Society of Actuaries in Ireland

Predictive Modeling for Customer Targeting – A Banking Example

23rd February 2017

Disclaimer

The views expressed in this presentation are those of the presenter and not necessarily of the Society of Actuaries in Ireland

Agenda

- What is customer targeting?
- Why bother?
- How do you do it?

What Is Customer Targeting?

"*The identification of individuals interested in a product*".

Not exactly, but you get the idea.

From a modelling point of view:

- Binary classification (buyer vs. non-buyer)
- Probability of buying
- Purchase propensity score

Some Examples

Some Examples - Amazon

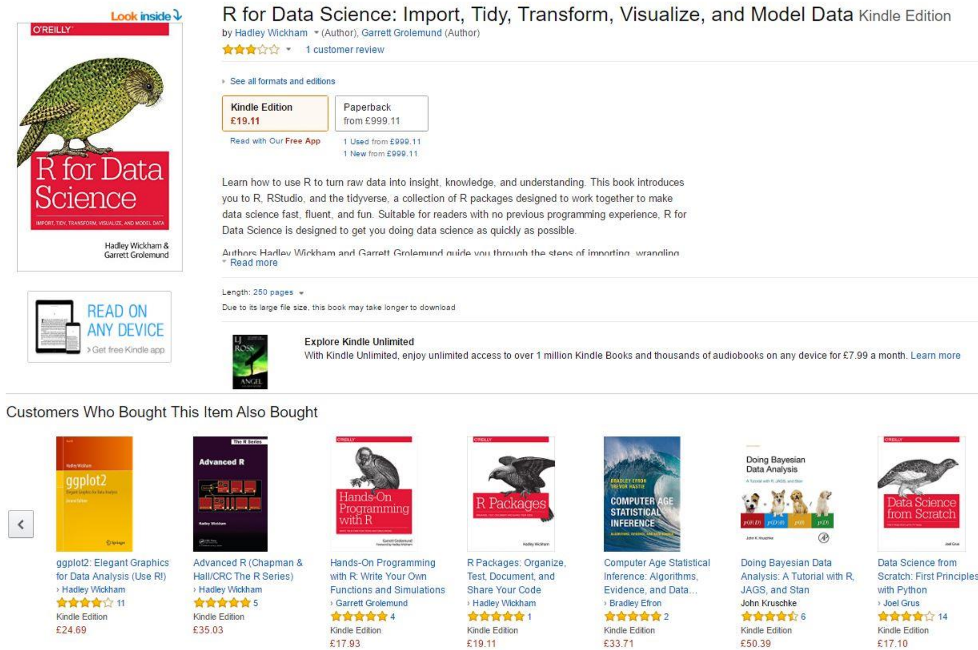

- *"Customers Who Bought This Item Also Bought"*
- Is this customer targeting? *"Is Pedro likely to purchase these products?"*
- Or is this product targeting? *"What other products is Pedro likely to purchase?"*

Some Examples – Netflix

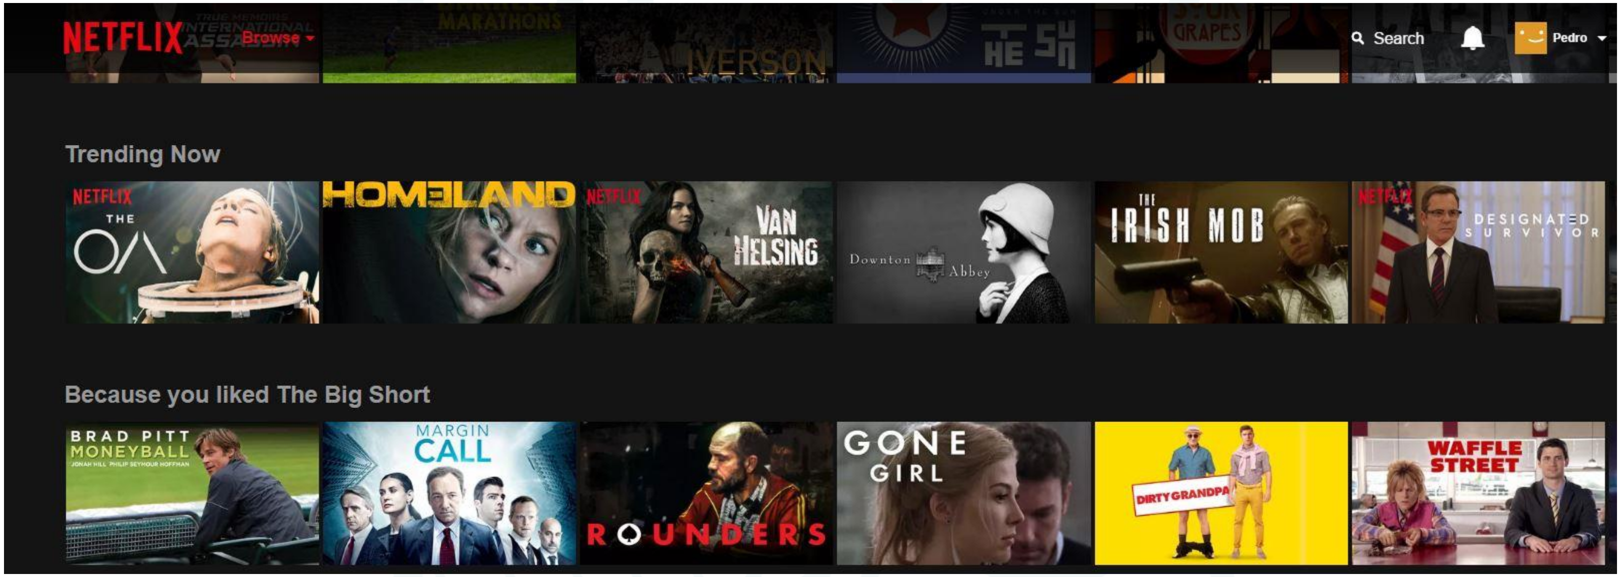

- “*Trending Now*” – Targeting based on popularity (majority, mode, top quartile, etc.)
- “*Because you liked The Big Short*” – Based on individual preferences (same director, actors, etc.)

Some Examples – Netflix

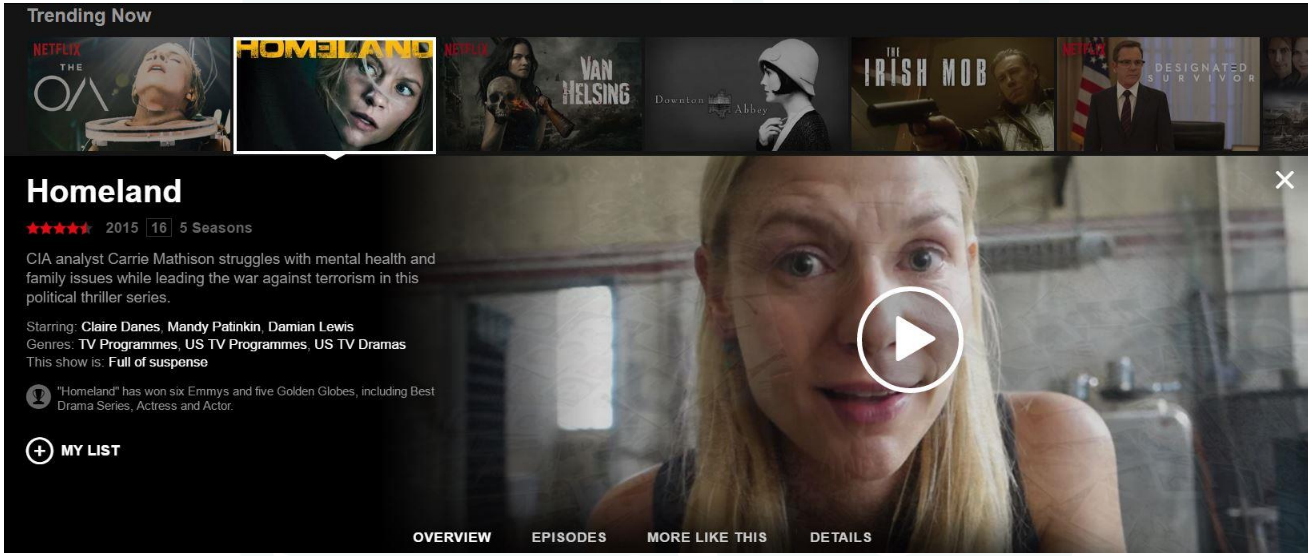

- “*Trending Now*” but note the estimated rating: 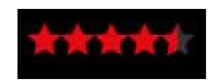 This is a propensity score. “*With how many stars would Pedro rate this very popular show?*”
- In fact, Netflix usually recommends movies (even in “Trending Now”) that you are expected to like **very much**.

Why Bother?

Why Bother? – Cross-Selling amazon

- On-site Recommendations*: "Recommended for you, Pedro"*
- *"Frequently bought together"*
- *"Your recently viewed items and featured recommendations – inspired by your browsing history"*
- *"Related to items you've viewed"*
- *"Customers who bought this item also bought"*
- *"There is a newer version of this item"*
- *"Recommended for you based on a previous purchase"*

And more on-site and e-mail based targeting.

They are jointly responsible for **35% of Amazon's sales**.

Why Bother? – Product Relevance

21st September 2009: Netflix awards a $1,000,000 prize to "*BellKor's Pragmatic Chaos*" after three years of ongoing competition.

The winning model provided significantly more accurate estimated ratings for Netflix's recommender system.

Why Bother? – Cross-Selling

TheActuary

The magazine of the actuarial profession

LIFE PENSIONS GI INVESTMENT RISK & ERM HEALTH REINSURANCE ENVIRONMENT REGULATION M

HOME NEWS FEATURES INTERVIEWS OPINION INTERNATIONAL EVENTS

The digital revolutionary

Andrew Brem, chief digital officer of Aviva, talks with Richard Purcell and Stephen Hyams about the challenges of digital transformation to create a better customer experience

01 DECEMBER 2016 | RICHARD PURCELL & STEPHEN HYAMS

Andrew Brem

"***Our industry has been quite traditional in terms of marketing and cross-selling, but consumers now expect you to pop up at the most relevant point and in the most relevant way to them, which is not necessarily on your own website***." He clarifies that this targeted approach to customer engagement uses customer analytics and internal or external data.

http://www.theactuary.com/features/2016/12/the-digital-revolutionary/

How Do You Do It?

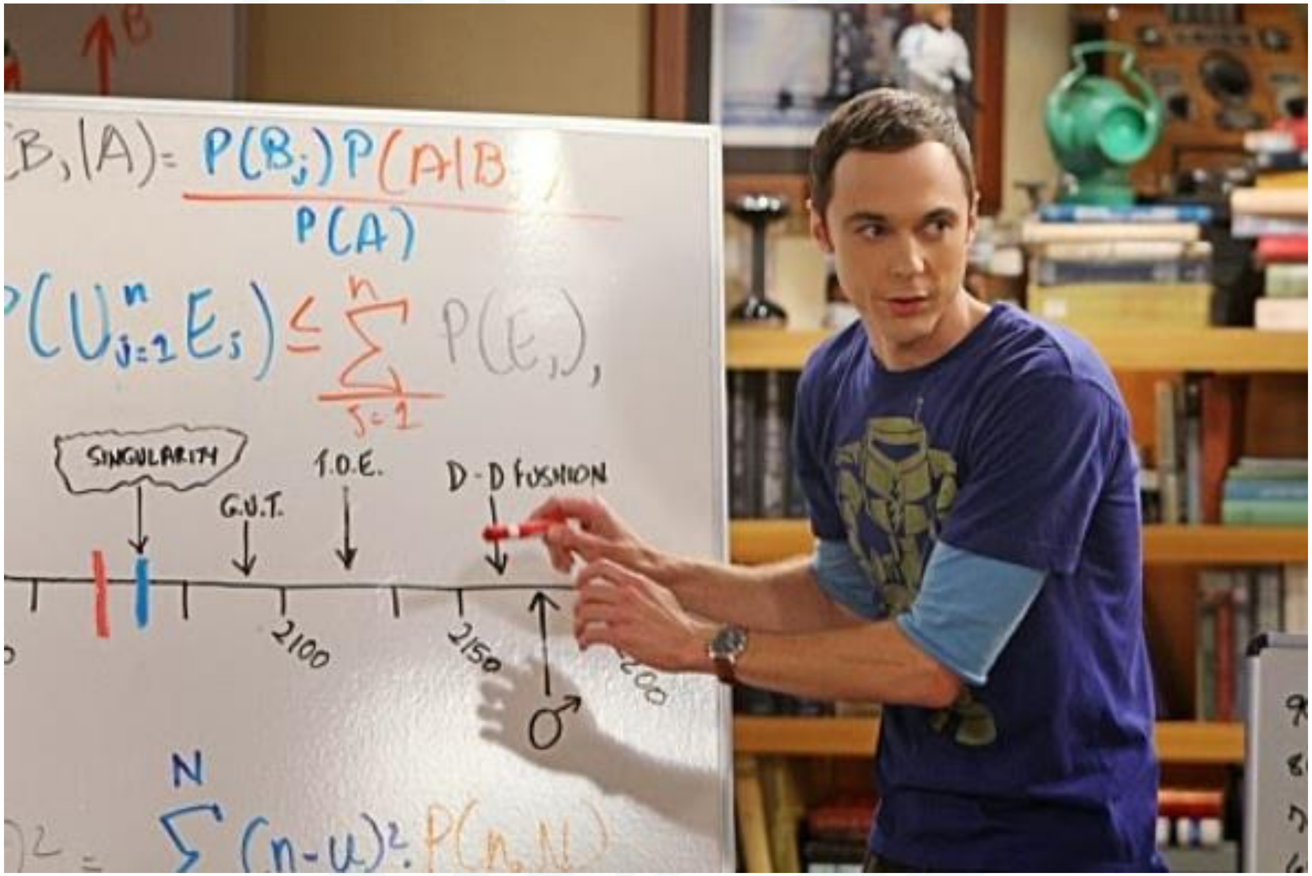

How Do You Do It? – Banking Telemarketing

- Acknowledgement: *[Moro et al., 2014] S. Moro, P. Cortez and P. Rita. A Data-Driven Approach to Predict the Success of Bank Telemarketing.* They made the database public.
- Lichman, M. (2013). *UCI Machine Learning Repository [http://archive.ics.uci.edu/ml]. Irvine, CA: University of California, School of Information and Computer Science*.

How Do You Do It? – Banking Telemarketing

41,188 calls were previously made.

11% customers (4,530) bought the product.

Can we do better or do we need to call everyone again the next time?

Class imbalance.

Training dataset (2/3)

Testing dataset (1/3)

How Do You Do It? – Banking Telemarketing

Under Sampling

Over Sampling

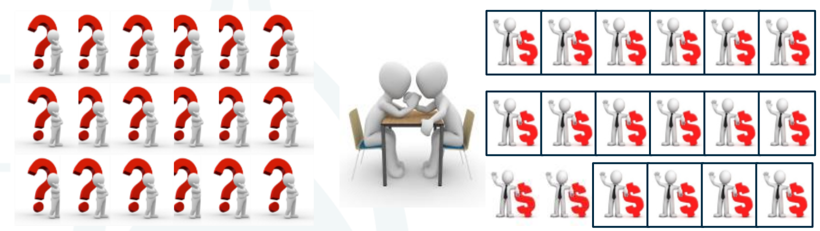

Balanced Sampling

Synthetic Data

We could also do nothing…

How Do You Do It? – Banking Telemarketing

What model(s) should we use?

I will focus on **Naïve Bayes** and **Support Vector Machines** but there are many more that could be useful:

•Logistic Regression
•Decision Trees
•Random Forest
•XGBoosting
•K - Nearest Neighbours
•And more…

And what about class imbalance, do we try a particular approach or do we not do anything at all? I had time so I tried them all.

How Do You Do It? – Banking Telemarketing

Naïve Bayes

$$\hat{y} = \underset{k \in \{1,\dots,K\}}{\operatorname{argmax}} \; p(C_k) \prod_{i=1}^{n} p(x_i | C_k).$$

It uses conditional probabilities, based on Bayes theorem, allocating an observation to its most probable class.

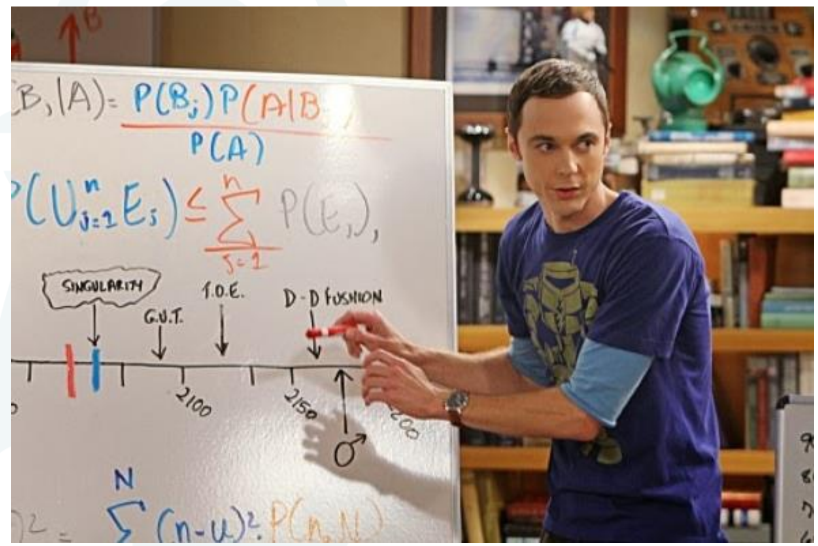

It assumes variables are normally distributed and not correlated, which is rarely true. However the classifier can be very effective even when assumptions are not met.

How Do You Do It? – Banking Telemarketing

Support Vector Machines (SVM)

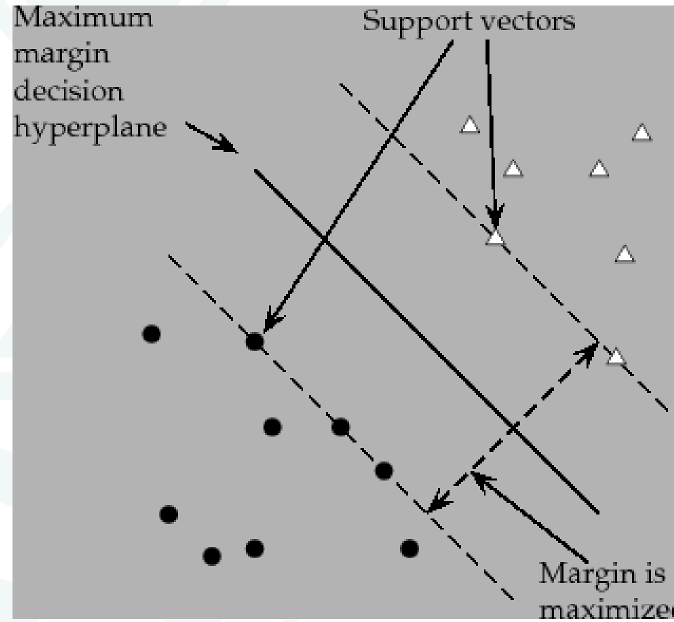

It constructs hyperplanes in a multidimensional space that separates cases of different class labels.

To find an optimal hyperplane, SVM uses an iterative training algorithm used to minimise an error function.

The hyper plane does not need to be a straight line. The kernel trick allows for non-linear classification. Possible kernels are:

•Linear
•Polynomial
•Radial
•Sigmoid

How Do You Do It? – Banking Telemarketing

This chart shows the proportion of customers contacted vs. the proportion of buyers correctly identified by the model. **Results based on the testing**

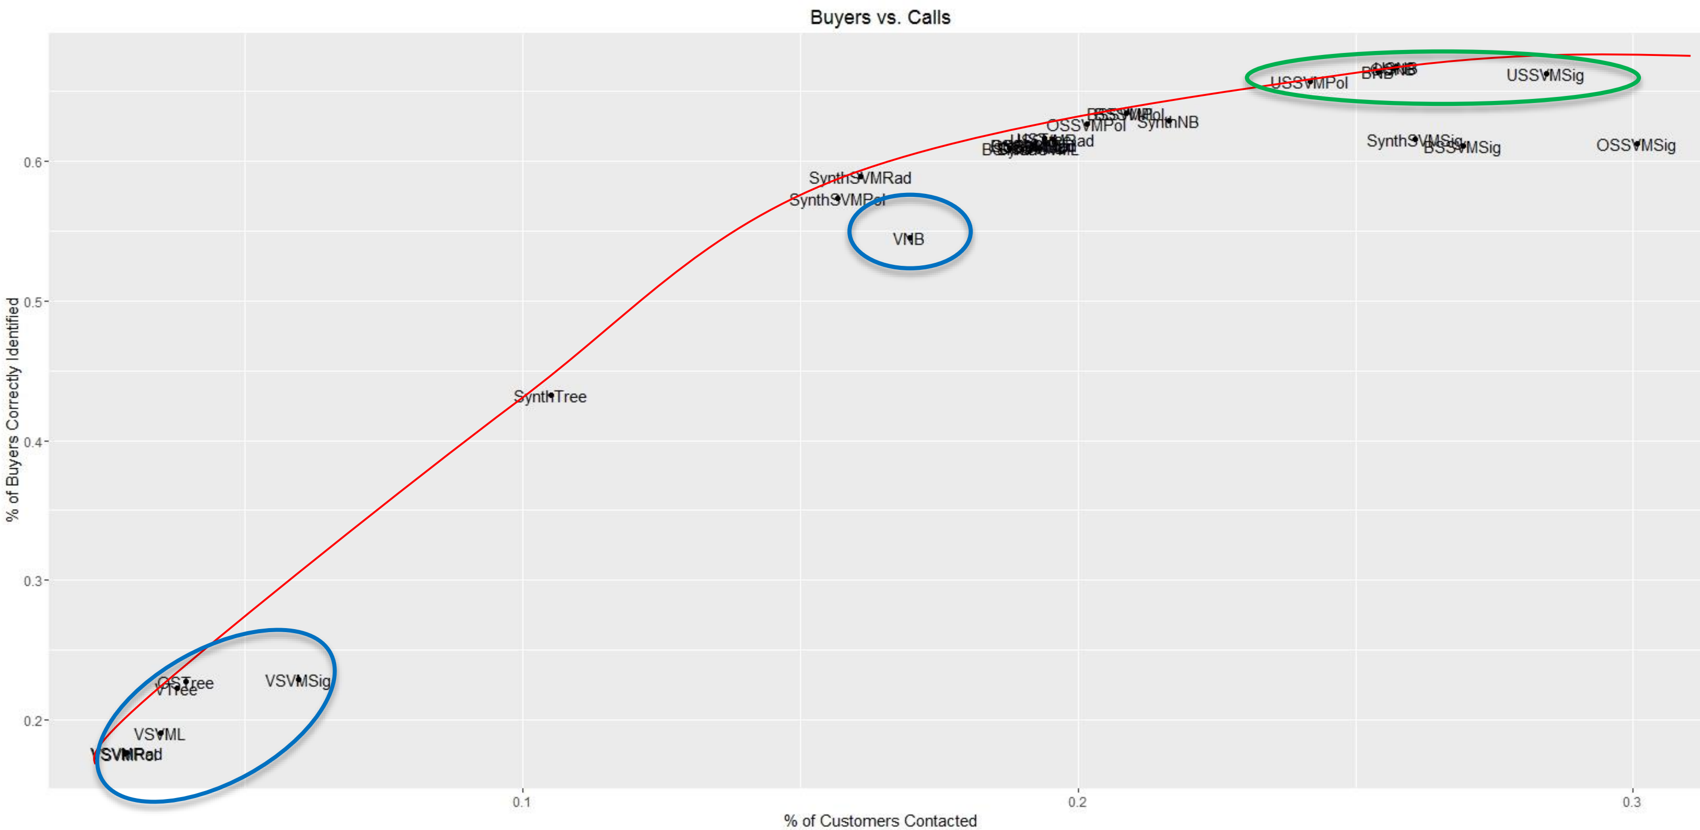

Blue circles show results with no action for class imbalance.
The **Green** circle show results with under sampling for class imbalance.

How Do You Do It? – Banking Telemarketing

Results for Naïve Bayes.

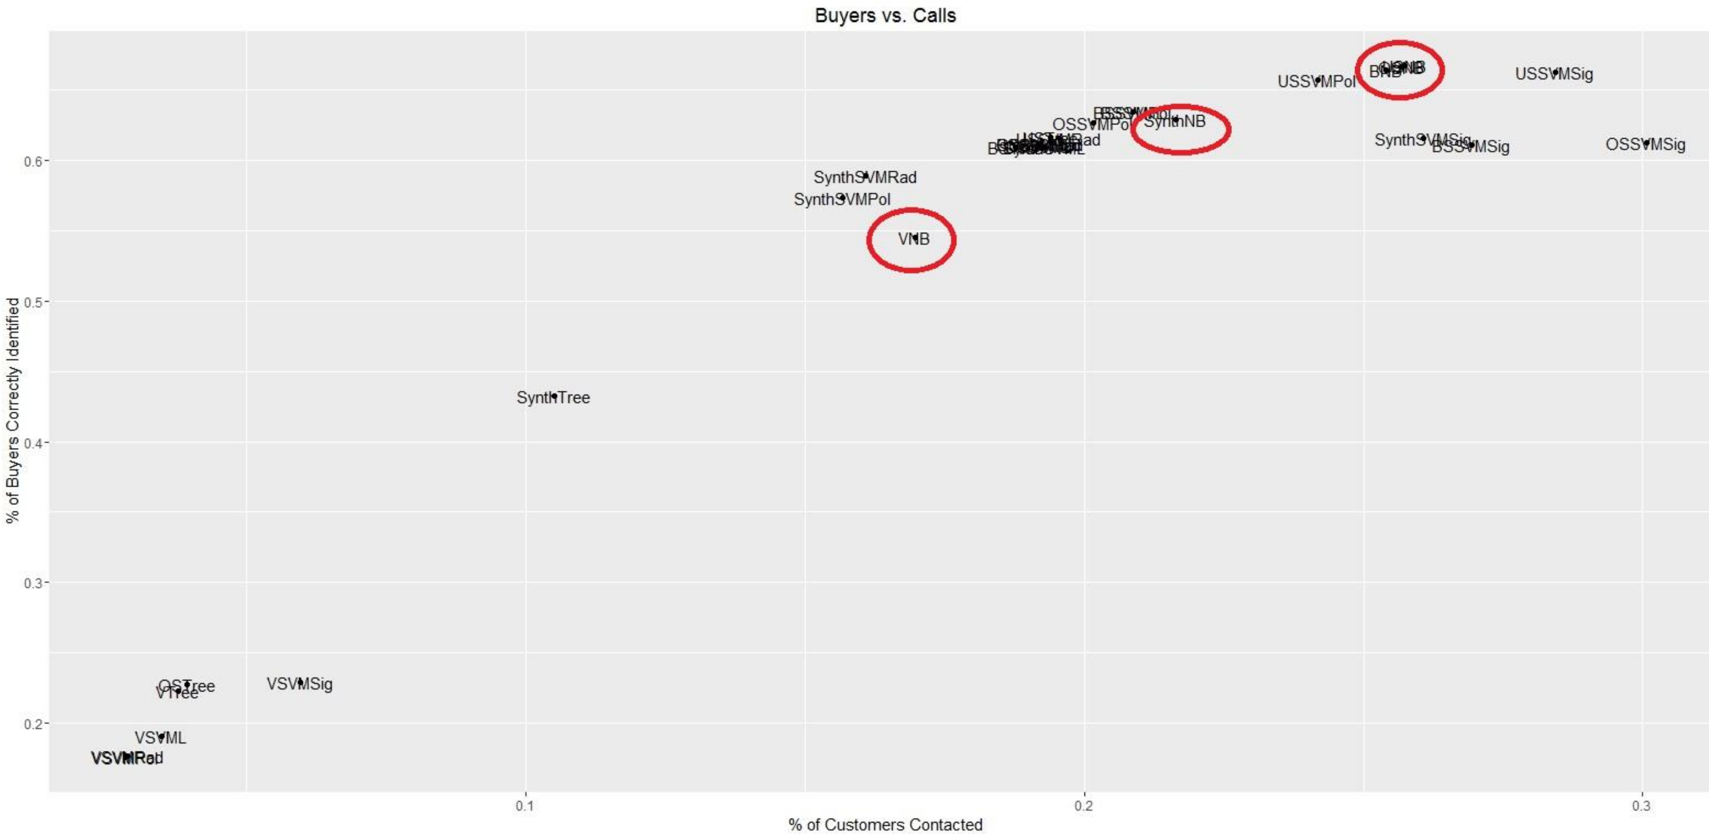

How Do You Do It? – Banking Telemarketing

Results for SVM Sigmoid.

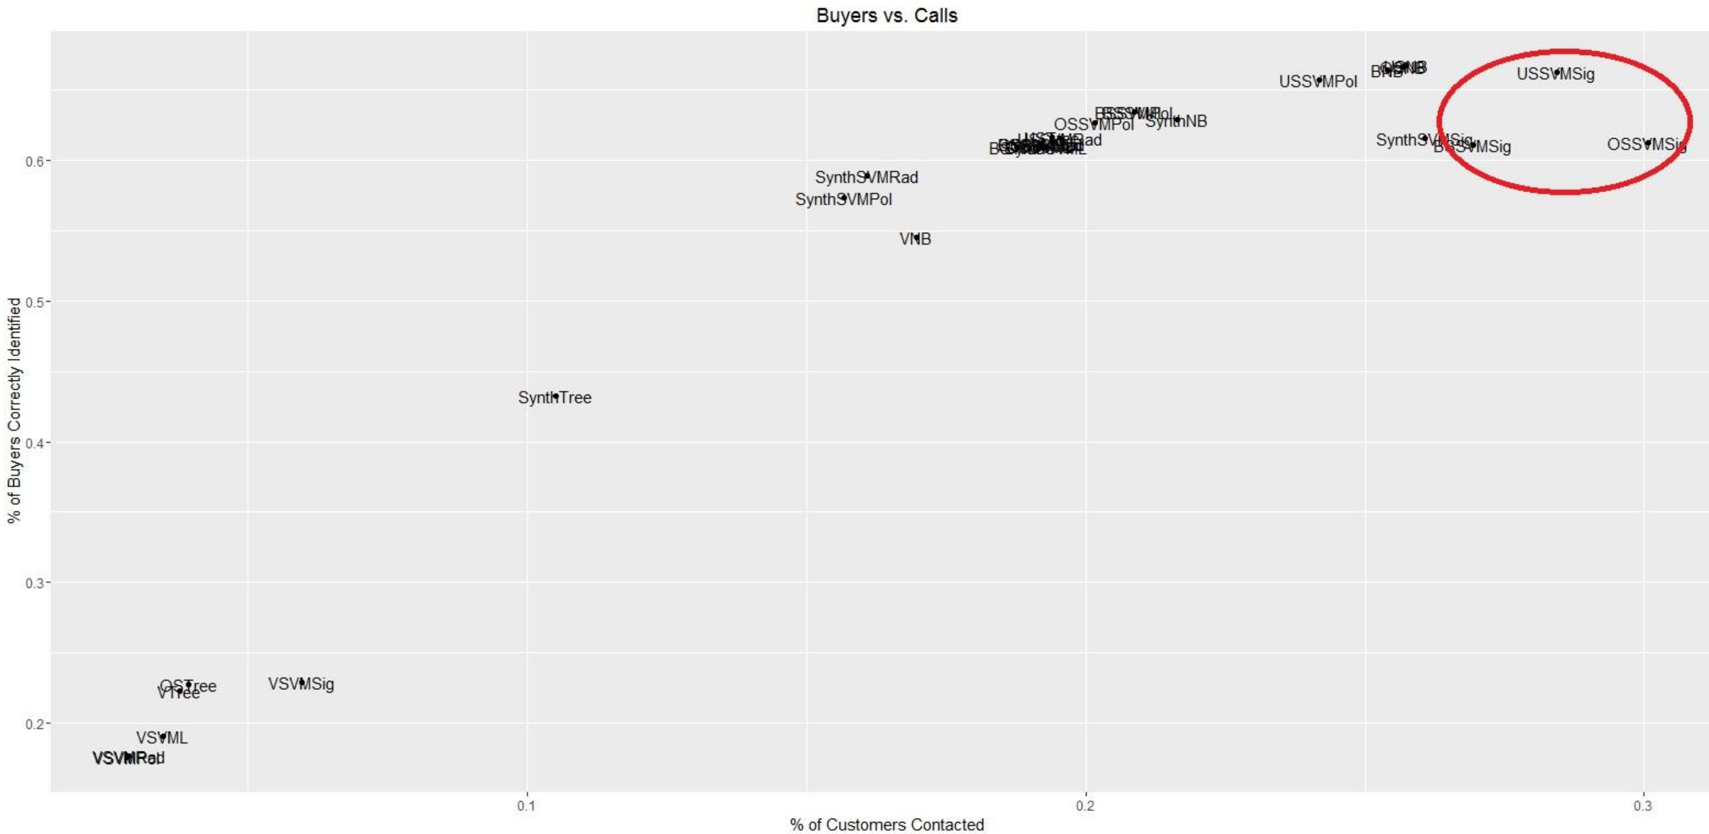

How Do You Do It? – Banking Telemarketing

Naïve Bayes provided better results but SVM gave good predictions too.

Could a combination of the two result in a better model?

•Ensemble 1: everyone is a buyer unless a majority of USNB, USSVMSig and USSVMPol votes against.

•Ensemble 2: everyone is a buyer unless USNB and USSVMSig agree to the opposite.

How Do You Do It? – Banking Telemarketing

Ensemble 2 can predict more buyers than Ensemble 1.

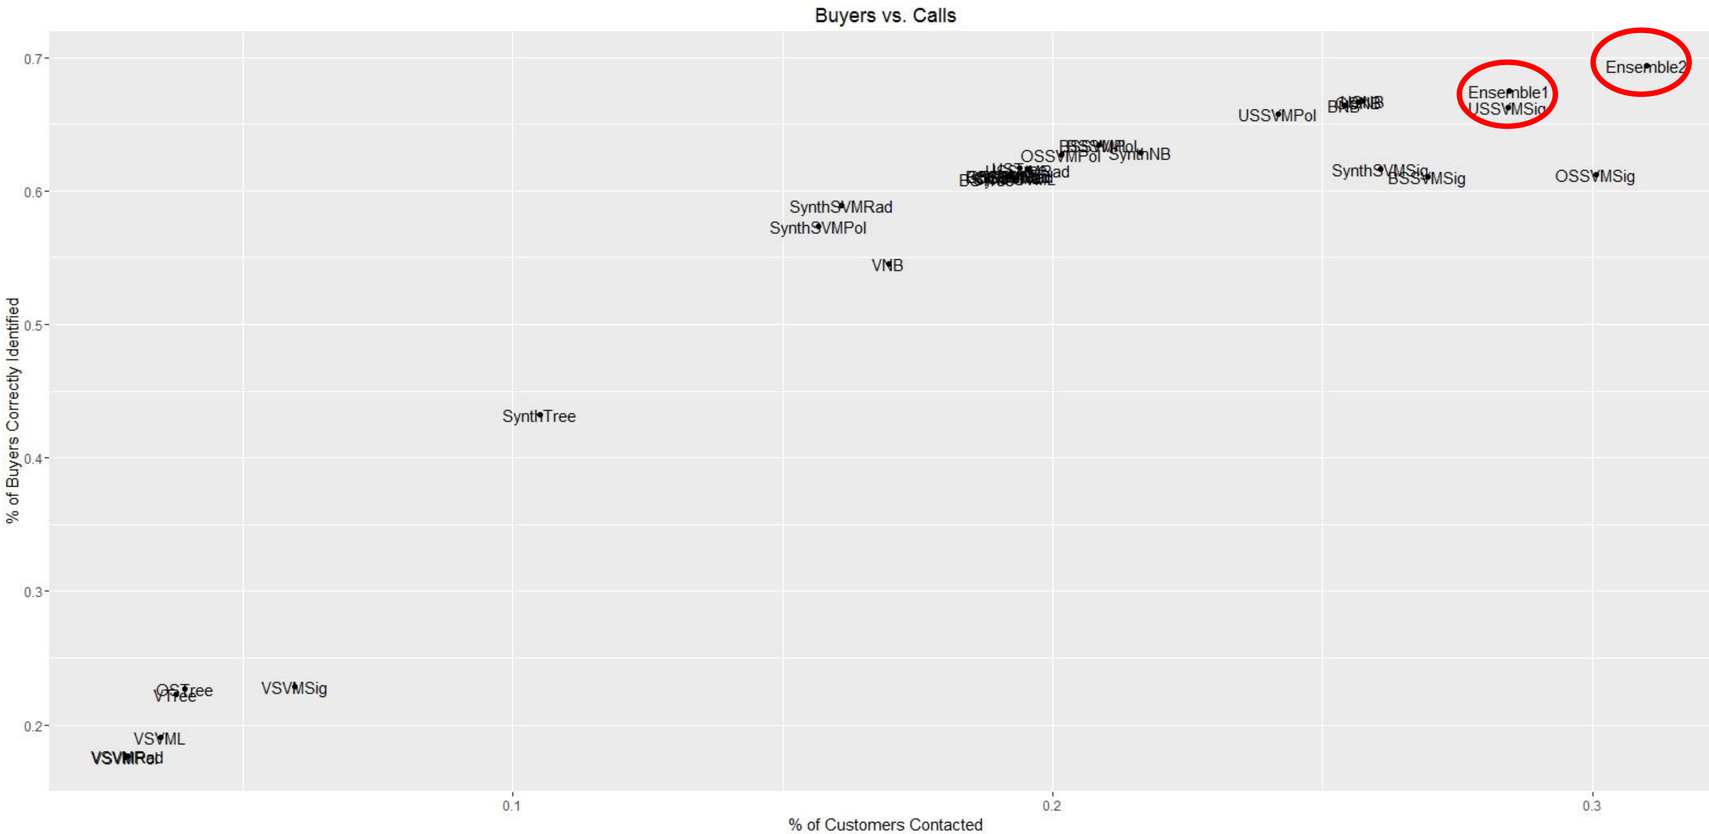

How Do You Do It? – Banking Telemarketing

So what? Putting results into context with the testing dataset:

- **13,729 customers**
- **1,510 buyers**

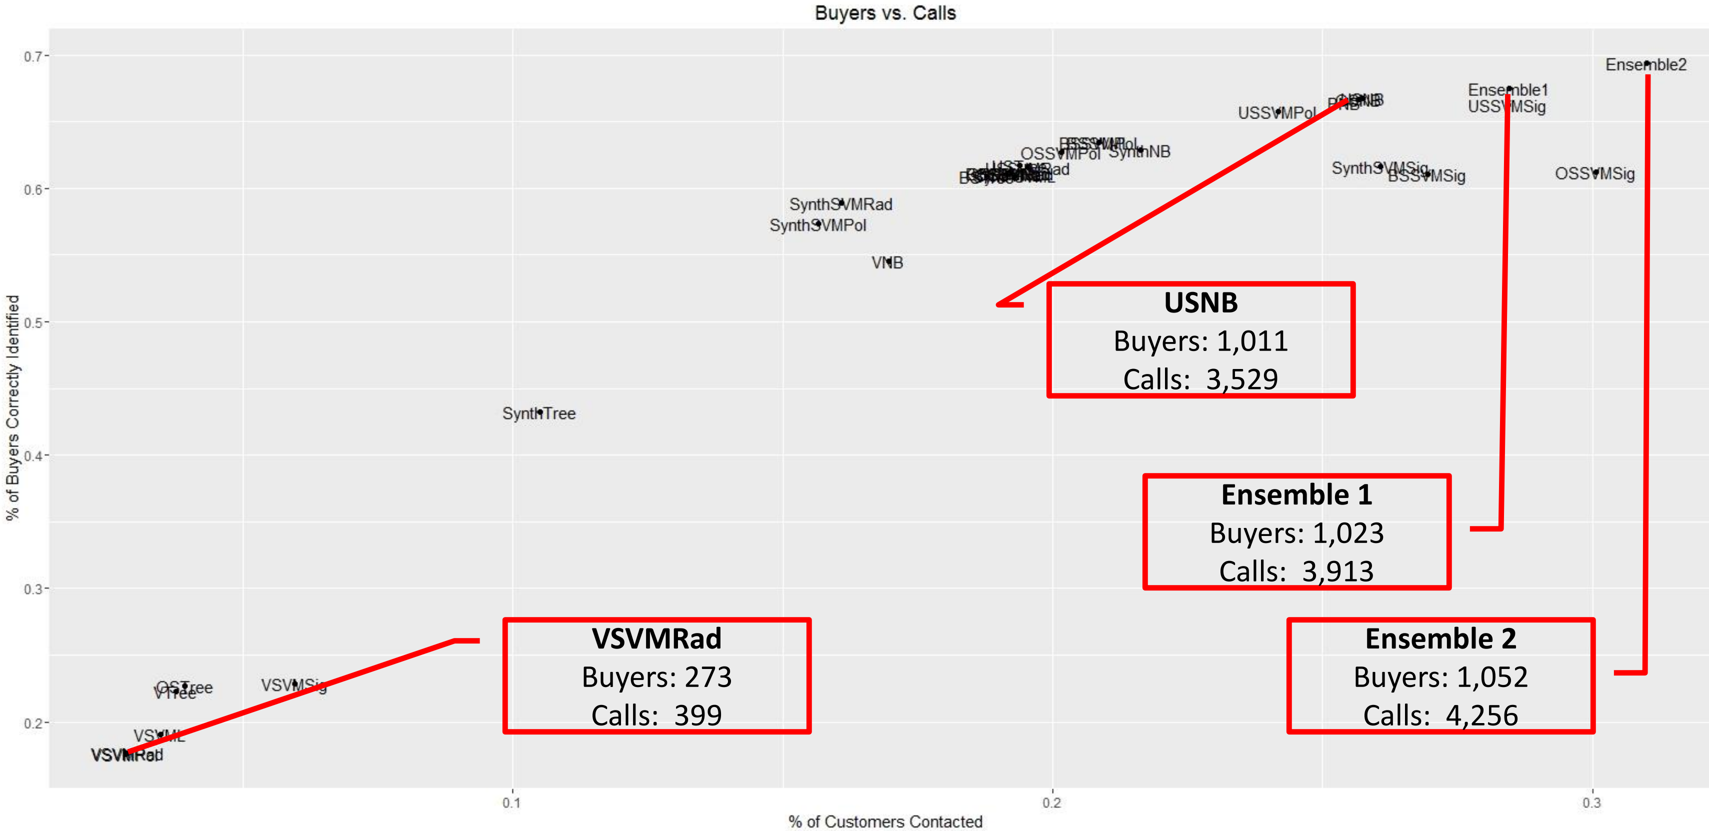

How Do You Do It? – Banking Telemarketing

SVM Rad with no action for class imbalance was the most accurate model with a 68% success rate. However, it would identify less than 20% of buyers. That is **20% of the buyers with 3% of the original calls**.

Ensemble 2 would have a 25% success rate but correctly identify 70% of buyers.
It would identify **70% of the buyers with 30% of the original calls**.

Method	Buyers	Calls
VSVMRad	801	1,195
Ensemble 2	3,143	12,769
Calling everyone	4,530	41,188

A cost-benefit matrix would help optimise the results.

How Do You Do It? – Banking Telemarketing

After all that, what can we say about buyers?

SVM will not give us anything so we focus on the conditional probabilities from Naïve Bayes and an analysis of the dataset. Just a few examples.

Image source: Pixabay / Gerd Altmann

How Do You Do It? – Banking Telemarketing

Age – Younger and older customers are more likely to buy.

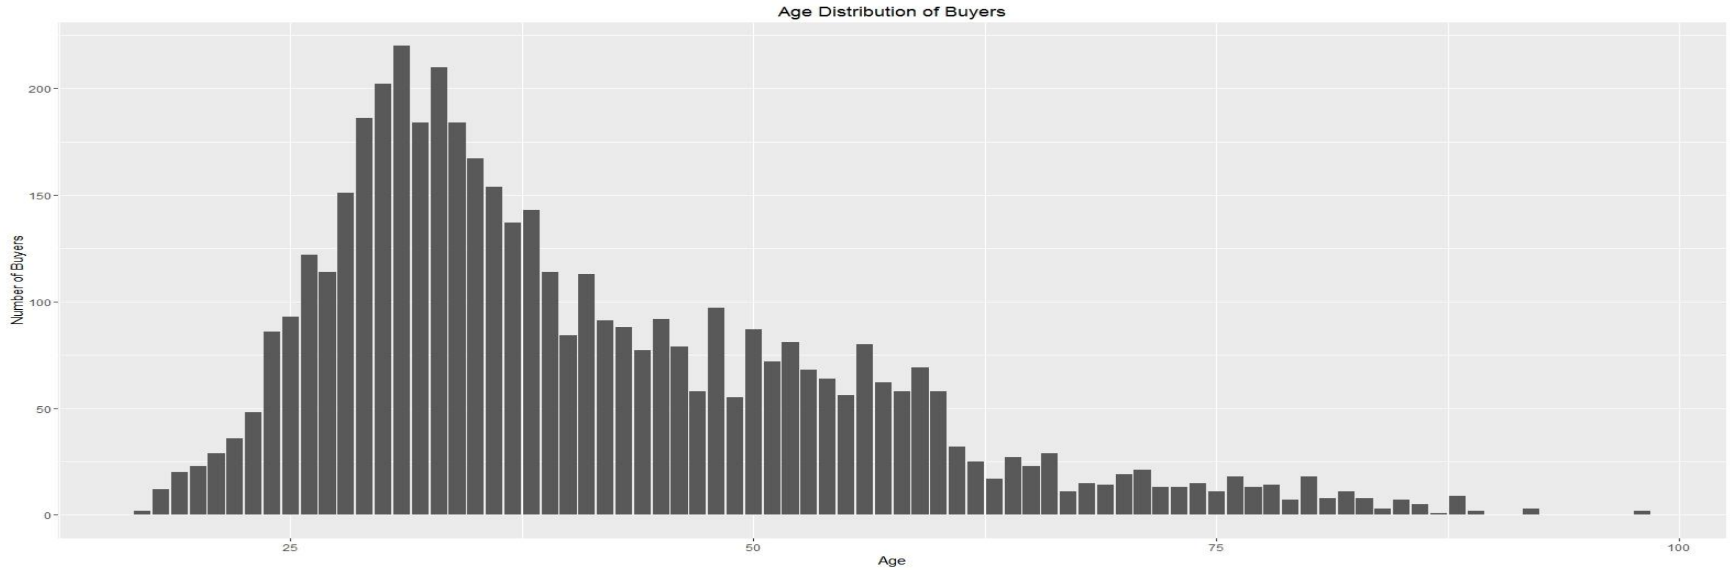

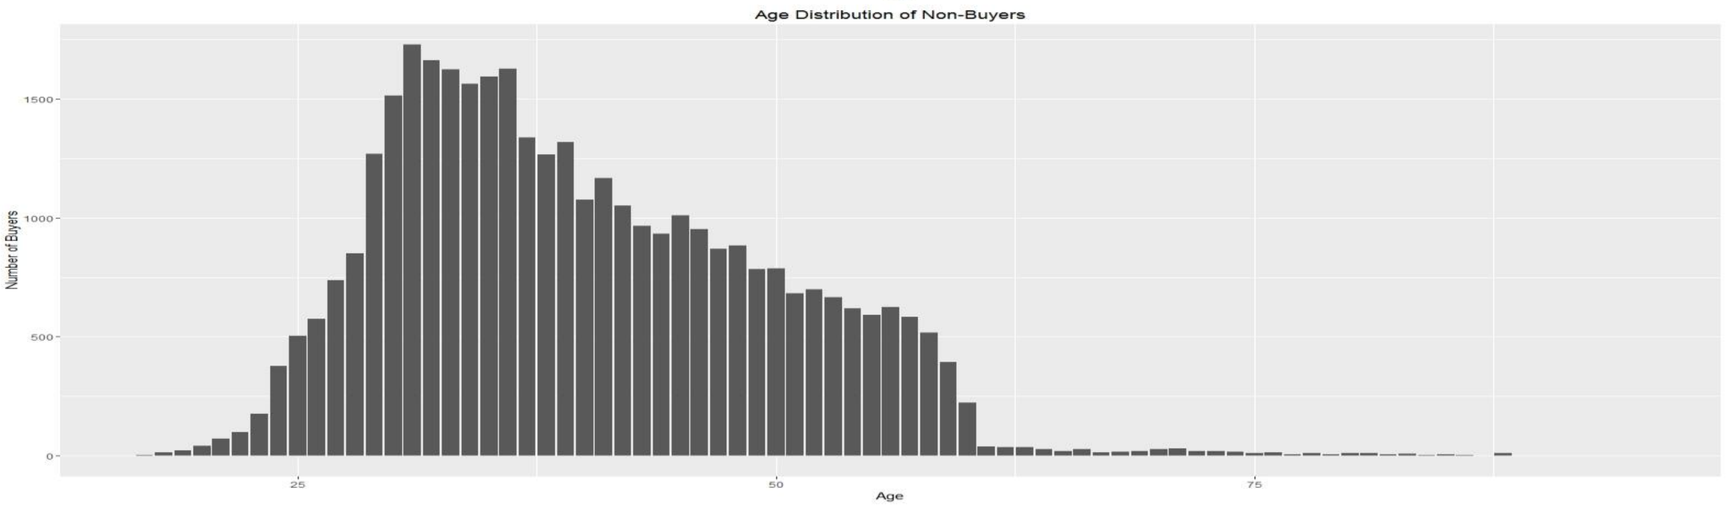

How Do You Do It? – Banking Telemarketing

Number of previous contacts – The more, the merrier.

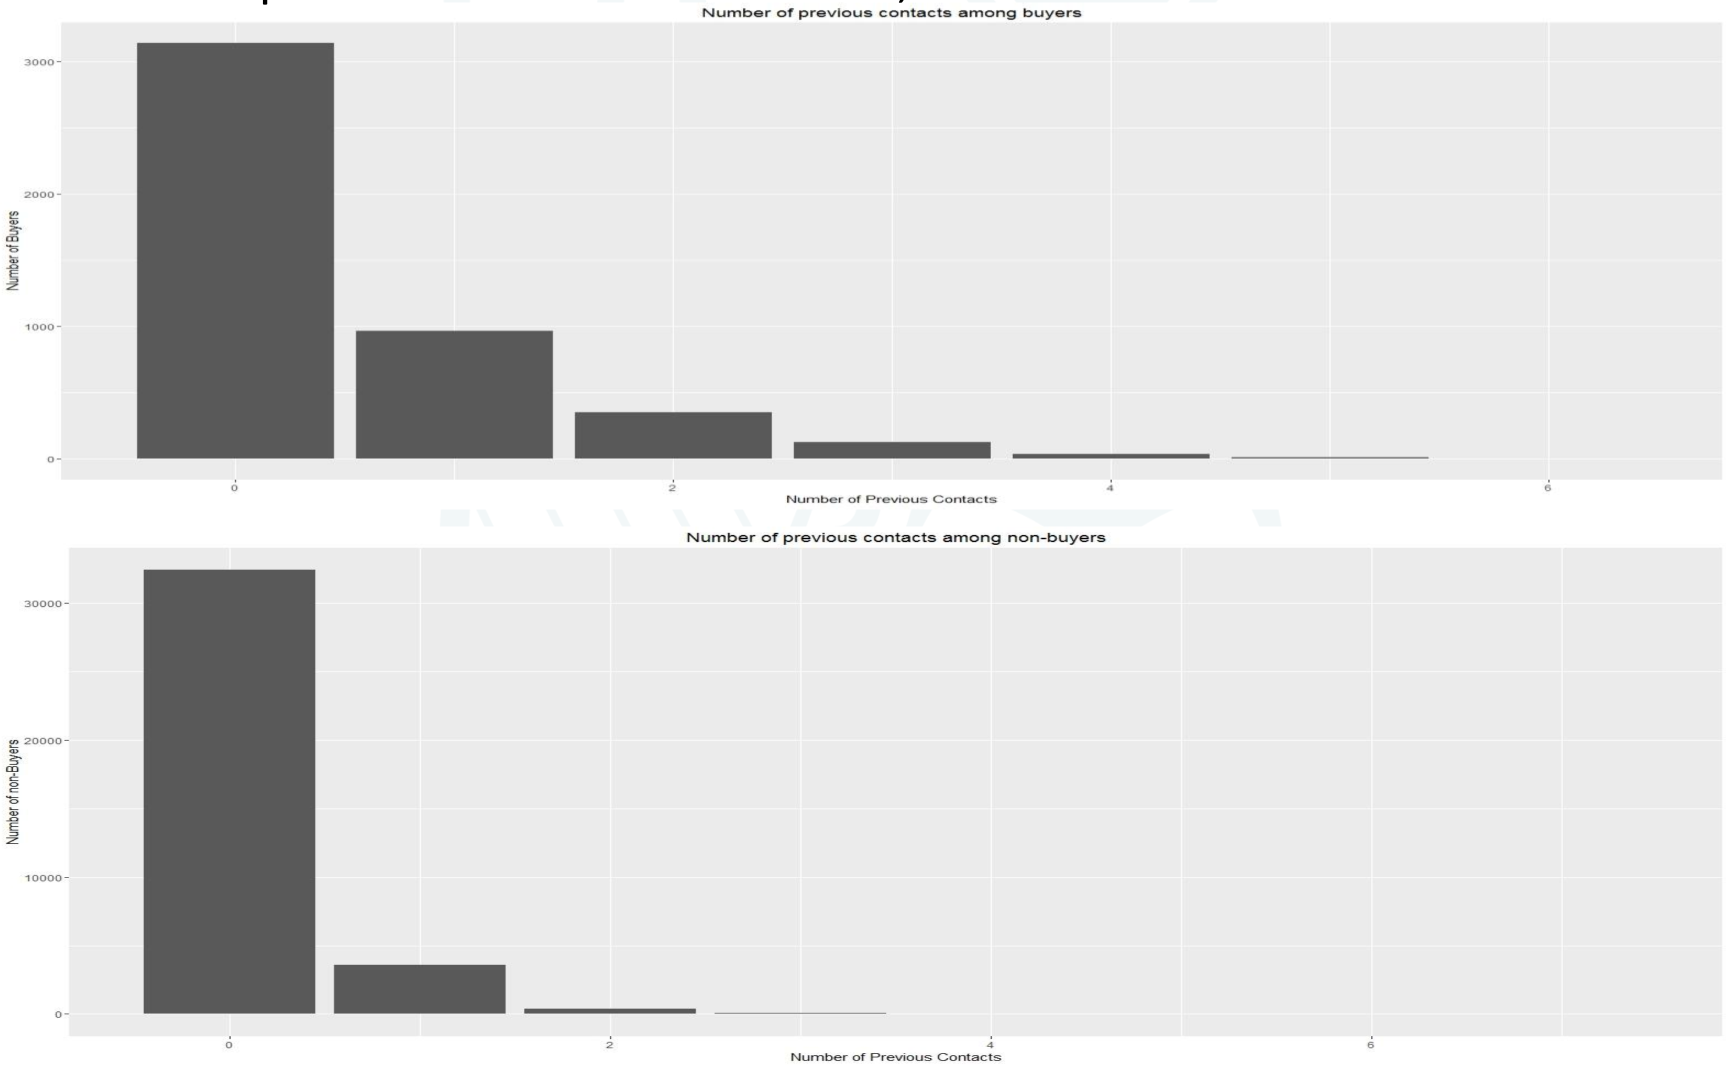

How Do You Do It? – Banking Telemarketing

Previous outcome – Past buyers make good future buyers.

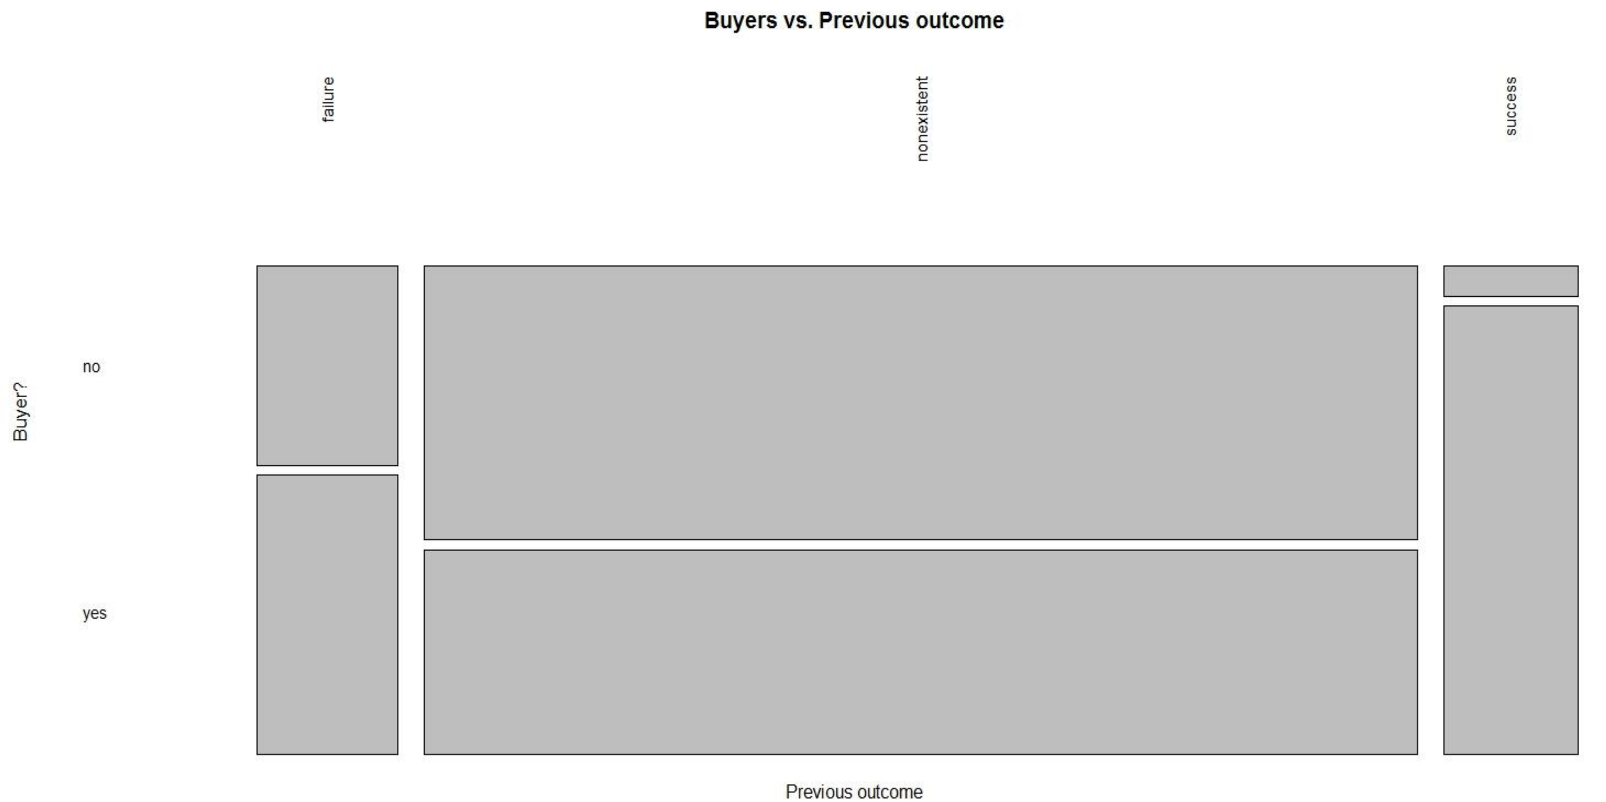

How Do You Do It? – Banking Telemarketing

Employment variation rate – Keener to buy when employment falls.

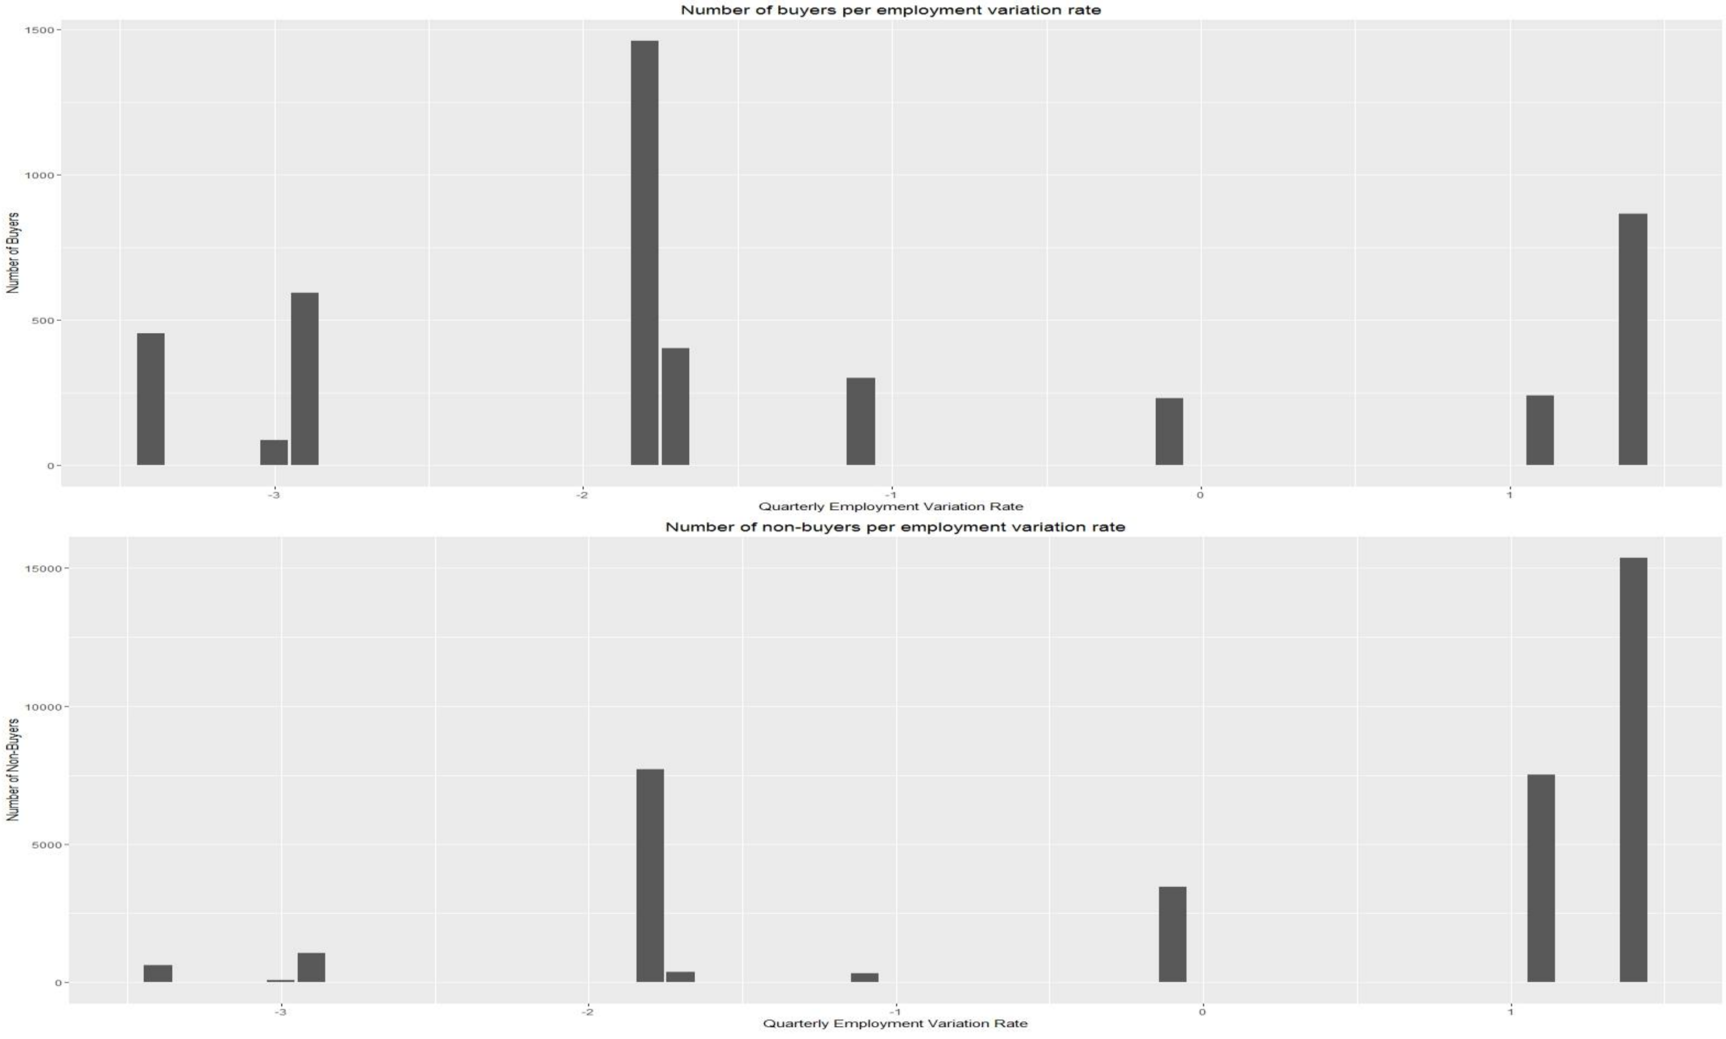

How Do You Do It? – Banking Telemarketing

What did we learn?

Timing is a very important factor: customers are more inclined to buy the product when economic conditions deteriorate (increasing unemployment and falling interest rates) **In fact, without these variables, the model struggles to provide reliable predictions.**

The prototype buyer is an existing customer, well educated, mobile phone user, with no family responsibilities (student or retired) and who has not defaulted on a loan before.

Summary

- Amazon and Netflix.
- Customer targeting adds value to the company, enhances the customer experience and provides actionable information about customers.
- Modelling approaches vary and there are many options. They all have strengths and weaknesses and a combination of models might provide better results than a single one.

Thank You