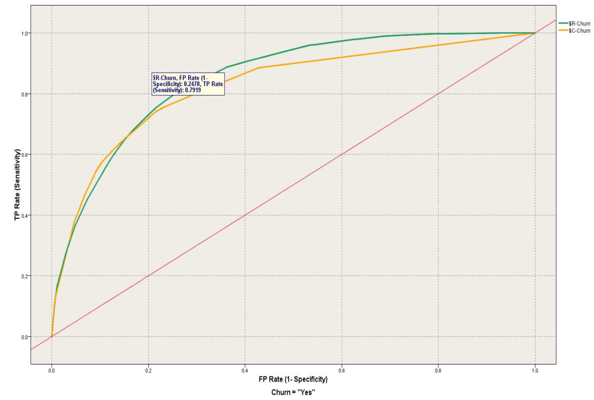

Evaluating and deploying predictive models

Jarlath Quinn – Analytics Consultant

FAQ's

- Is this session being recorded? Yes
- Can I get a copy of the slides? Yes, we'll email a PDF copy to you after the session has ended.
- Can we arrange a re-run for colleagues? Yes, just ask us.
- How can I ask questions? All lines are muted so please use the chat facility – if we run out of time we will follow up with you.

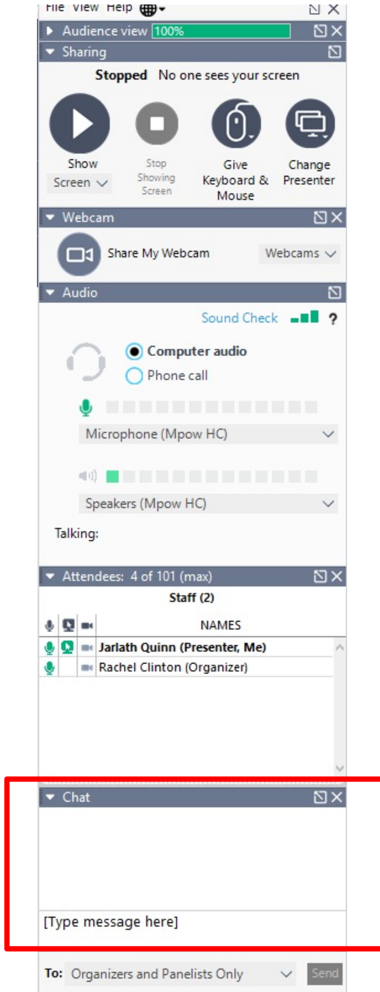

- Premier accredited partner to IBM and Predictive Solutions specialising in advanced analytics & big data technologies
- Work with open source technologies (R, Python, Spark etc.)
- Team each has 15 to 30 years of experience working in the advanced and predictive analytics industry

- Deep experience of applied advanced analytics applications across sectors
  - Retail
  - Gaming
  - Utilities
  - Insurance
  - Telecommunications
  - Media
  - FMCG

Agenda

- What does ‘good’ look like?
- Comparing measures of predictive accuracy
- Using charts to evaluate model performance
- Testing model performance
- Selecting models based on profitability
- Deploying models

What does ‘good’ look like?

What does all this stuff have in common?

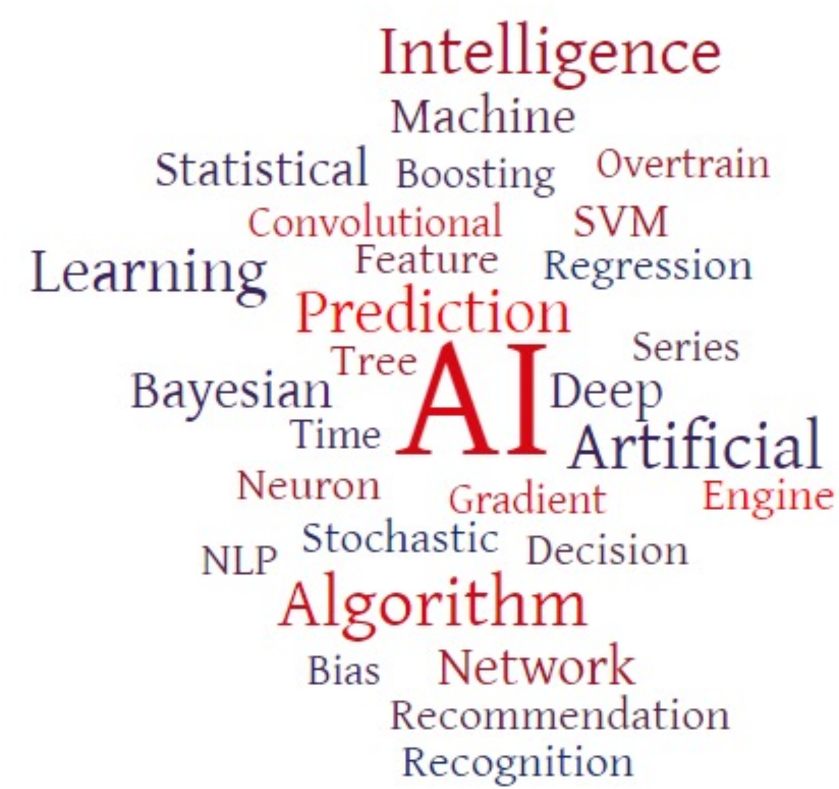

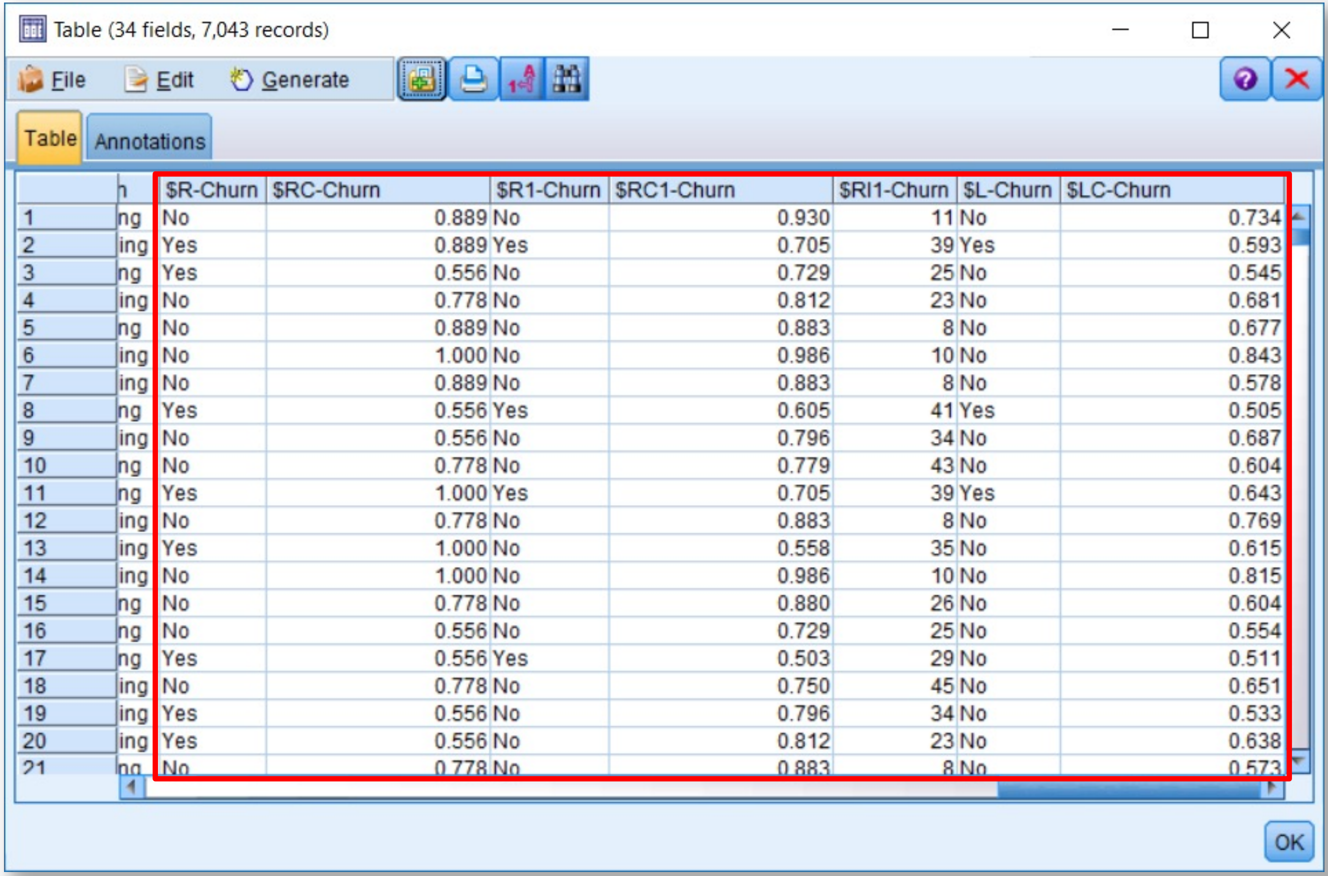

	h	$R-Churn	$RC-Churn	$R1-Churn	$RC1-Churn	$RI1-Churn	$L-Churn	$LC-Churn
1	ng	No	0.889	No	0.930	11	No	0.734
2	ing	Yes	0.889	Yes	0.705	39	Yes	0.593
3	ng	Yes	0.556	No	0.729	25	No	0.545
4	ing	No	0.778	No	0.812	23	No	0.681
5	ng	No	0.889	No	0.883	8	No	0.677
6	ing	No	1.000	No	0.986	10	No	0.843
7	ing	No	0.889	No	0.883	8	No	0.578
8	ng	Yes	0.556	Yes	0.605	41	Yes	0.505
9	ing	No	0.556	No	0.796	34	No	0.687
10	ng	No	0.778	No	0.779	43	No	0.604
11	ng	Yes	1.000	Yes	0.705	39	Yes	0.643
12	ing	No	0.778	No	0.883	8	No	0.769
13	ing	Yes	1.000	No	0.558	35	No	0.615
14	ing	No	1.000	No	0.986	10	No	0.815
15	ng	No	0.778	No	0.880	26	No	0.604
16	ng	No	0.556	No	0.729	25	No	0.554
17	ng	Yes	0.556	Yes	0.503	29	No	0.511
18	ing	No	0.778	No	0.750	45	No	0.651
19	ing	Yes	0.556	No	0.796	34	No	0.533
20	ing	Yes	0.556	No	0.812	23	No	0.638
21	ng	No	0.778	No	0.883	8	No	0.573

They generate new data

The output from predictive models can offer…

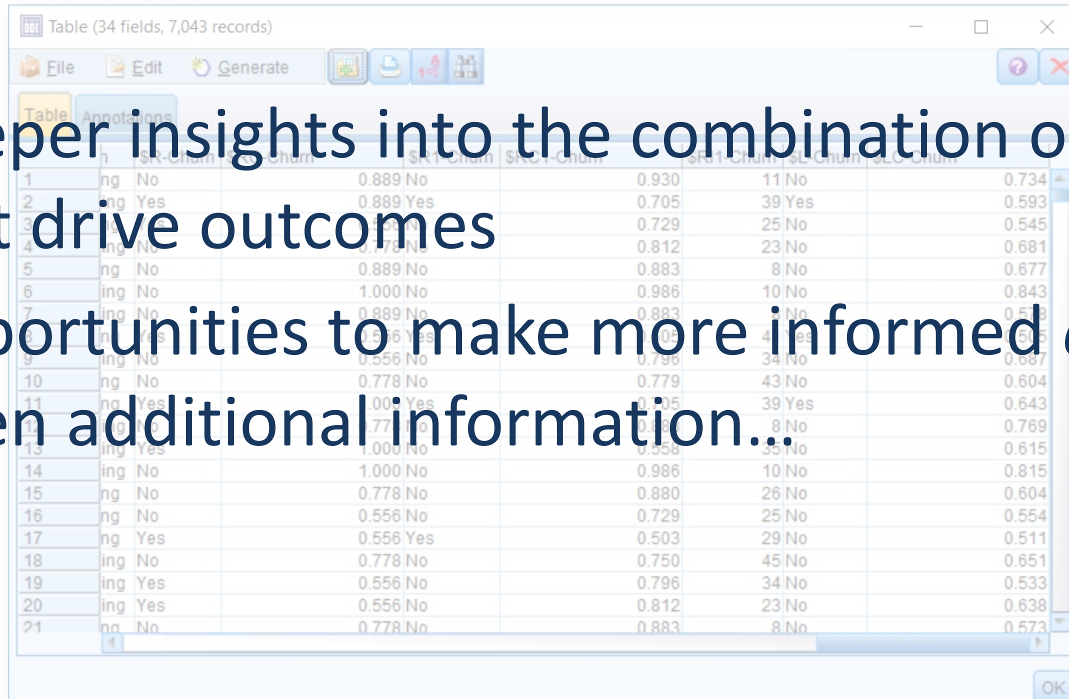

- Deeper insights into the combination of factors that drive outcomes
- Opportunities to make more informed *decisions* given additional information…

Insight vs Action

However, data analysts are often guilty of failing to acknowledge that a predictive model is not the same thing as a useful application

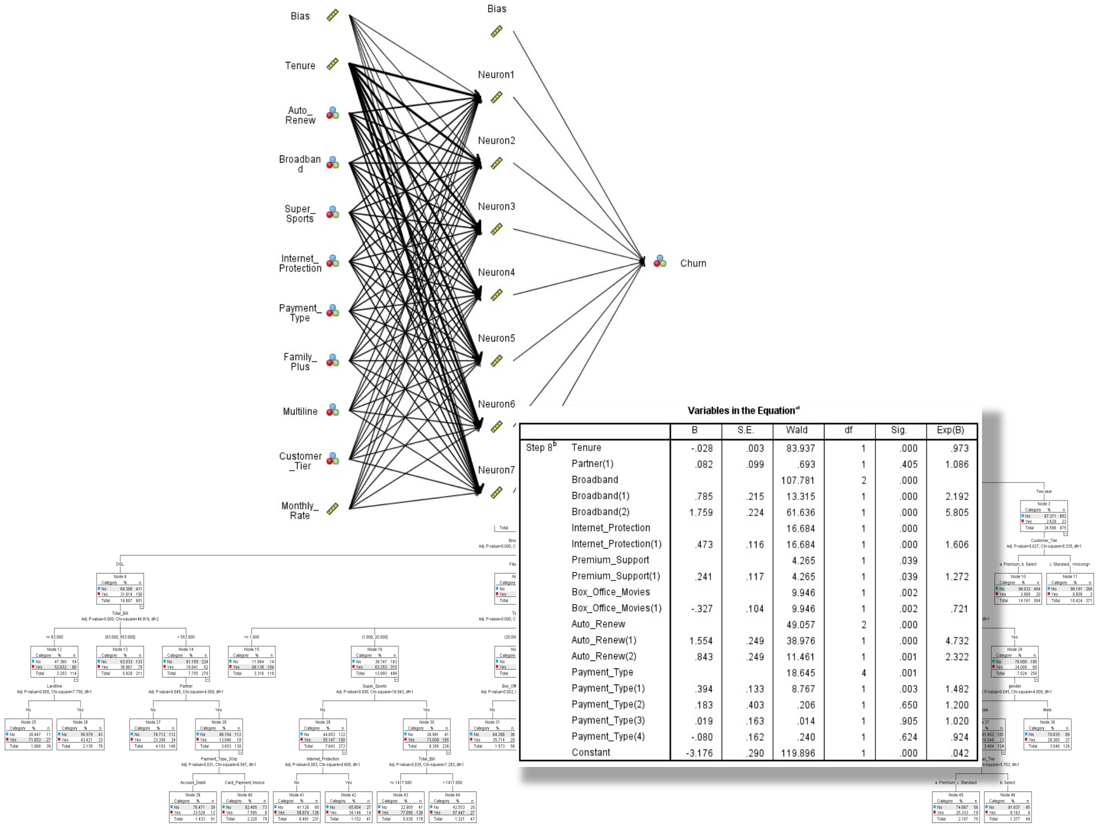

Variables in the Equation[a]

		B	S.E.	Wald	df	Sig.	Exp(B)
Step 8[b]	Tenure	-.028	.003	83.937	1	.000	.973
	Partner(1)	.082	.099	.693	1	.405	1.086
	Broadband			107.781	2	.000	
	Broadband(1)	.785	.215	13.315	1	.000	2.192
	Broadband(2)	1.759	.224	61.636	1	.000	5.805
	Internet_Protection			16.684	1	.000	
	Internet_Protection(1)	.473	.116	16.684	1	.000	1.606
	Premium_Support			4.265	1	.039	
	Premium_Support(1)	.241	.117	4.265	1	.039	1.272
	Box_Office_Movies			9.946	1	.002	
	Box_Office_Movies(1)	-.327	.104	9.946	1	.002	.721
	Auto_Renew			49.057	2	.000	
	Auto_Renew(1)	1.554	.249	38.976	1	.000	4.732
	Auto_Renew(2)	.843	.249	11.461	1	.001	2.322
	Payment_Type			18.645	4	.001	
	Payment_Type(1)	.394	.133	8.767	1	.003	1.482
	Payment_Type(2)	.183	.403	.206	1	.650	1.200
	Payment_Type(3)	.019	.163	.014	1	.905	1.020
	Payment_Type(4)	-.080	.162	.240	1	.624	.924
	Constant	-3.176	.290	119.896	1	.000	.042

A SELECT INTERNATIONAL COMPANY

Insight vs Action

- Insights are about achieving deeper understanding
- Applications are about decisions and actions

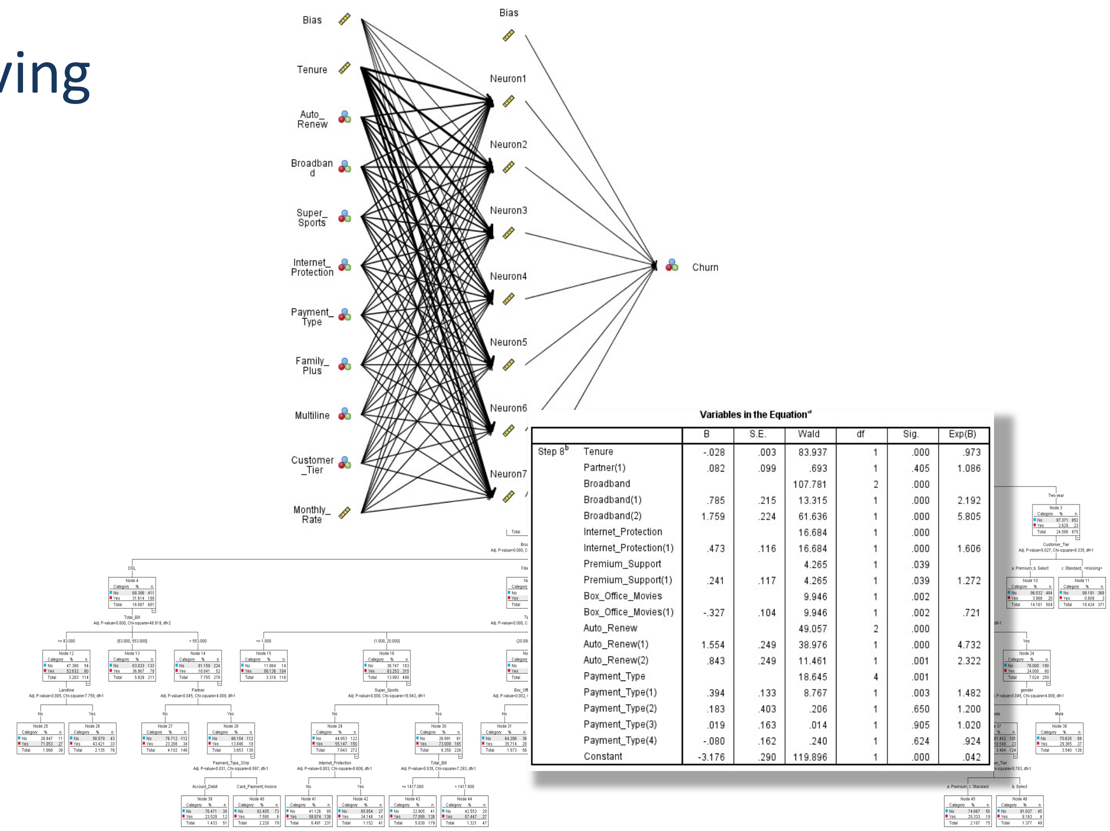

Variables in the Equation[a]

		B	S.E.	Wald	df	Sig.	Exp(B)
Step 8[b]	Tenure	-.028	.003	83.937	1	.000	.973
	Partner(1)	.082	.099	.693	1	.405	1.086
	Broadband			107.781	2	.000	
	Broadband(1)	.785	.215	13.315	1	.000	2.192
	Broadband(2)	1.759	.224	61.636	1	.000	5.805
	Internet_Protection			16.684	1	.000	
	Internet_Protection(1)	.473	.116	16.684	1	.000	1.606
	Premium_Support			4.265	1	.039	
	Premium_Support(1)	.241	.117	4.265	1	.039	1.272
	Box_Office_Movies			9.946	1	.002	
	Box_Office_Movies(1)	-.327	.104	9.946	1	.002	.721
	Auto_Renew			49.057	2	.000	
	Auto_Renew(1)	1.554	.249	38.976	1	.000	4.732
	Auto_Renew(2)	.843	.249	11.461	1	.001	2.322
	Payment_Type			18.645	4	.001	
	Payment_Type(1)	.394	.133	8.767	1	.003	1.482
	Payment_Type(2)	.183	.403	.206	1	.650	1.200
	Payment_Type(3)	.019	.163	.014	1	.905	1.020
	Payment_Type(4)	-.080	.162	.240	1	.624	.924
	Constant	-3.176	.290	119.896	1	.000	.042

Ideally, a predictive model is:

- Accurate
- Stable
- Transparent
- Simple
- Coherent
- Deployable

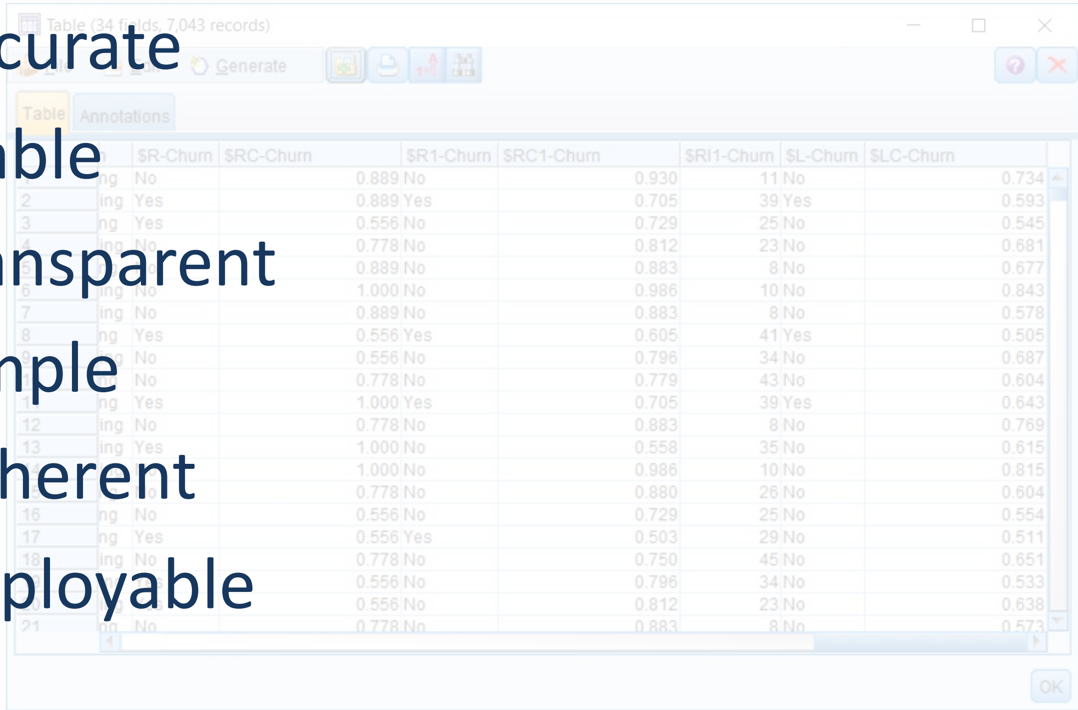

But what is the right decision if our ideal model indicates…

- A patient has a 20% chance of being readmitted to hospital within 30 days
- A customer has a 33% likelihood of not renewing a contract
- 213 transactions are flagged as ‘anomalous’
- A student has a 62% chance of accepting an offer of a university place
- 1,500 guests are estimated to be unsatisfied with their previous stay at a hotel

Our ideal predictive application is:

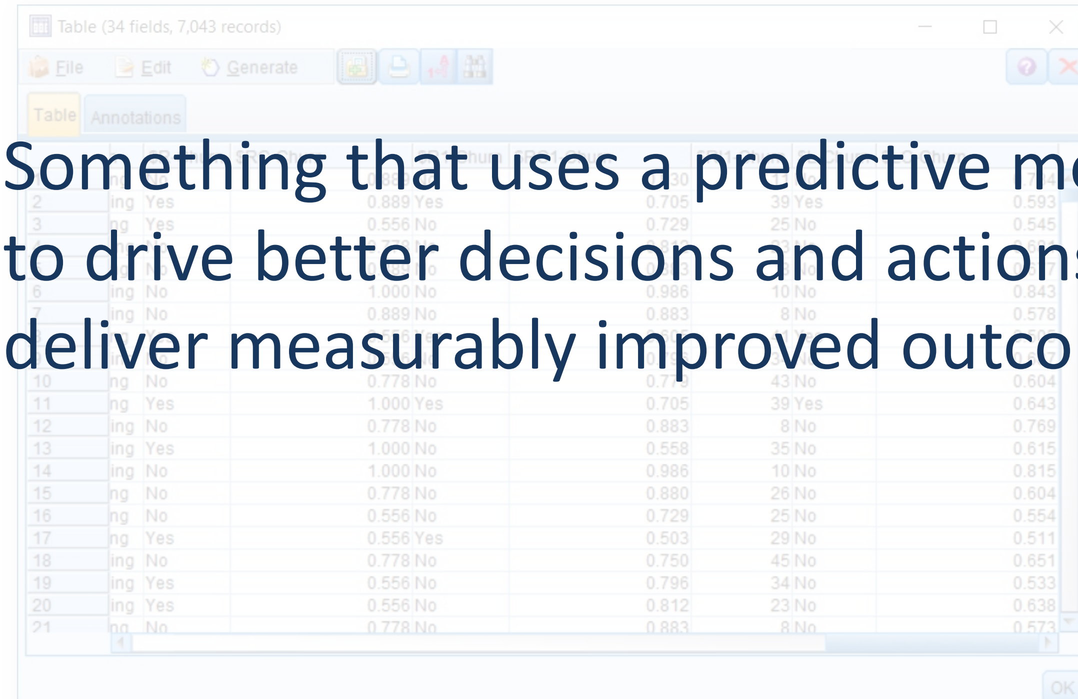

- Something that uses a predictive model to drive better decisions and actions to deliver measurably improved outcomes.

Given that…

- Decisions drive actions
- Actions have consequences
- …and therefore costs
- The context in which these new actions are taken is of critical importance. Especially if:
  - They are simply unfeasible
  - They impact others in the organisation
  - They use up existing resources
  - They are not monitored or measured
  - They are unfair

Analytical Lifecycle Methodologies: tools for developing applications

- Microsoft's Team Data Science Process (TDSP)

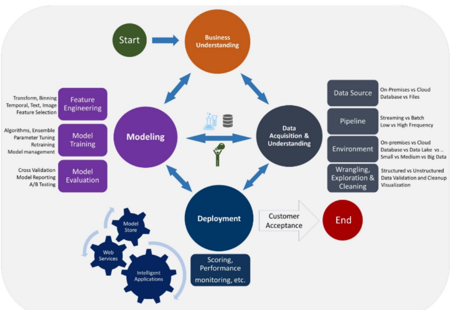

- EMC's Data Analytics Lifecycle

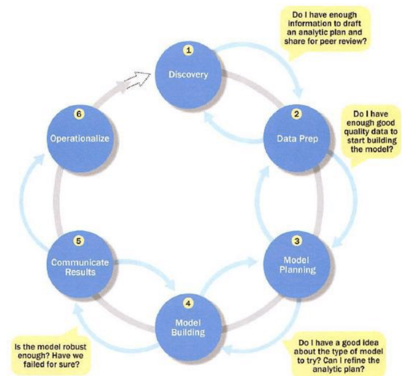

IBM's Analytics Solution Unified Method (ASUM-DM)

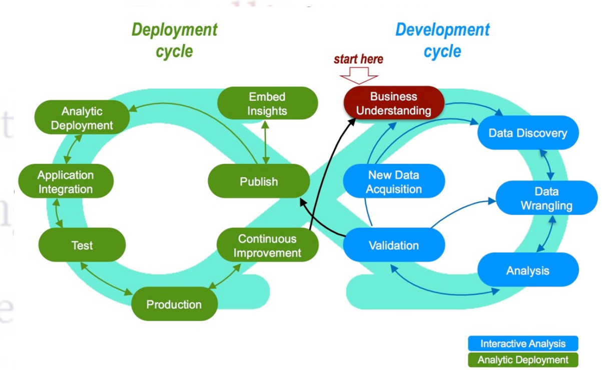

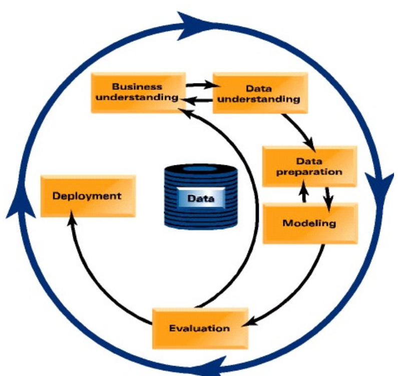

- Cross-Industry Standard Process for Data Mining (CRISP-DM)

Analytical Lifecycle Methodologies

- Are not designed just to make analytics more understandable for non-technical audiences
- They are also supposed to remind data analysts that activities like predictive modelling exist within a wider context

The Cross-Industry Standard Process for Data Mining (CRISP-DM)

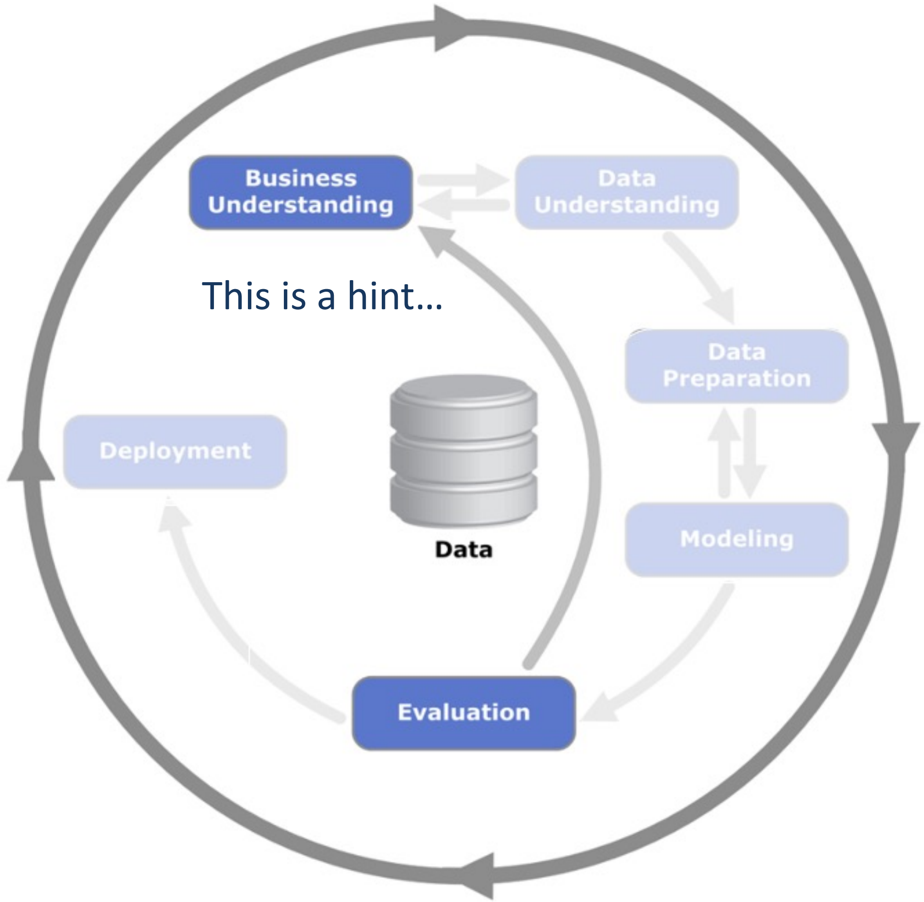

Poorly Documented Business Understanding

- “Aim of the project is to build an accurate model to predict subscribers at risk of cancelling their contracts”
- “By identifying at risk subscribers, effective action in the form of additional offers can be taken to reduce the likelihood of churn”

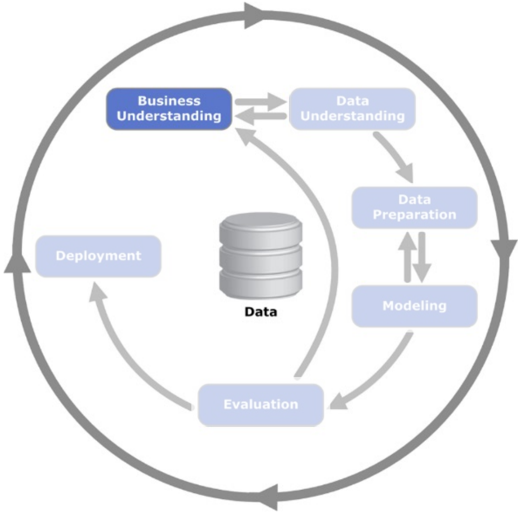

Well Documented Business Understanding

- Primary Business objective: **Increasing customer loyalty**
- "…stipulated as a strategic goal for the company. It's recognised that costs associated with customer acquisition have been rising in recent years and that increasing pressure from competitor activity has led to a slowdown in the growth of market share". **So this is regarded as a *valuable objective*.**
- "Due to the contractual basis of our business model, we can identify precisely when new customers are acquired and when they cease to transact with us". **So this is a *measurable outcome***.
- "Previous analysis using satisfaction surveys has shown that if we are able to identify customers likely to cancel their contracts three months before their termination date, we can persuade around 50% of them to renew their contracts with us". **So this can be regarded as an *actionable result*.**

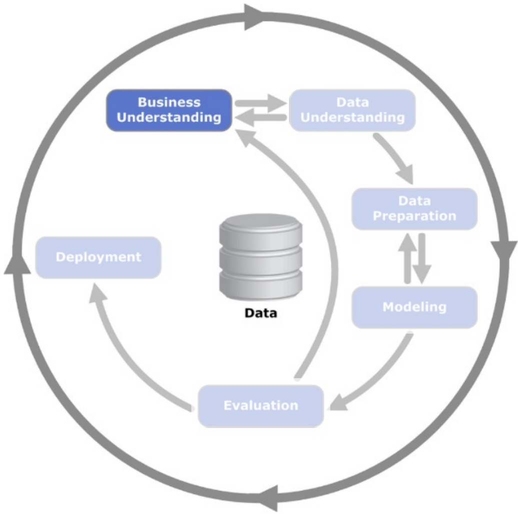

Well Documented Business Understanding

- The average cost incurred with persuading customers to renew an existing contract is **$35** (this is only incurred by those who agree to renew)
- The average annual revenue received from a contract is **$132**
- The cost per head of acquiring and onboarding new customers is **$45**
- Assume that it takes one month to replace a customer, then the lost revenue on average is only **$11**
- However, to replace the customer, the acquisition and onboarding costs mean that this value jumps to **$56**
- If the company is losing 100,000 customers a year (a not unreasonable number) then the total costs are **$5.6 million**
- Let's assume a modest-performing model identifies what it thinks are the *top **30,000 customers who are likely to leave*** (or 'churn') annually. So that is only 30% of the 100,000 churners

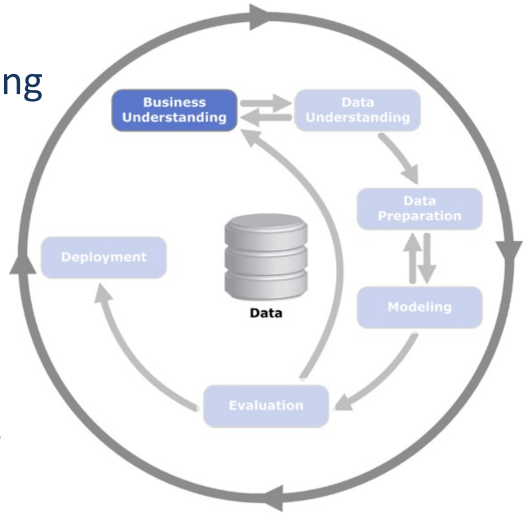

Well Documented Business Understanding

- Of course, we can't assume that the model is completely accurate, so let's assume that it's only right 2 out of 3 times. Meaning that we may only identify **20,000** customers who will churn annually.
- But we can only hope to persuade **50%** of them to remain customers. So that's **10,000** customers we have retained and **90,000** customers who cancel their contracts
- It now costs the company **$5.04** million to replace the lost customers and they will incur additional retention costs of **$350K** to persuade **10,000** customers to renew their contracts bringing the total costs to **$5.39** million.
- This represents a relatively modest but worthwhile *cost reduction of **$210K***.
- Crucially though, in doing so they will have managed to retain 10,000 customers with total annual revenues of **$1.32 million**.

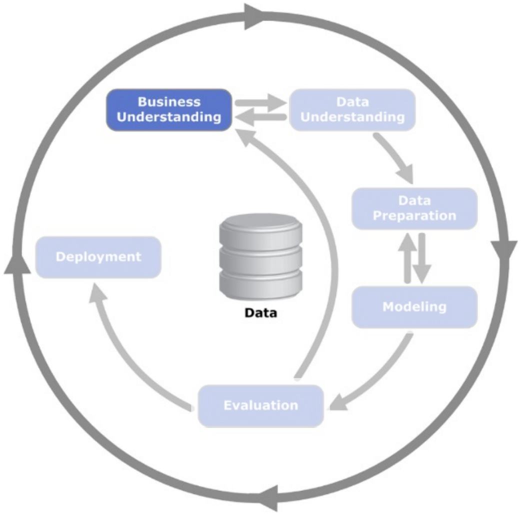

Comparing measures of predictive accuracy: Continuous Targets

Accuracy Measures for Continuous Targets

- Continuous data consist of numerical outcomes like revenue in dollars, temperature in Celsius, number of insurance claims, weight in kg or time in minutes
- There are a number of predictive accuracy measures that are designed to be used when the target (or dependent) variable in a predictive modelling application is continuous

claim_amount	predicted_claim_amount
[illegible]	[illegible]
243	240
179	183
264	125
334	317
274	284
208	208
150	136
383	427
305	369
116	250
636	153
236	258
270	235
160	181
110	124
259	255
226	233
161	179
107	124
340	306
260	275
189	203
104	134
400	407
349	354
147	243
65	151
207	209
129	194
157	155
113	111
208	207

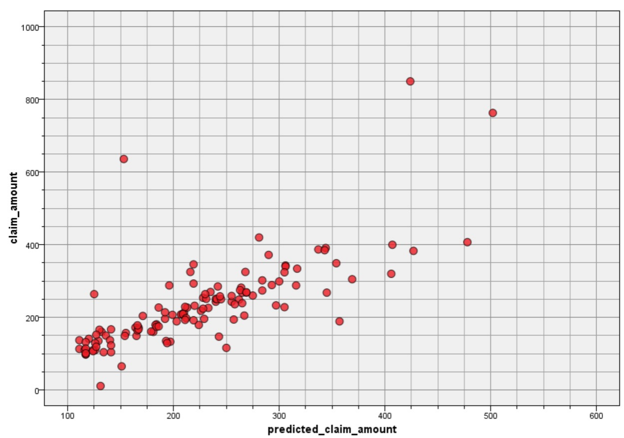

Accuracy Measures for Continuous Targets

- A scatterplot is an easy way to visualise predictive accuracy with continuous targets

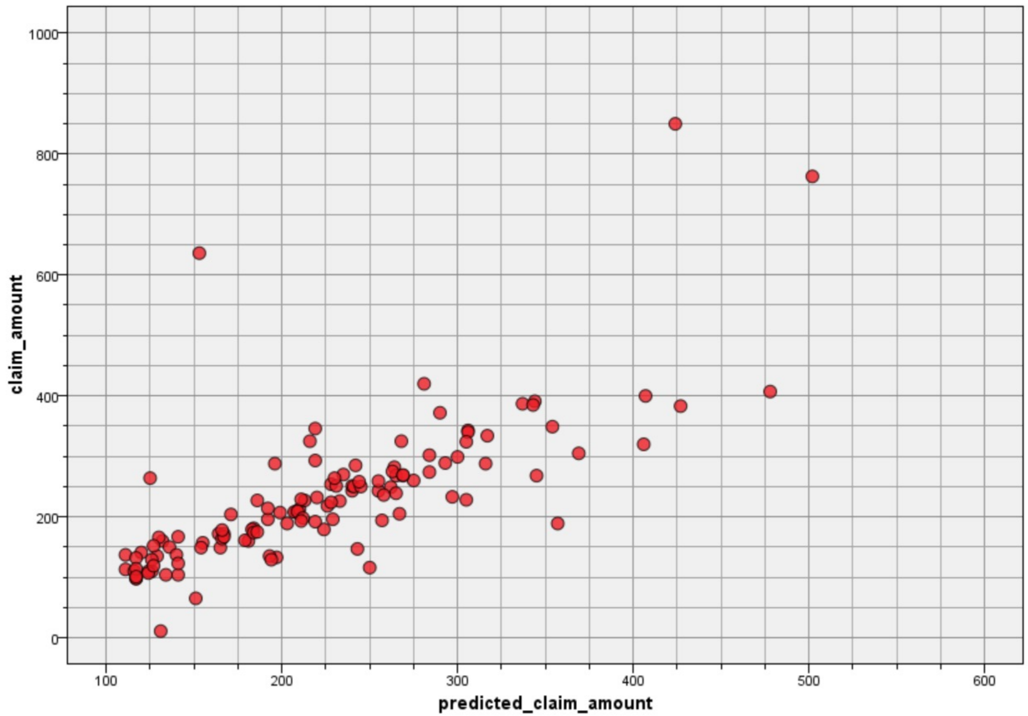

claim_amount	predicted_claim_amount
[illegible]	[illegible]
243	240
179	183
264	125
334	317
274	284
208	208
150	136
383	427
305	369
116	250
636	153
236	258
270	235
160	181
110	124
259	255
226	233
161	179
107	124
340	306
260	275
189	203
104	134
400	407
349	354
147	243
65	151
207	209
129	194
157	155
113	111
208	207

Simple Correlations

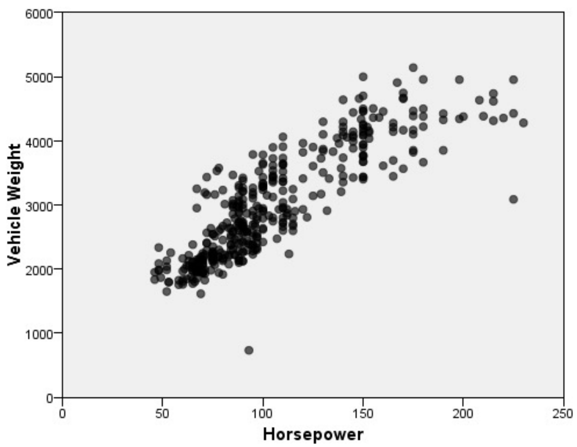

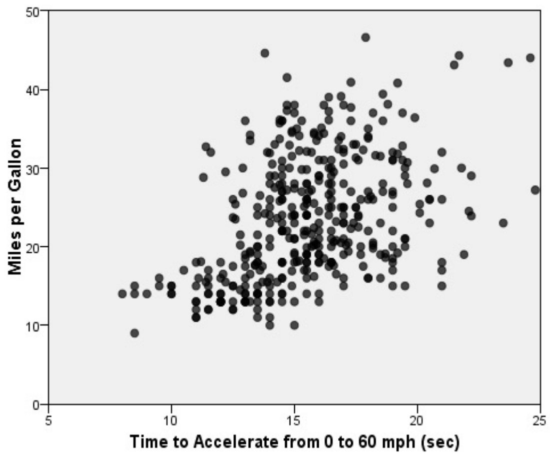

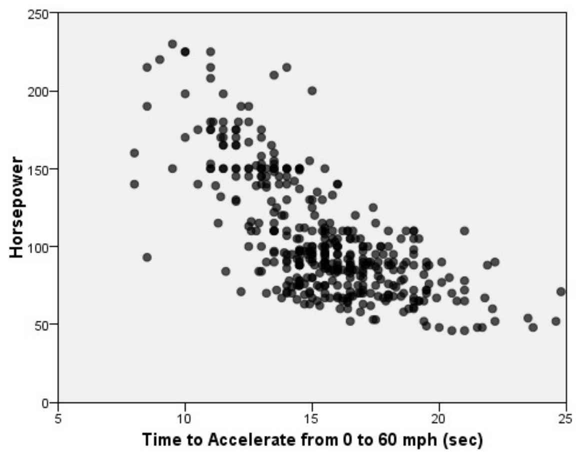

0.859

0.434

-.701

Pearson Correlation Values range between -1 and +1

Simple Correlations

- Pearson's r = 0.75
- If we square this number we can get an 'R-Squared' value
- So 0.75 squared becomes 0.56
- An R-Square value of 0.56 means that 56% of the variation in one field can be accounted for by another
- So larger R-Square values mean a better 'model fit' i.e. greater predictive accuracy

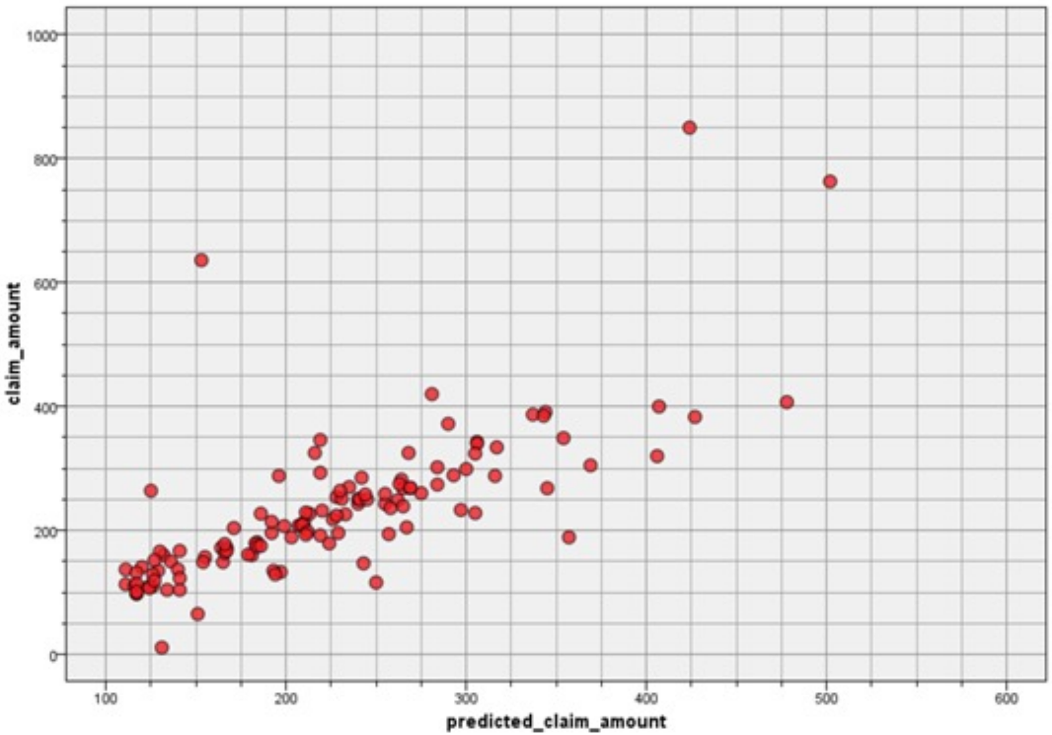

Model Fit Statistics for Continuous Targets

Larger Values = better fit
Smaller values = better fit

- Correlations values such as **R**
- Regression fit statistics like **R^2**
- **Mean Error** - average (mean) of errors across all records (i.e. positive and negative values)
- **Mean Absolute Error (MAE)** - average of the absolute values of the errors across all records
- **Mean Squared Error (MSE)** - average of squared errors
- **Root Mean Squared Error (RMSE)** - square root of average squared errors
- **Akaike information criterion (AIC)** - penalized-likelihood criterion
- **Bayesian information criterion (BIC)** - penalized-likelihood criterion

Comparing measures of predictive accuracy: Categorical Targets

Predicting Categorical Targets

- Categorical data consist of values that represent categories such as fault type, churn status, response to offer or movie genre
- So measures like R^2 don't make sense
- Predictions for categorical targets normally consist of the predicted category and one or more values showing the prediction likelihood.

Churn	Predicted Churn	Confidence_Score	Likelihood of Churn
Yes	No	0.779	0.221
Yes	No	0.585	0.415
No	No	0.978	0.022
Yes	No	0.757	0.243
No	No	0.709	0.291
No	No	0.933	0.067
Yes	No	0.899	0.101
No	No	0.573	0.427
No	No	0.964	0.036
No	No	0.980	0.020
Yes	Yes	0.793	0.793
No	No	0.661	0.339
No	No	0.885	0.115
Yes	No	0.715	0.285
No	No	0.671	0.329
No	No	0.563	0.437
No	No	0.938	0.062
No	No	0.792	0.208
No	No	0.563	0.437
Yes	Yes	0.631	0.631
Yes	Yes	0.709	0.709

Predicting Categorical Targets

- In the screenshot we can see:
  1. The predicted category
  2. A confidence score showing the likelihood of that record's prediction being true.
  3. A probability score showing the likelihood that a record will belong to a specific target group such as churn. These scores can be used to select the records with the highest risk of churning.

	1	2	3
Churn	Predicted Churn	Confidence_Score	Likelihood of Churn
Yes	No	0.779	0.221
Yes	No	0.585	0.415
No	No	0.978	0.022
Yes	No	0.757	0.243
No	No	0.709	0.291
No	No	0.933	0.067
Yes	No	0.899	0.101
No	No	0.573	0.427
No	No	0.964	0.036
No	No	0.980	0.020
Yes	Yes	0.793	0.793
No	No	0.661	0.339
No	No	0.885	0.115
Yes	No	0.715	0.285
No	No	0.671	0.329
No	No	0.563	0.437
No	No	0.938	0.062
No	No	0.792	0.208
No	No	0.563	0.437
Yes	Yes	0.631	0.631
Yes	Yes	0.709	0.709

Note that the highlighted records have the same scores for **confidence** and **likelihood** as they both measure the same outcome: the likelihood of being in the 'Yes' group

Predicting Categorical Targets

- A simple way to evaluate the accuracy of a model with a categorical target is with a crosstab or 'confusion' matrix

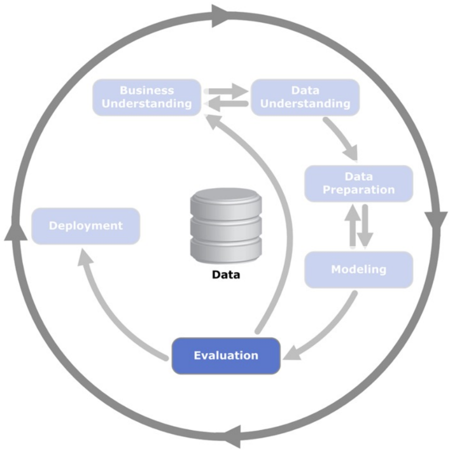

Model A: Actual Churn by Predicted Churn

			Churn_Predicted	
			No	**Yes**
Churn_Actual	**No**	Frequency	2330	278
		Percent Correct	**89.3%**	10.7%
	Yes	Frequency	473	493
		Percent Correct	49.0%	**51.0%**

Predicting Categorical Targets

- Ok. So Which model is the best one?

Model A: Actual Churn by Predicted Churn

			Churn_Predicted	
			No	Yes
Churn_Actual	No	Frequency	2330	278
		Percent Correct	**89.3%**	10.7%
	Yes	Frequency	473	493
		Percent Correct	49.0%	**51.0%**

Model B: Actual Churn by Predicted Churn

			Churn_Predicted	
			No	Yes
Churn_Actual	No	Frequency	2129	479
		Percent Correct	**81.6%**	18.4%
	Yes	Frequency	307	659
		Percent Correct	31.8%	**68.2%**

Model C: Actual Churn by Predicted Churn

			Churn_Predicted	
			No	Yes
Churn_Actual	No	Frequency	1919	689
		Percent Correct	**73.6%**	26.4%
	Yes	Frequency	214	752
		Percent Correct	22.2%	**77.8%**

Predicting Categorical Targets

- What are the *costs* of misclassifying cases?

Model A: Actual Churn by Predicted Churn

			Churn_Predicted	
			No	**Yes**
Churn_Actual	**No**	Frequency	2330	False 78
		Percent Correct	89.3%	Positive 10.7%
	Yes	Frequency	False 73	493
		Percent Correct	Negative 49.0%	51.0%

Predicting Categorical Targets

- What are the *costs* of misclassifying cases?

Model A: Actual Churn by Predicted Churn

			Churn_Predicted No	Churn_Predicted Yes
Churn_Actual	No	Frequency	2330	278
		Percent Correct	**89.3%**	10.7%
	Yes	Frequency	473	493
		Percent Correct	49.0%	**51.0%**

Model B: Actual Churn by Predicted Churn

			Churn_Predicted No	Churn_Predicted Yes
Churn_Actual	No	Frequency	2129	479
		Percent Correct	**81.6%**	18.4%
	Yes	Frequency	307	659
		Percent Correct	31.8%	**68.2%**

Model C: Actual Churn by Predicted Churn

			Churn_Predicted No	Churn_Predicted Yes
Churn_Actual	No	Frequency	1919	689
		Percent Correct	**73.6%**	26.4%
	Yes	Frequency	214	752
		Percent Correct	22.2%	**77.8%**

Predicting Categorical Targets

- Gains charts are used to evaluate model performance

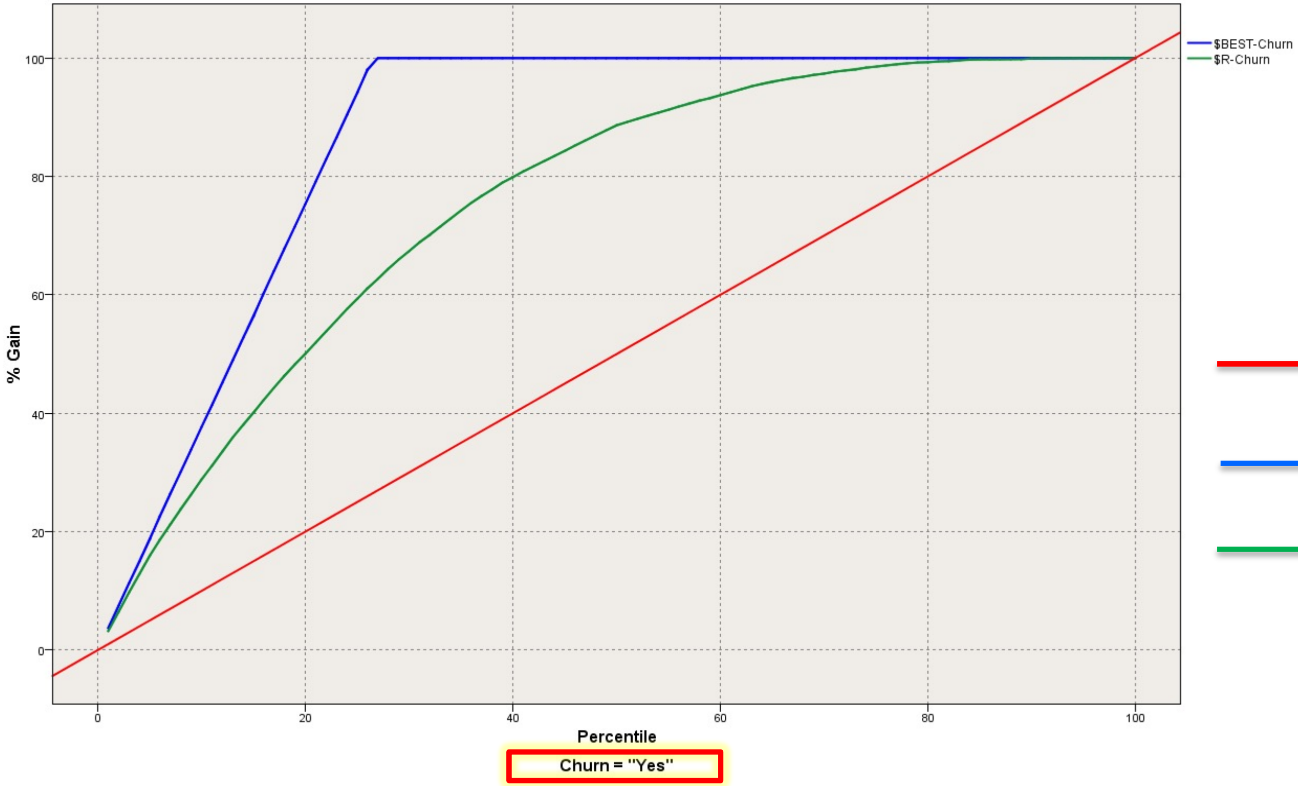

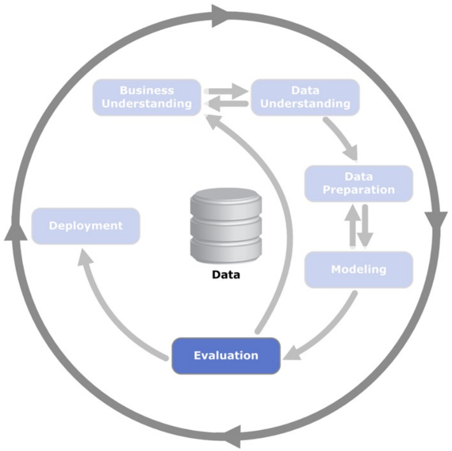

Random performance

Perfect performance

Model performance

Predicting Categorical Targets

- Think of the model as a kind of filter

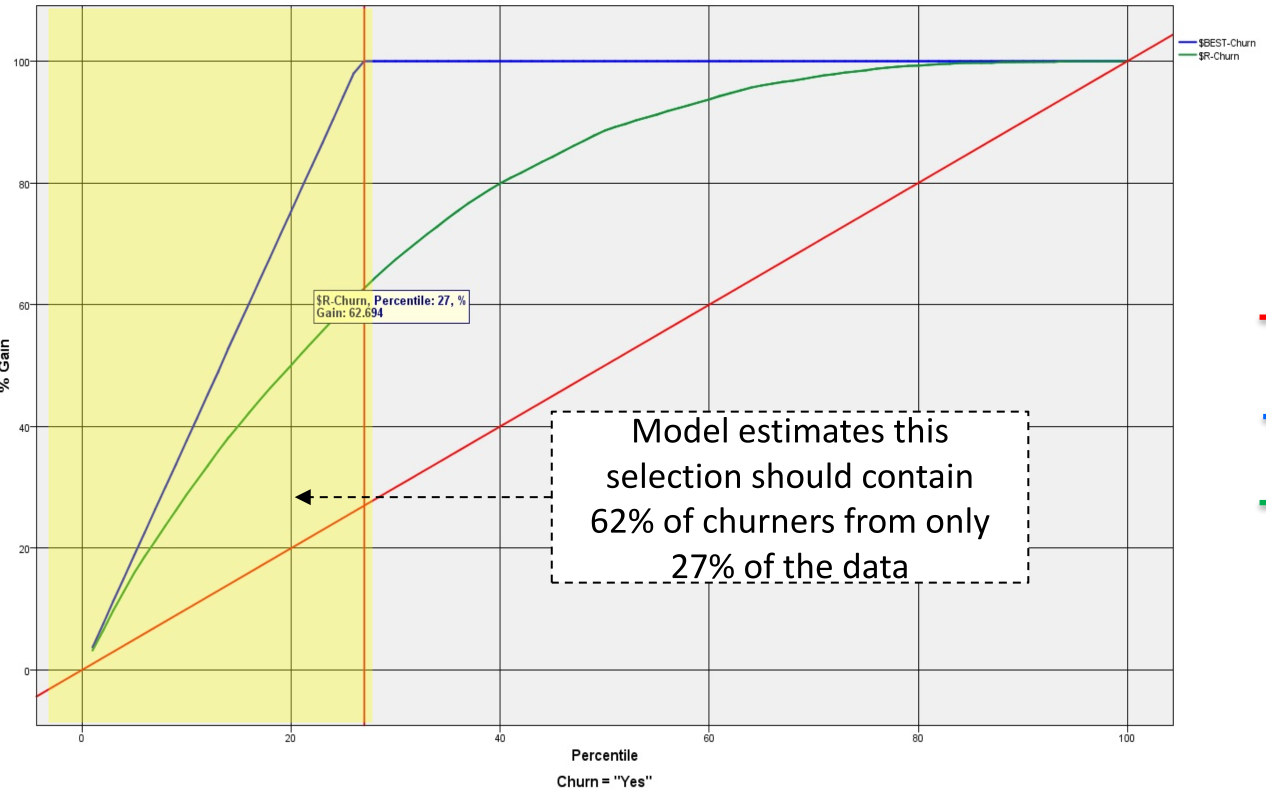

Random performance

Perfect performance

Model performance

Predicting Categorical Targets

- These kinds of charts make it possible to compare model performance

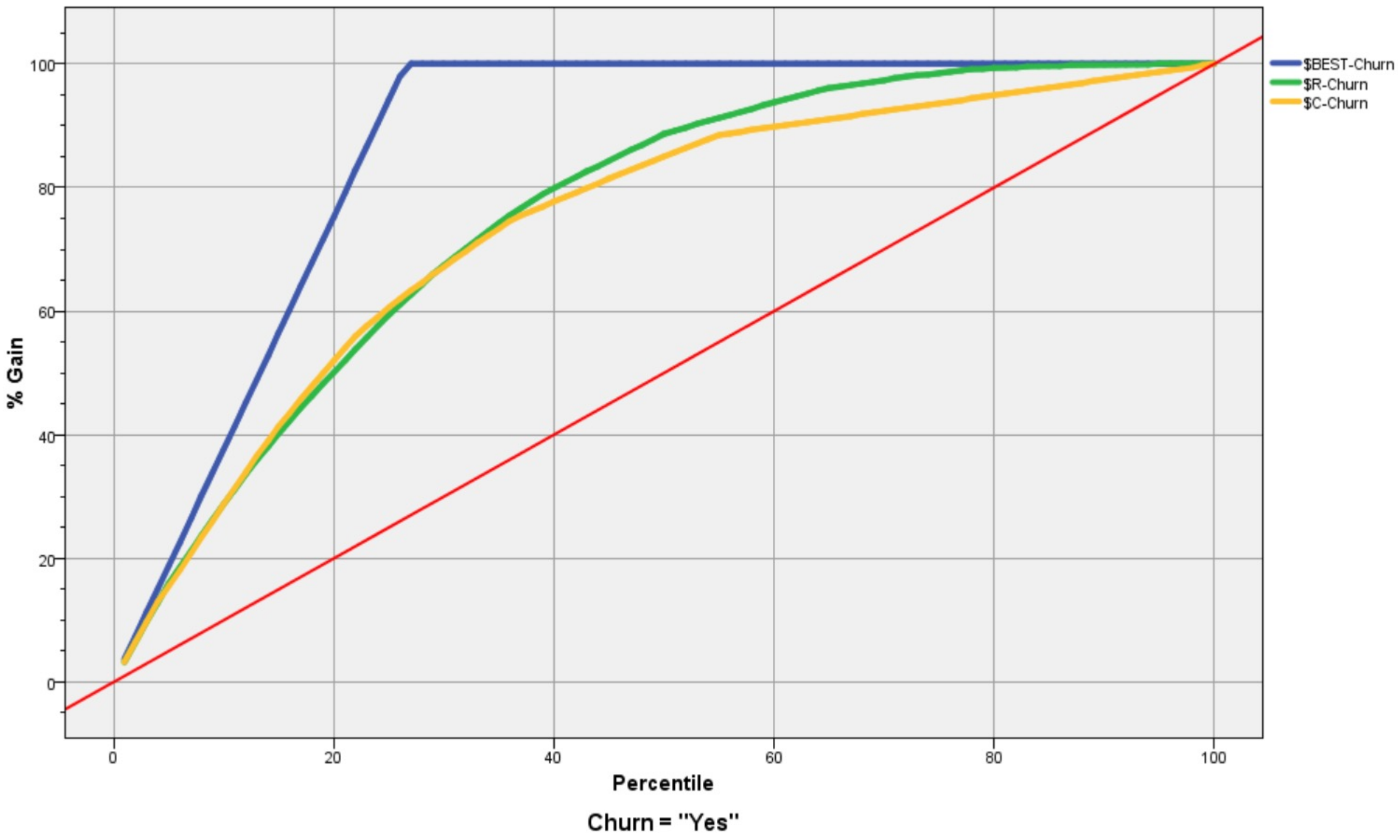

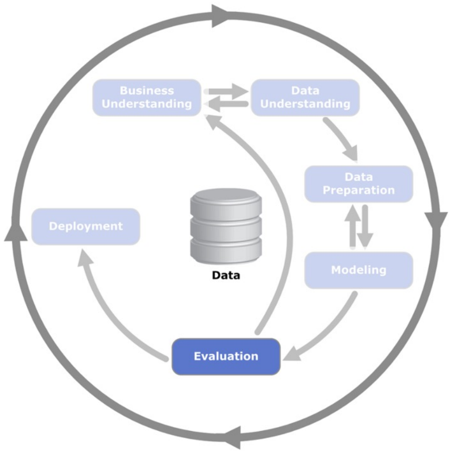

Predicting Categorical Targets

- An ROC (Receiver Operator) chart is commonly used

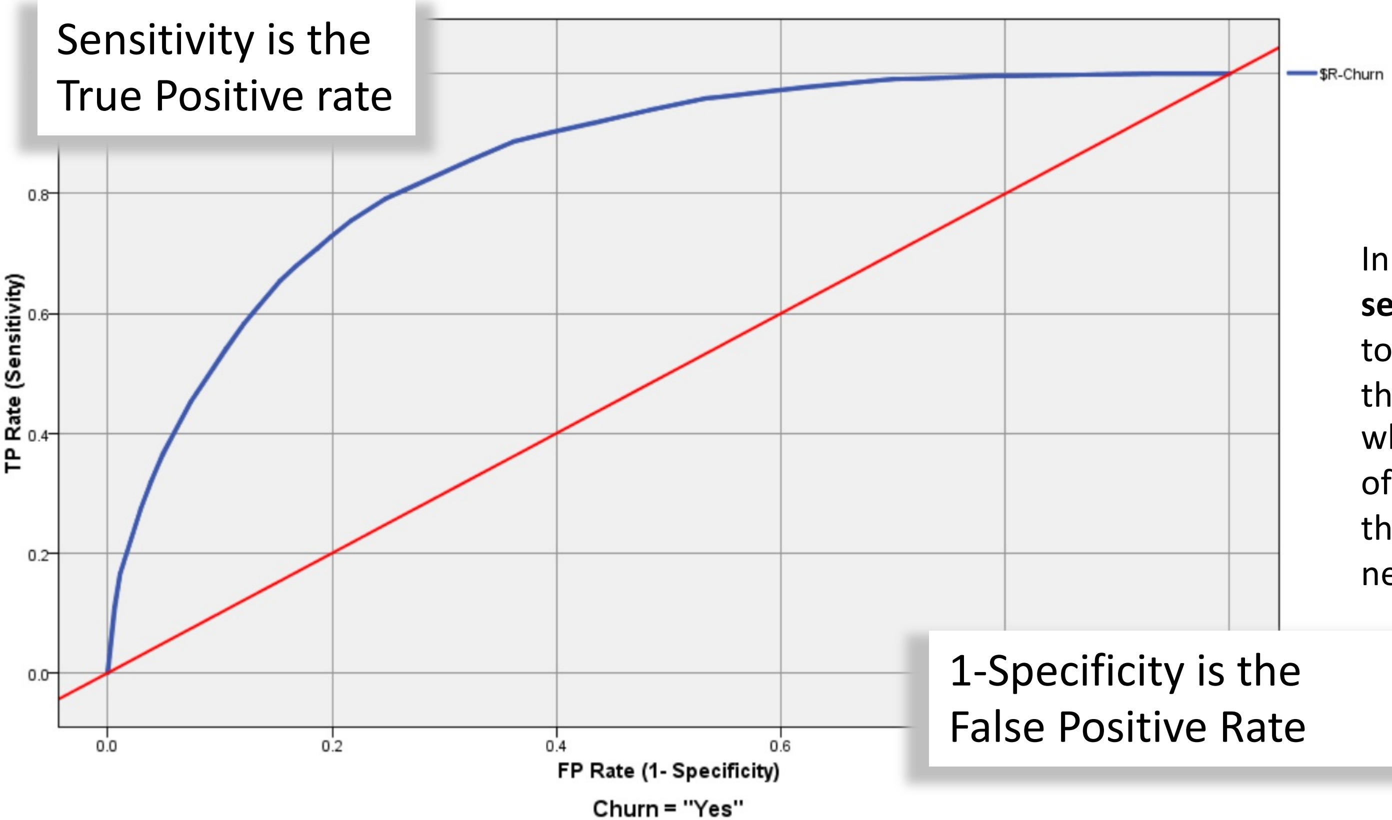

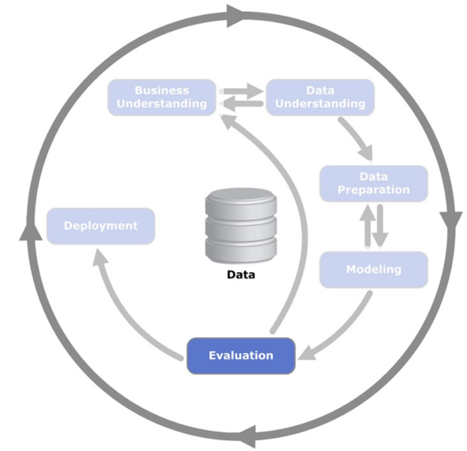

In science and medicine **sensitivity** is the ability of a test to correctly identify those with the disease (true positive rate), whereas **specificity** is the ability of the test to correctly identify those without the disease (true negative rate)

Predicting Categorical Targets

- Selecting cases with a ROC (Receiver Operator) Chart

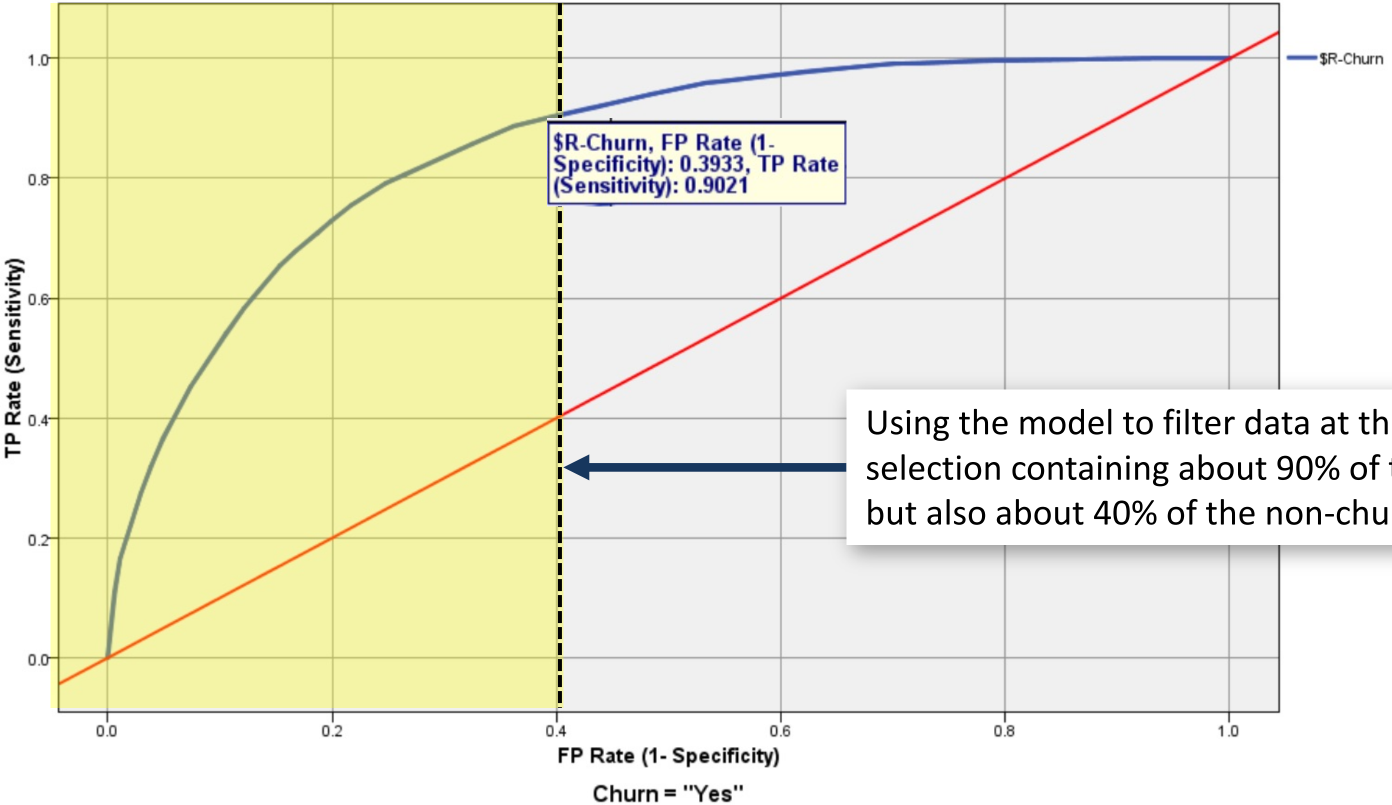

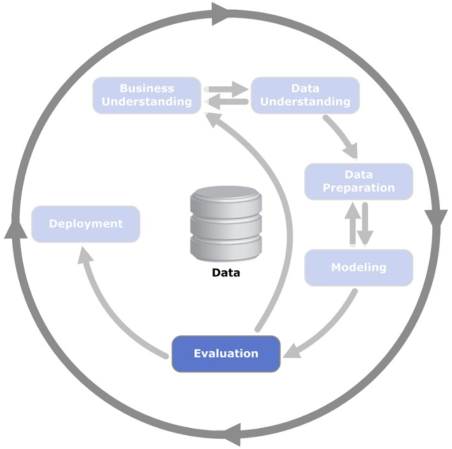

Using the model to filter data at this point would result in a selection containing about 90% of the Churners (TP = 0.9021) but also about 40% of the non-churners (FP = 0.3933)

Model Fit Statistics for Categorical Targets

- **Overall Accuracy –** Percentage of predictions that are correct. Best for when the targets groups have roughly equal proportions (e.g. 50/50) or when false positives and false negatives have equal importance
- **Precision –** Proportion of outcomes *predicted to be true* that are actually true
- **Recall –** Proportion of outcomes that *are true* that the model correctly predicted
- **F_1 –** An attempt to balance the *Precision and Recall* metrics. Larger values are indicate greater accuracy.
- **Area Under the Curve (AUC) –** Based on the ROC charts we saw earlier. Random models have AUC values of 0.5 since we would expect to find 50% of the predicted occurrences from 50% of the data by random sampling. Therefore only models worse than random would have scores less than 0.5. In general AUC scores of 0.8 and above indicate reasonably good predictive accuracy.

Model Fit Statistics for Categorical Targets

- **Gini Coefficient –** A commonly used measure in credit scoring applications. It's actually based on the Area Under the Curve (AUC) measure. But unlike the AUC value, with Gini random models and *worse than random* models both equal 0. Which some people feel is a more intuitive range of values.

- **Lift value -** A measure of how much better the model is at predicting an outcome compared to a random approach. If 10% of the data is comprised of customers who have churned, a random model predicting every case to be a churner would be right about 10% of the time. This would generate a lift value equal to 1.0. If however the model was better than random and was able to predict the proportion of churners with 20% accuracy, it would be twice as accurate as the random approach and would generate a lift value of 2.0. Often Lift values are calculated on a proportion of the data, say the top 30%, where the model is most confident.

Testing Models

Testing Model Performance

- Predictive models should always be tested. Ideally on a separate random data selection.

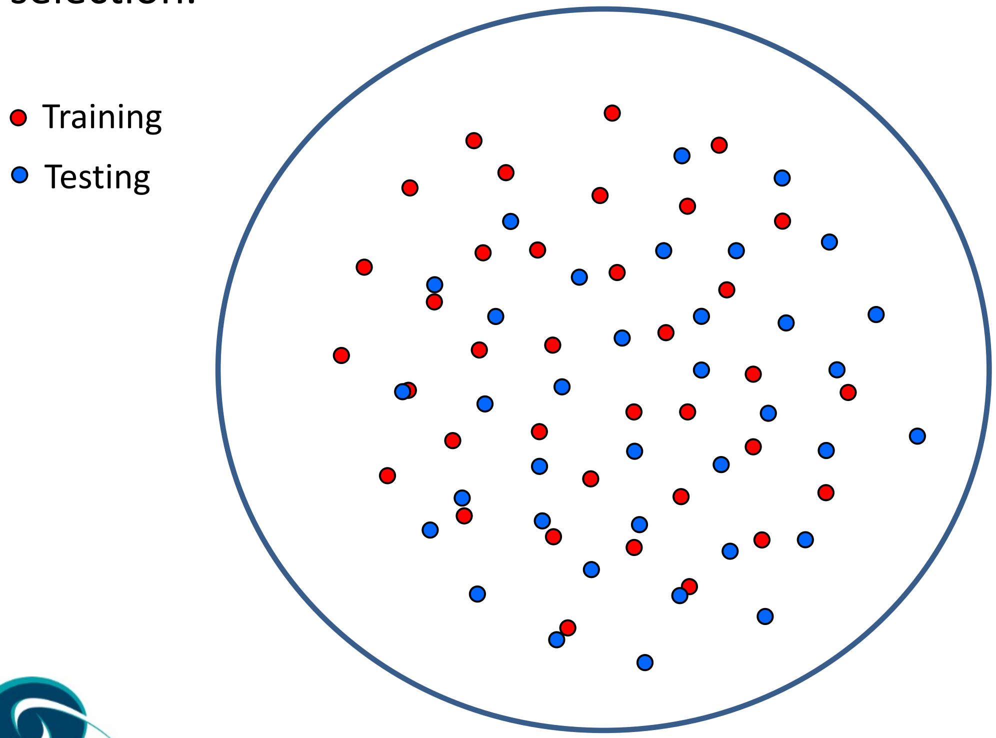

Churn	Partition
No	1_Training
No	1_Training
No	1_Training
Yes	1_Training
No	2_Testing
Yes	2_Testing
No	1_Training
Yes	2_Testing
Yes	1_Training
Yes	2_Testing
Yes	2_Testing
No	1_Training
Yes	2_Testing
No	2_Testing
No	2_Testing
No	1_Training
No	2_Testing
Yes	1_Training
No	1_Training
Yes	2_Testing

Testing Model Performance

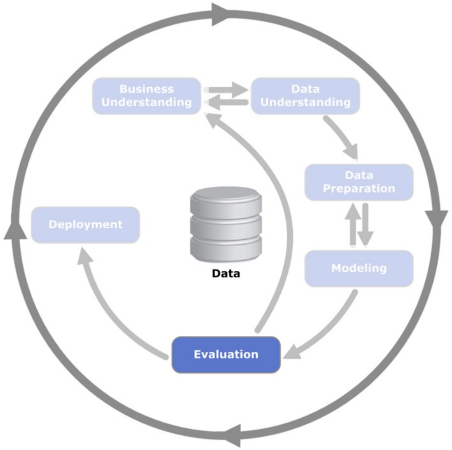

- Compare model performance on ‘Train’ and ‘Test’ sample

Graph	Model	Area Under Curve
	CHAID 1	0.855
	Logistic regres...	0.849
	Neural Net 1	0.848
	Bayesian Netw...	0.838
	C5 1	0.826

Training Sample Performance

Graph	Model	Area Under Curve
	Logistic regres...	0.846
	Neural Net 1	0.837
	CHAID 1	0.835
	Bayesian Netw...	0.832
	C5 1	0.812

Testing Sample Performance

Testing Model Performance

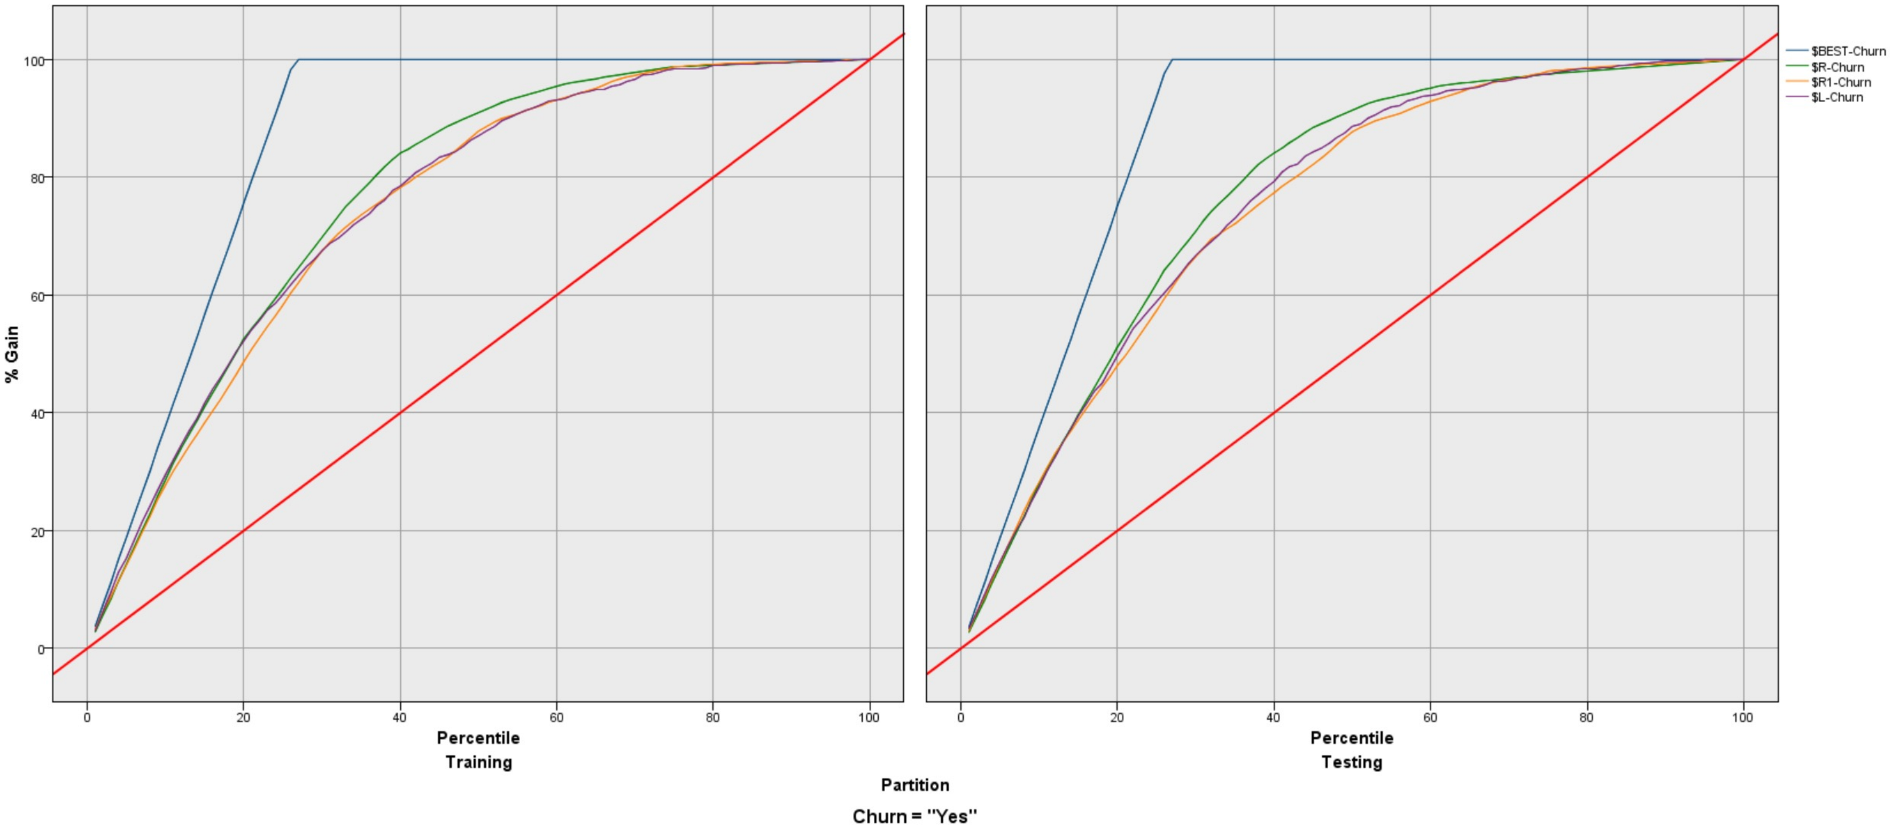

Testing Model Performance

- **Cross Validation**
  - Useful for when there isn't enough data to split into training and testing groups
  - Cross Validation splits the existing dataset into separated groups or 'folds' (e.g. 10)
  - It then trains a number of models against each fold and tests the performance against a small withheld proportion (e.g. 10%)
  - The results are aggregated to show overall performance

SMARTVISION
Europe

Deployment

Deployment happens in the real world

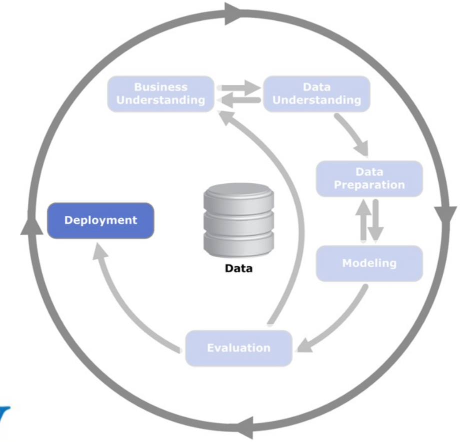

- How the models scores are used depends on the context on the application

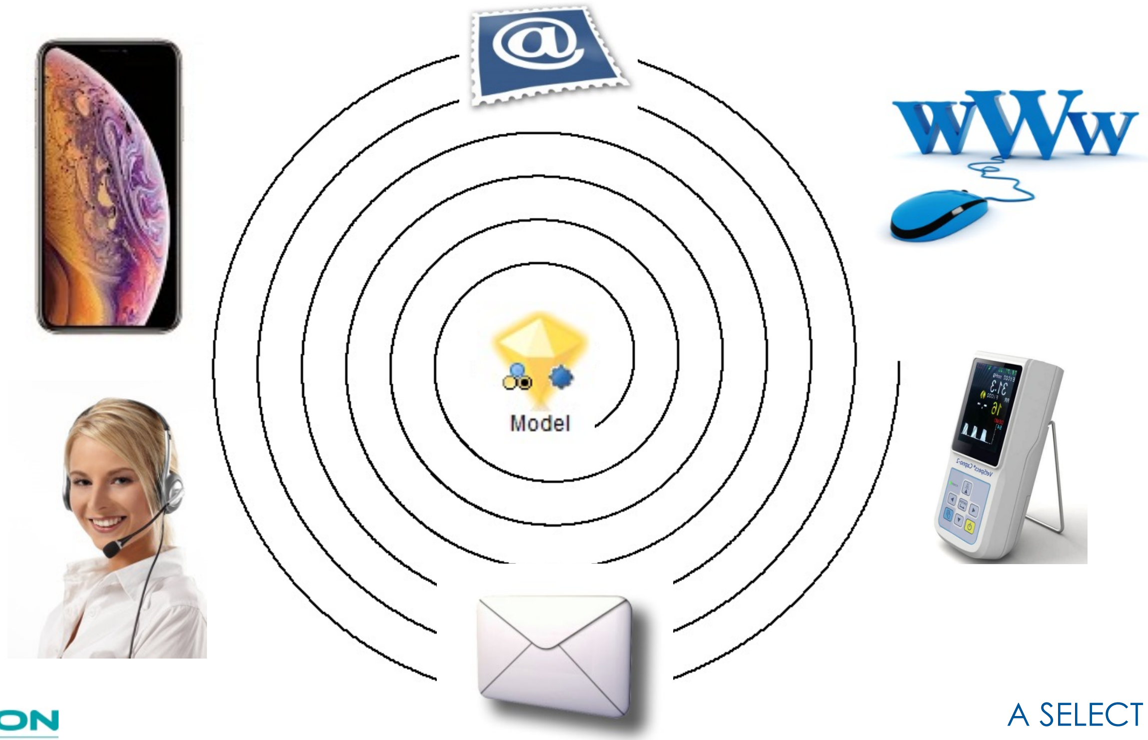

Deployment

	customerID	Churn_Likelihood
7472	6003	0.720
7473	5663	0.720
7474	4548	0.720
7475	5267	0.720
7476	5637	0.720
7477	6024	0.720
7478	2365	0.720
7479	3740	0.720
7480	523	0.720
7481	2835	0.720
7482	5261	0.720
7483	6106	0.720
7484	1762	0.720
7485	4711	0.720
7486	2441	0.720
7487	1036	0.720
7488	6883	0.720
7489	2588	0.720
7490	115	0.720
7491	4088	0.720
7492	1024	0.720
7493	19	0.720
7494	2999	0.720
7495	662	0.720
7496	5397	0.720
7497	3949	0.720
7498	6296	0.720
7499	3172	0.720
7500	5039	0.720
7501	5785	0.720
7502	4480	0.720
7503	1614	0.720
7504	1446	0.720
7505	6854	0.720
7506	5007	0.720
7507	4806	0.720

From our Business Understanding example earlier:

- “Let’s assume a modest-performing model identifies what it thinks are the *top **30,000 customers** who are likely to leave* (or ‘churn’) annually”
- So if we apply the model on a quarterly basis that would be **7,500** customers each quarter
- Therefore, having built the best model we can, we use it to select the **7,500** customers with the highest likelihood churning
- Also, remember that we’re only expecting the model to be right **2 out of 3** times
- And even then, we only expect to retain **50%** of the people where the model was correct

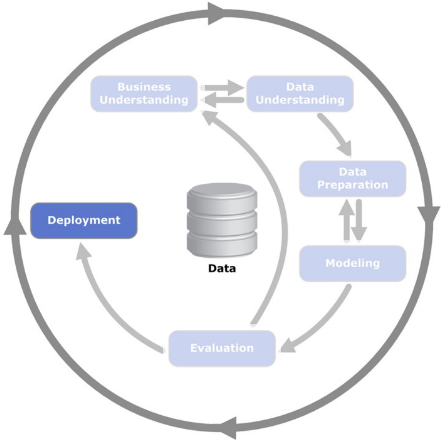

Deployment

- Deployment might take the form of simply choosing people based on a threshold likelihood value

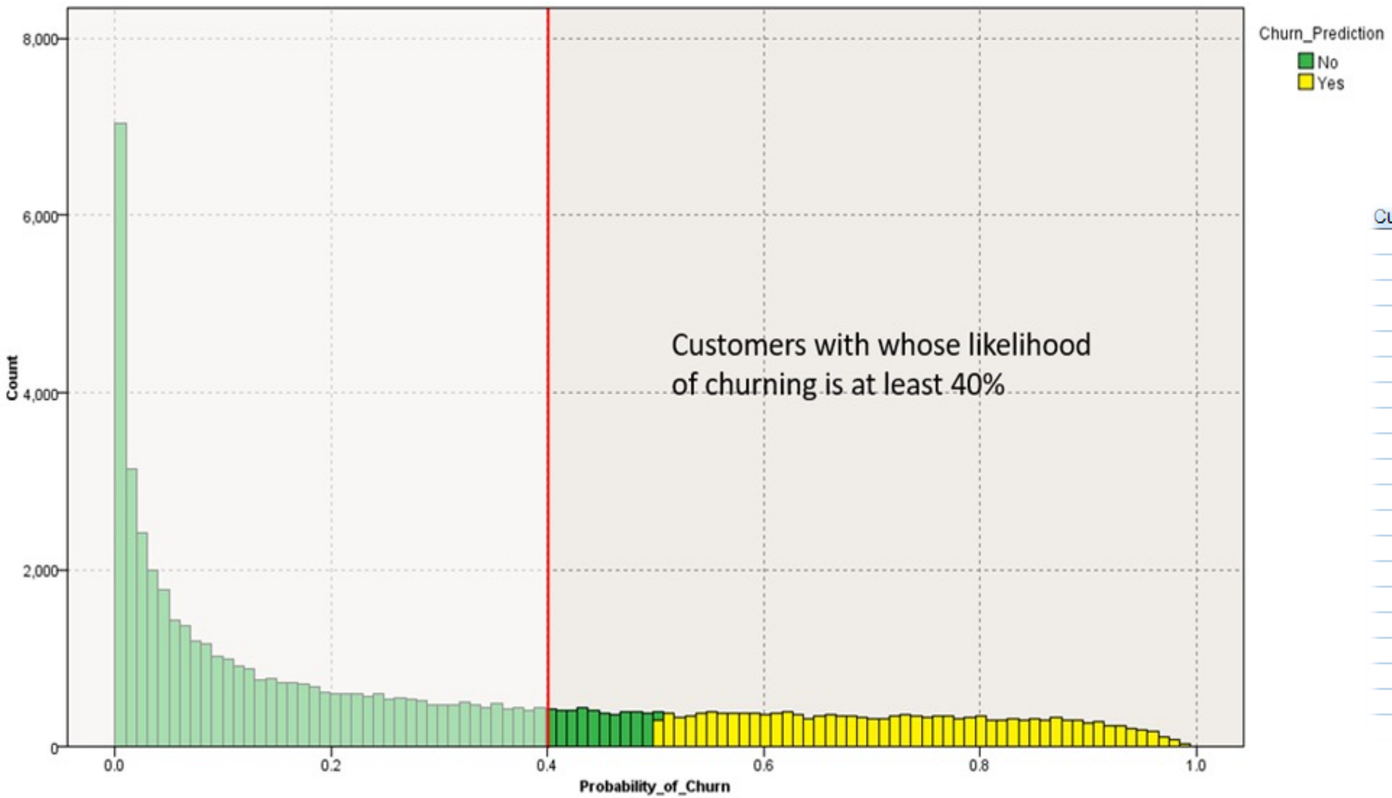

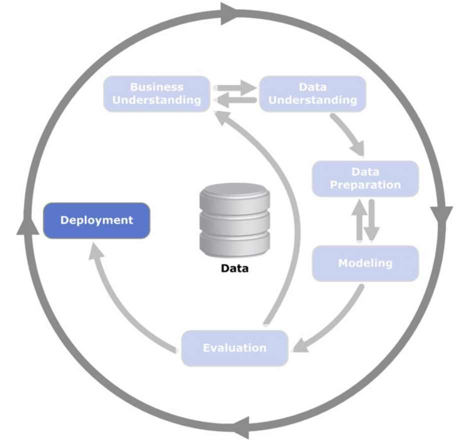

CustomerID	Churn_Prediction	Probability_of_Churn
335.000	No	0.402
6814.000	Yes	0.513
3801.000	Yes	0.607
5653.000	Yes	0.538
784.000	Yes	0.608
6778.000	Yes	0.860
5862.000	Yes	0.928
2971.000	Yes	0.928
5268.000	Yes	0.695
6856.000	No	0.465
5912.000	Yes	0.580
1524.000	Yes	0.779
801.000	Yes	0.888
5756.000	Yes	0.692
1494.000	Yes	0.793
6026.000	Yes	0.846
2031.000	Yes	0.909
692.000	No	0.476
2195.000	Yes	0.620
1401.000	Yes	0.840

Deployment

- Alternatively, we could choose to select cases based on profit rather than simple likelihood

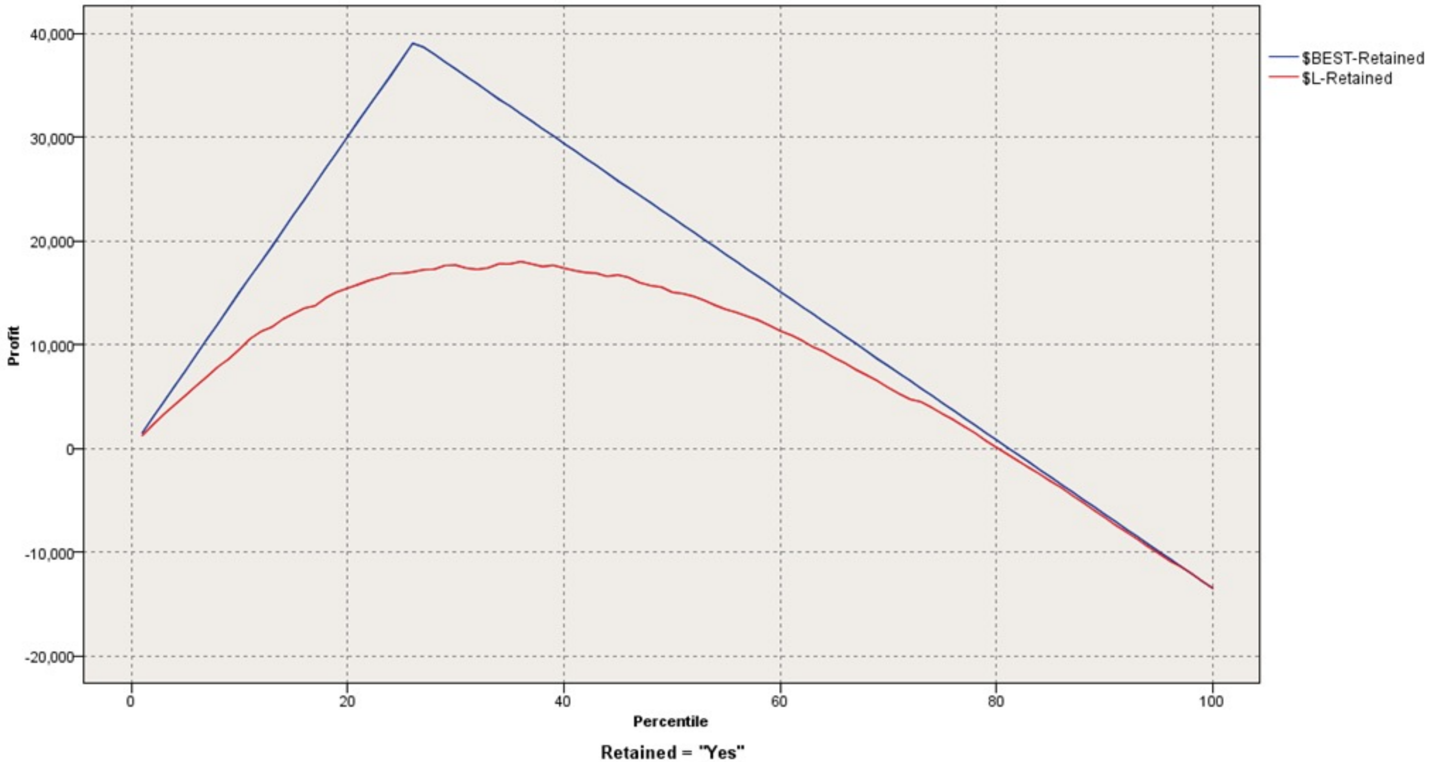

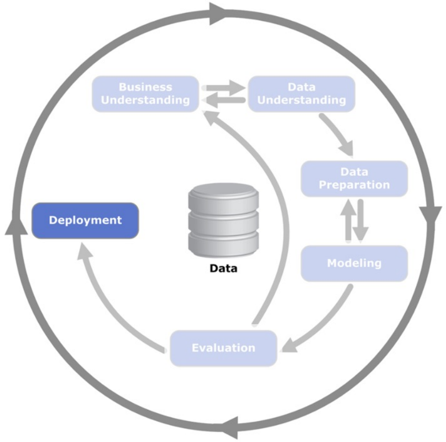

Deployment

- The apex of the model line shows the selection that should generate the maximum profit

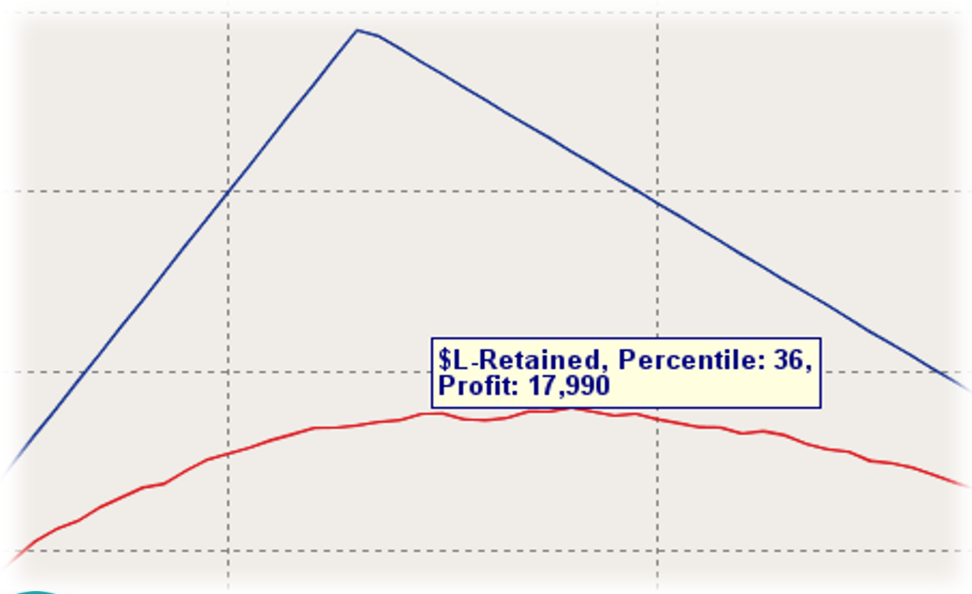

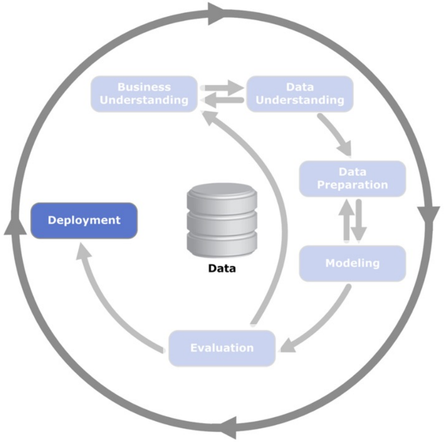

SMARTVISION
Europe

Deployment

- *Proving* that the application works may mean that we have to think of the entire exercise as an experiment with testable results

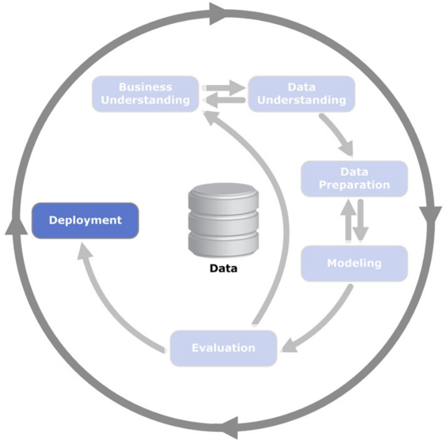

Does our retention strategy work?

Advice and Conclusions

How do we maximise successful deployment?

- Creating predictive models is the job of the data analyst
- Creating predictive *applications* is a team effort rooted in scrupulous planning
- Evaluating and deploying models is only ever effective with a thorough business understanding

Download our new e-book for free

The insider's guide to predictive analytics

£0.00

- 1 +

Add to basket

Category: books

A SELECT INTERNATIONAL COMPANY

Working with Smart Vision Europe Ltd.

- **Sourcing Software**
  - You can buy your analytical software from us often with discounts
  - Assist with selection, pilot, implementation & support of analytical tools
  - http://www.sv-europe.com/buy-spss-online/
- **Training and Consulting Services**
  - Guided consulting & training to develop in house skills
  - Delivery of classroom training courses / side by side training support
  - Identification & recruitment of analytical skills into your organisation
- **Advice and Support**
  - offer 'no strings attached' technical and business advice relating to analytical activities
  - Technical support services

Contact us:

+44 (0)207 786 3568
info@sv-europe.com
Twitter: @sveurope
Follow us on Linked In
Sign up for our Newsletter

Thank you