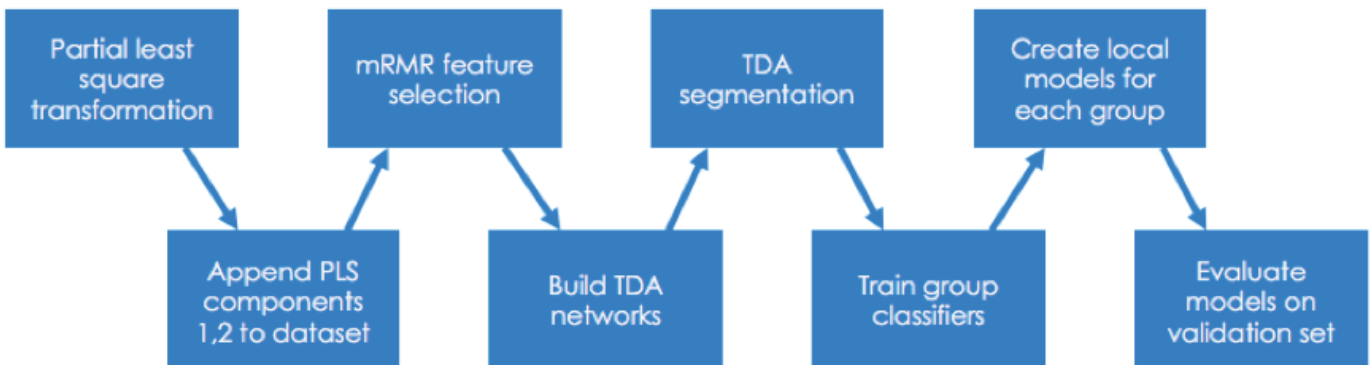
First a partial least square (PLS) transformation was done on the data matrix with respect to the outcome, which in this case was churn. The first two PLS components were appended back to the training data set.

Next, we utilized the minimum redundancy and maximum relevance (MRMR) algorithm on our platform to pick the most highly relevant features. By identifying the most relevant features we enhanced our TDA networks because we were able to isolate noise – producing better groups or clusters of customers.

The network was then coloured by the outcome variable and the locality of the outcome could be checked.

This network supports operational deployments where groups would be created by our auto grouping algorithms. This facilitates the creation of topological segmentations on the data set allowing for the creation of local models for each segment. A group classifier would also be trained and comes into the prediction pipeline before a local model.

When a new data point comes in, the group classifier return the probabilities of falling into each group and each group’s local model gives the probability of churn. The final churn probability is then the local model probability conditioned by the group probability. This workflow reduces systematic bias compared to a global model for the whole population.



While certainly not an ahh-ha moment, it was immediately clear that clients churned for a wide variety of reasons demonstrated by their distribution across the network in Figure 1. Our auto grouping algorithm identified the high churn rate groups within a few seconds. Furthermore, the characteristics of each group was immediately apparent via the explain tables.

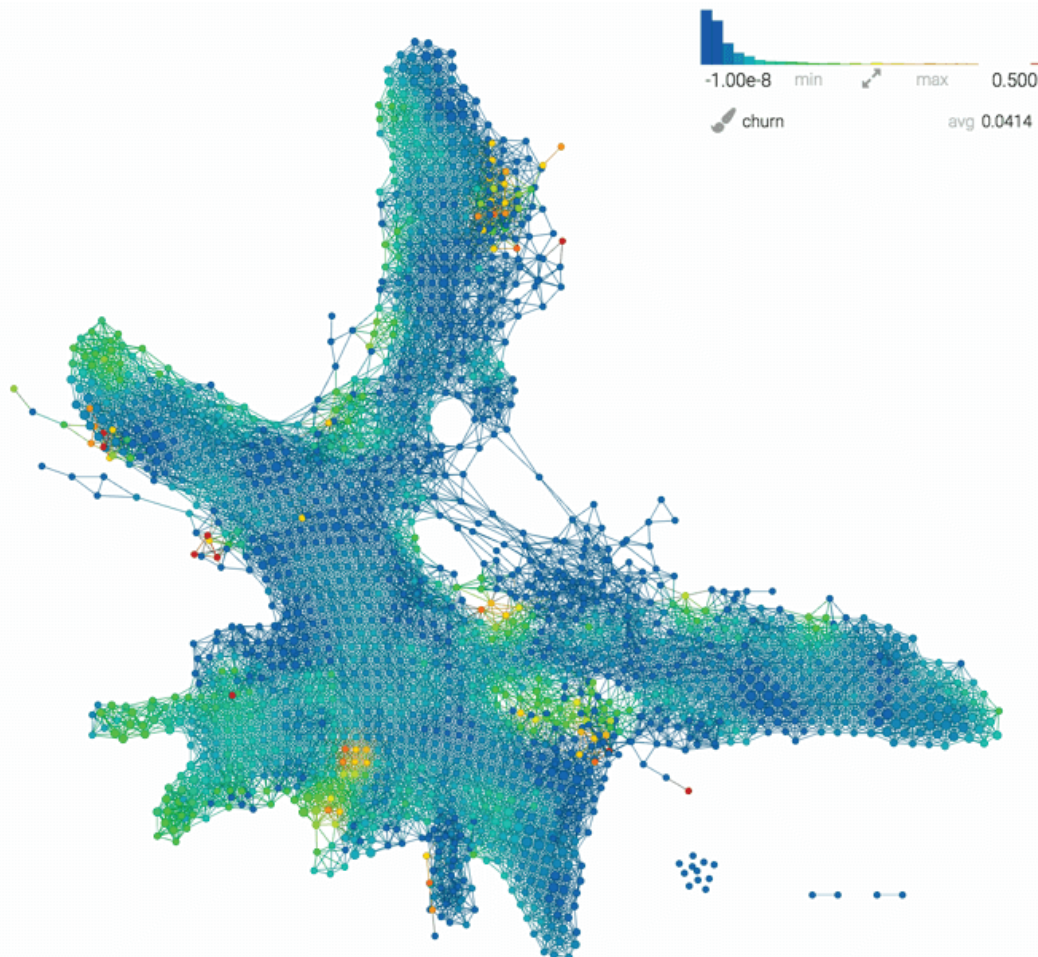
The Kolmogorov-Smirnov test on continuous variables and the hypergeometric test on the categorical variables sheds light on the nuances from one churner group to another.

Take the high churn rate group on top of the network for example (coloured red + yellow). This client group tended to have lots of missing information on their marital status (80% in group vs. 20% in the rest of the data) and they were


all women clients. They also tended to have low balances in their asset and investment accounts while also having low numbers on transfers and trading activities.

If we move to the lower left corner of the network and focus on the high churn rate group there (again coloured in red and yellow), the pattern is significantly different. Gender was 100% male and one of the outstanding reasons they churned was the underperformance of their assets.

In the current context, TDA not only helps our customer to predict the clients at risk but also to understand the variety of underlying drivers for leaving the bank. This client profiling significantly alleviates the painful process of random guessing on the appropriate and relevant retention proposals. Without TDA, these client insights typically take months or years to develop. We surfaced them in a matter of days.



Remember this is an imbalance data set (base churn rate 3.4%) and if the classifier for some reason labels every client as non-churners one would get an accuracy of 96.6%, which is misleading. The appropriate metric to evaluate the final model will be the ROC curve and the area under the curve (AUC), which is show below.



The dash line in the chart represents a random guessing model and the blue line is our final model, which produces a 24% lift.

As one moves along the line, it's basically a trade-off between recall and precision. In the context of client churn, the recall (or sensitivity) gives the gauge of how many churners our customer will catch out of all true churners and the precision gives our customer the actual churners out of all predicted churners.

From a private banking operation stand point, it's a trade-off between client asset retention (higher recall = more client outreach, more expense but more assets retained) and customer relationship maintenance cost saving (higher precision = less client outreach, therefore cheaper but some churners get through). Our model provides our customer with a comprehensive view into their trade-offs and the expected return of their client retention efforts.

The aforementioned modelling, segmentation, profiling and prediction workflow was based on our customer's historical data. In practice, one can set up an automatic process in operation to retrain the TDA model once new data are available so that the model always stays up to date and adapts to pattern shift in client behaviour.

Case study - 2:

Customer Churn at Wealth Management

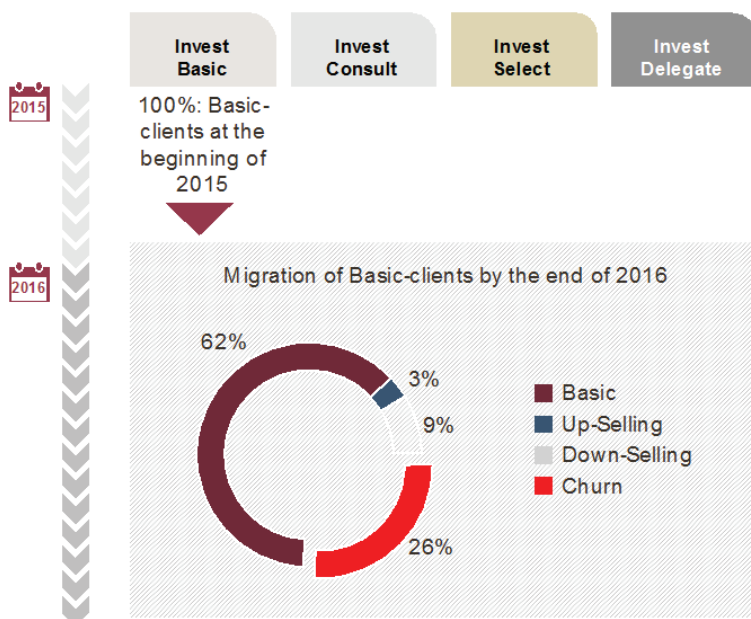
Initial State of the Client

	Invest Basic	Invest Consult	Invest Select	Invest Delegate
Advisory contract	✓	✓	✓	✓
Risk profile	✓	✓	✓	✓
Custody account	✓	✓	✓	✓
Transactional advice	4 incl. p.a.	12 incl. p.a.	✓	✓
Transactions	Charged separately	12 incl. p.a.	24 incl. p.a.	✓
Proactive recommendations	x	x	✓	✓
Portfolio review	1 incl. p.a.	2 incl. p.a.	2 incl. p.a.	3 incl. p.a.
Account statement	1 incl. p.a.	1 incl. p.a.	2 incl. p.a.	2 incl. p.a.
Performance statement	Charged separately	✓	✓	✓
Tax statement	Charged separately	✓	✓	✓
Pricing	0.60% - 0.30% p.a.	0.84% - 0.54% p.a.	1.08% - 0.78% p.a.	1.20% - 0.90% p.a.

Key shortcomings of initial offering and pricing landscape

- 1 Very retail-like client base with few and undifferentiated portfolio holdings migrated to Basic advisory solution
- 2 Insufficiently differentiated offering landscape. No focus on real value drivers as differentiation factors (i.e. access to experts [bankers vs. asset class specialists / investment consultants]; Degree of safety through portfolio monitoring, etc.)
- 3 Too high pricing for same service as former execution only proposition

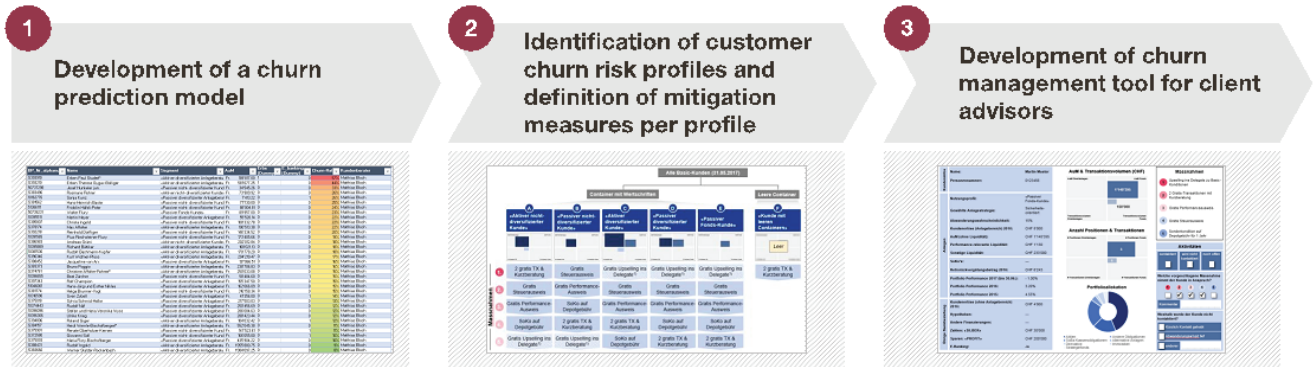
Client churn & downgrading specifically visible in the advisory service landscape



26% churn rate of Invest Basic clients within two years

- After two years, a large share of Invest Basic clients has either left the bank or has downgraded to an execution only solution
- 9% of clients who used Invest Basic in January 2015 had opted out to move into an execution only proposition
- More than 50% of new clients have left the bank within 2 years

TCG's approach to churn management

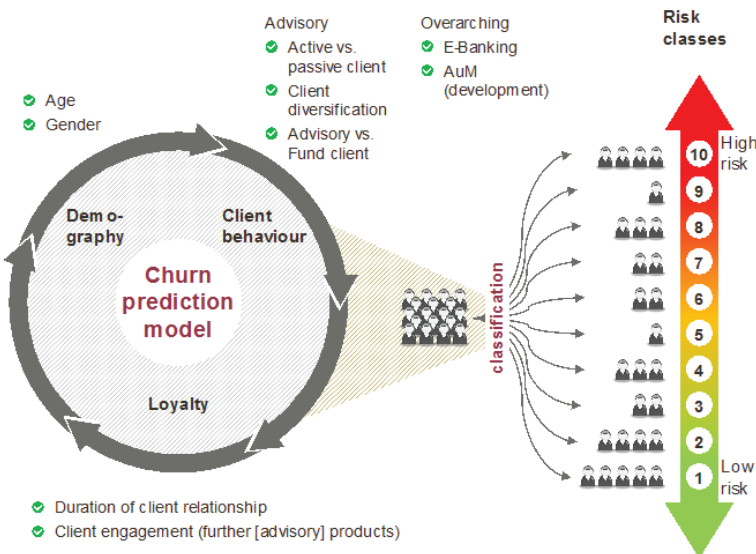


- Development of a regression model (churn prediction model) to identify risk of client attrition
- Factors identified were number of products used, number of lines in portfolio, type of investment (fund vs. single lines), etc.
- Validation of prediction factors with bankers and senior management
- Identification of customer churn risk profiles based on churn prediction model
- Definition of mitigation measures for all risk profiles
- Development of bespoke churn management tool for client advisors
- Tool aided client advisors to work through their list of target clients and adopt proposed measures
- 82% of all Basic clients were contacted, and churn rates decreased significantly to below 1% p.a.

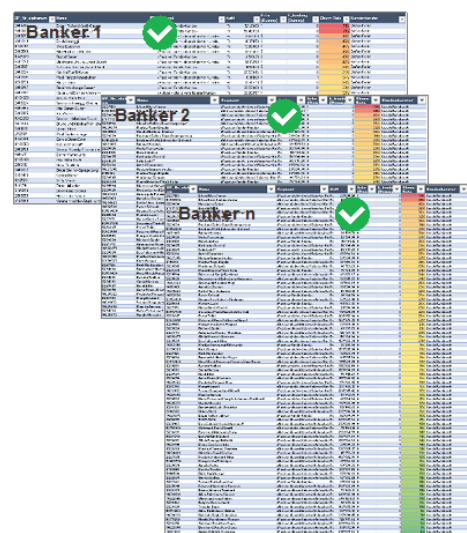
At TCG, we have developed a tool for client advisors to firstly identify customers at risk of leaving the bank and secondly to maintain them as clients through specific measures

A churn prediction model enabled risk classification of clients based on CRM data

Development of churn prediction model



Sounding with bankers




The churn probability of clients was measured based on a set of churn predictors. The results of the regression model were then verified with relationship managers and they were able to maintain them as clients through specific measures

A bespoke Banker Tool allowed for a more effective management of client relationships

The goal was to enhance client engagement with clear measures especially for high value clients in high risk classes.

Client info	
Name	Martin Muster
Client ID	0123456
Advisory	
Client behaviour profile	Passive funds client
Selected investment strategy	Risk-averse
Churn probability	15%
Client revenue (investment segment) 2016	£ 8,000
AuM (without cash)	£ 1,148,395
Performance-relevant liquidity	£ 1,130
Other liquidity	£ 233,000
Discounts	—
Retrocession refund amount 2016	£ 6,243
Portfolio performance 2017	-1.00%
Portfolio performance 2016	3.20%
Portfolio performance 2015	4.55%
Client relationship	
Client revenue in further segments 2016	£ 4,000
Mortgages	—
Further financing	—
Payments	£ 30,000
Savings	£ 200,000
E-Banking	Yes



Measures

- Up-Selling measure
- Service for free measures
- Special condition measure
- ..
- ..

Activities

Contacted	Will not be contacted	Still open
<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Which proposed measures will be applied?

-
-
-
-
-

Comments:

Why was the client not contacted?

- Only just recently contacted
- Low churn risk
- Other:

A longlist of measures has been developed, reduced to a shortlist and specified for each client profile

Author



Amitabha Sinha
VP - Delivery and Analytics,
India and SAARC

Amitabha has over 20 years of experience in technology and strategic consulting. Both on the technology as well as the digital strategy side, he has run transformation projects with Fortune 100 companies in geographies like US, Canada, UK, and Singapore. Amitabha started his career with IBM Watson Lab and has 5 patents with IBM. Thereafter, he worked for consulting majors like PwC (UK) and EY (US) and has run large BFSI ODCs based out of Singapore and Hong Kong.

Amitabha is a First Class First President's Gold Medal Winner from IIT Delhi, has an MS from Illinois Tech and a PGP from IIM Bangalore.