OPTIMISE MARKETING CAMPAIGNS

with Predictive Analytics

Andrew Addison, CEO, Purple Square CX

CONTENTS

1. Today's marketing challenge

All around the world, in every territory and every vertical, marketers are being asked to achieve more with less.

We, the marketers face ever increasing competition globally for products and services, and more niche companies are emerging that provide everything the consumer wants at a lower cost base. Consumers are expecting us to provide relevant offers to them, no matter how deep our relationship; we are supporting new products or services being released by our own organisations and internal competition for share of our customer and prospect wallet, all of this alongside tightening budgets.

More than ever, we need to ensure we present the right offers to the right people by the right channel at the right time. To do this well, we need to combine the twin pillars of **Predictive Analytics** and **Marketing Optimisation. [1]**

Whilst I am no marketing statistician or data scientist, I have been fortunate to work with some brilliant minds over the last 25 years, and some of it has rubbed off. I have seen teams and individuals that can estimate likely offer uptake and responsiveness to a given offer, to within a fraction of a percent and be right a lot more than they are wrong. However, at a granular level, it is rare to find one that will say "Person X will buy product Y" – they look instead at a wider holistic picture, a wider average, a group of people and predict the likelihood of that group, on average, to perform a specific action, or not.

Organisations invest significant amounts of money into marketing communications, so it is important to not only understand how effective that marketing automation is, but also to ensure that money, along with effort and resources are being spent in the right areas. This is where analytics and reporting are used, to identify potential targets and predict their likelihood to buy, to optimise marketing activities, report on effectiveness, and improve selection activities, all with the holy grail aim of increasing your organisations share of customer wallet.

Placing these two marketing analytical processes into a broad marketing campaign workflow:

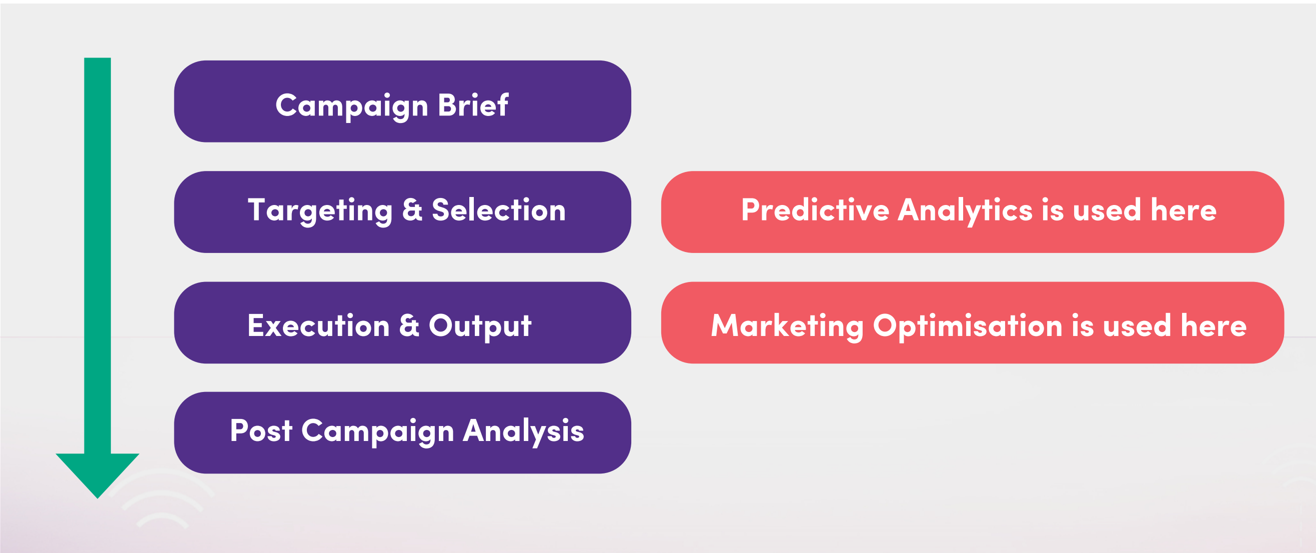

[1] We are focusing on the role of Marketing Optimisation in the context of customer communications and not the broader marketing mix

2. Identifying the right people

Overview

Predictive Analytics, within the scope of marketing, is a process by which an analytics or marketing team develop offers and targeted communications using data, propensity scores and metrics; to better refine and filter targeting and selection; to include only those with a higher likelihood to buy what is being promoted; whilst minimising costs and reducing waste.

That all sounds logical! But how does it all actually work?

Before we jump into the details and look at specific examples, as with most things, there is key terminology that you need to get your head around first. This book won't turn you into an analyst but at the very least, will be the bridge you need to have insightful conversations and meaningful planning sessions with an analyst; a partnership which is essential in any successful marketing automation strategy. Armed with an understanding of cluster modelling and analysis, propensity scores and clarity around next best offers, presenting the right offers to the right customers, at the right time, won't seem as daunting as it may do right now.

Throughout most of our careers as marketers we discuss the role of the customer in our business. We talk about them as if we know them individually! In reality, when working with hundreds of thousands to millions of customers, we cannot possibly know each individual customer, what we can understand are the groups of people that will behave or respond in a broadly similar fashion.

One approach is to create simple definitions, or personas, which make it easier to group our customers, such as High/Medium/Low value, or breakdowns by geographies, life stages or demographics. Whilst these do work and may often give you a broad mechanism for targeting, they often miss important nuances in the data.

This is where Cluster Analysis comes along. By using demographic, transactional, behavioural, attitudinal, and any other data points, from a variety of sources simultaneously, analysts can identify more meaningful and relevant groups (clusters) of customers for our business.

Once you understand your customer groupings, the next step is to identify who your product and services offers will appeal to most. This step is required to limit or prioritise the number of offers presented – there's no point in giving a 10% discount to a customer who is going to buy your product anyway. A score can be used as a filter criterion for a particular product category; meaning you not only stop spending money on people less likely (or overwhelmingly likely) to buy that product, but also allow those customers to be targeted for alternative, more relevant activities.

We call this a Propensity Score. This score enables businesses, based on strategic requirements and financial information, to filter the customers expected to be contacted.

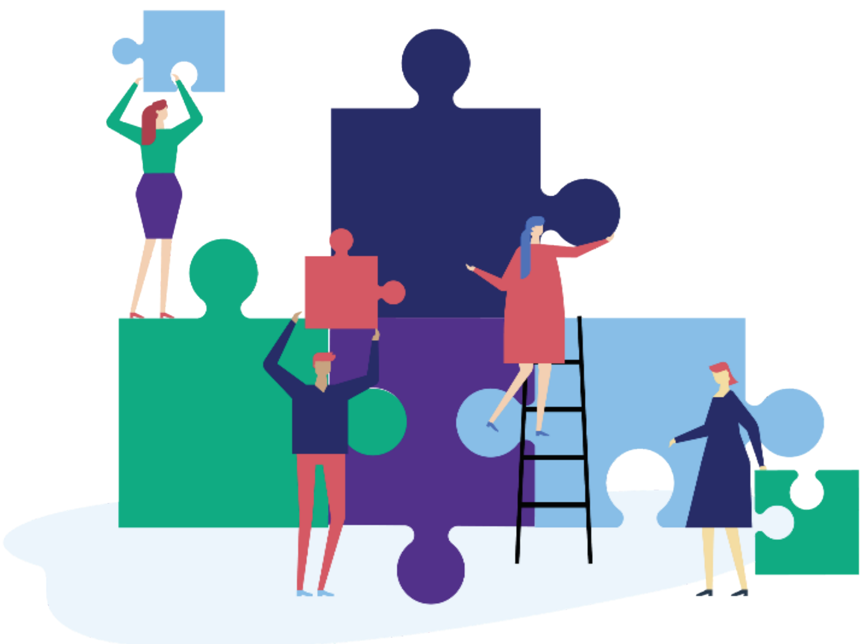

This score can be used to include or exclude groups or customers based on their likelihood to respond positively. However, it is a big decision as to where to draw the line. For example, assuming a score from 0 (very unlikely) to 1 (almost guaranteed). Removing 50% of our potential target audience for a promotion, (e.g. where the score is <0.6) could affect our campaign success by -1%? In isolation this could be a no-brainer, saving tens of thousands of spend on a minimal return, but if the margin value of the 1% sale is significantly higher than the cost of targeting everyone, then a different decision may be made.

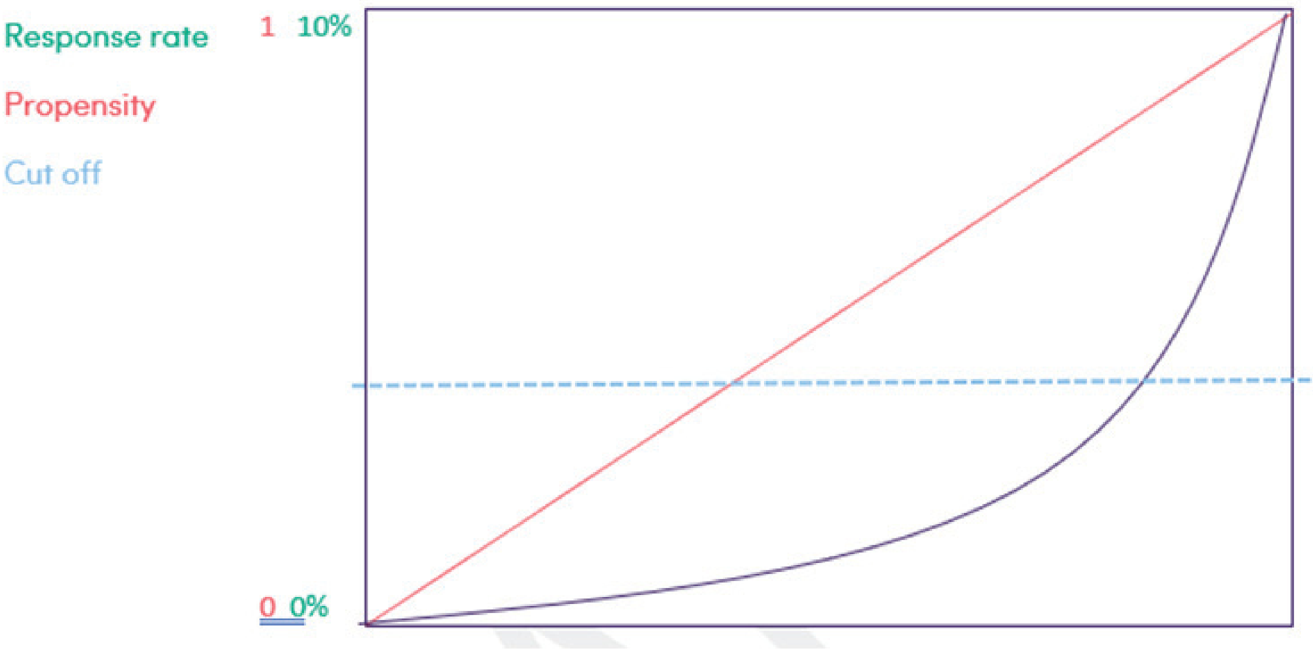

We've all heard the adage, "**it costs 5-7 times more to acquire a customer than retain one**", and whilst the specific numbers vary by organisation and industry, maintaining a customer by meeting their needs and expectations is always cheaper than acquisition.

Therefore, a significant aspect of predictive analytics efforts goes into Churn Analysis, to identify those customers that are most at risk of leaving or changing/reducing service level.

This is particularly important for organisations that provide ongoing or subscription services, where the lost customer may be lost for an extended period to a competitive service. As a marketer it is essential to protect and incentivise the "high propensity to lapse" customers to stay, with rich offers and new contracts.

Example

Company X has 1 million customers and is about to launch their latest product. Cost of the preliminary launch campaign communications is approximately £100k, plus £1 per contacted customer.

Sending this communication to everyone will cost £1.1mn, but the launch budgets are capped at £350k. Meaning only 250k customers can be contacted.

The company analysts have identified that the most likely 250k responders will generate 8% response rate, compared to a 3% for the whole customer base.

Average gross margin is estimated to be £40 per sale.

By contacting everyone, we estimate that there will be 30k responders, generating £1.2mn in gross margin, at a cost to market of £1.1mn.

By contacting only the top 250k customers, there will be 20k responders, generating £800k in margin at a cost to market of £350k.

We can see there has been a gross margin reduction of £300k, this is more than offset, by reduced cost of marketing.

Before we jump into the next level of detail, a reminder of why we use predictive analytics in marketing automation:

1. **To manage costs**
2. **To improve response rates**
3. **To prioritise the most valuable customers**

Cluster Modelling

Whilst we aim to create the feeling of a 1-to-1 relationship with our customers, with many hundreds of thousands of customers, if not millions, the ability to create truly unique customer experiences is almost impossible. The reality is, that when you are working at this scale, you need to sort your customers into more manageable groups, so that you can create experiences that resonate with the resources you have available.

Cluster modelling focuses on identifying groups of customers or prospects with very similar traits and behaviours. These groups could be your highest value customers, most engaged customers, customers with the lowest risk of churn, the lowest profitability customers, the least engaged customers, or those with the highest risk of churn. By creating these clusters of customers, you can start to direct specific activities and programs to supporting their needs and your goals.

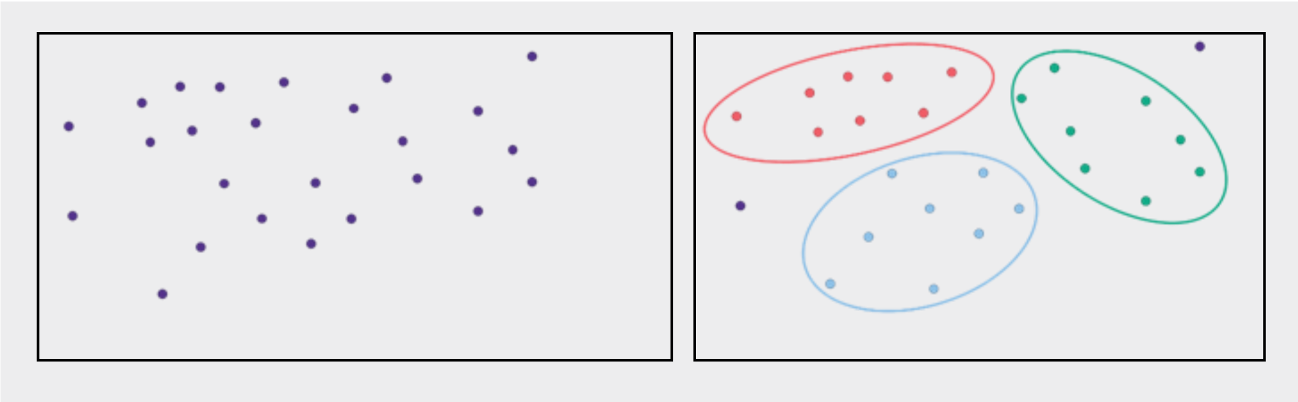

What was seemingly random data, can be interrogated, to identify consistent groups, with the occasional outliers.

Many marketers are already familiar with the use of clusters, we use them in our marketing and product/service planning, we just call them Personas. Personas provide an approximation of a customer type within our customer base, allowing us to understand who buys from us, what their motivations are and how they respond to the brand. These personas are then used to build relevant customer journeys to lead to a growth or improvement in our relationship with them.

Clusters are created in several different ways and often for different purposes, ensuring you understand the definition and construction of the cluster analysis is imperative to making good use of them. Clusters must be re-run on a regular basis to ensure they remain accurate and represent the customer base today.

We are not going to discuss the specific methods for performing a cluster analysis, there are plenty of tools and specialists that already do that hard work for you. But at a high level, a simple cluster analysis will use several different customer attributes, demographics, transactions, and behaviours to determine meaningful groups. The more attributes that are involved, the more fragmented the clusters are likely to be. The goal of the marketing analysts/statistician is to identify the meaningful or statistically significant attributes that make a difference. For example, do both customer revenue and margin correlate enough that one of them is effectively redundant?

Whilst there are a multitude of ways of interpreting your different clusters and segments, it is not unusual to prioritise them for attention in different ways. The example below is just one example, with A1's being the best, most engaged customers (who you'd have to actively annoy to lose) to C3's being the least engaged and least likely to purchase again.

There could be more than one cluster in each cell, or none at all, but the goal is always to move more of those cluster members up towards the A1 cell.

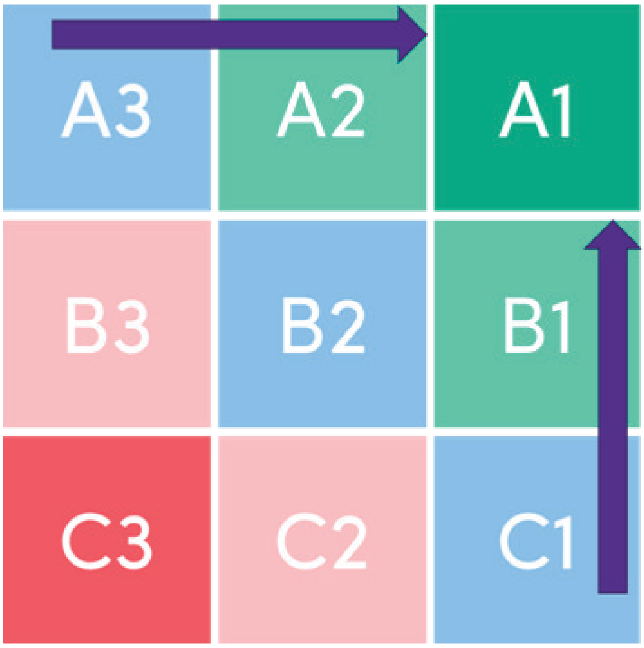

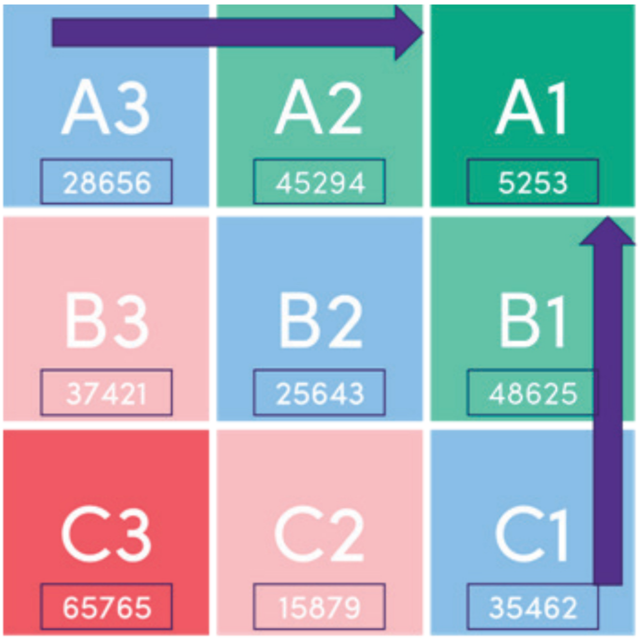

One of the big challenges comes when we start seeing numbers allocated.

Looking at the second image, would you really reduce your marketing to the 65,765 people in cell C3 to zero? Yes? No? Maybe?

Yes: Could you achieve more by allocating that spend to the 89,761 people sat in A3, B2 and C1, those that might benefit from the nudge up that additional investment creates.

No: Are C3 detractors? i.e. those that are negatively influencing our business and would that investment reduce the number of detractors.

Maybe: Would some investment on a different journey enable us to play a longer game to win them over, but most of our marketing spend is redirected to the short-term wins?

The adage of “it costs 5 times more to win new customers than retain an existing one”, applies even within the customer base. Loyal customers up in the top right are easier and cheaper to maintain than those who are less engaged or faithful to your brand (Bottom left).

Being able to quantify those meaningful clusters of people allows us to make meaningful decisions about the future.

One of the risks associated with Cluster Modelling is that marketers assume that they have an “average person” and treat the segment accordingly. Studies performed across a range of groups of people and requirements shows that an average person rarely exists. Australian researchers looked for the “Average Aussie” following the 2011 census, Female, aged 37, 2 kids (boy and girl aged 9 and 6), lived in a state capital with her husband in a 3-bed house, they found that person did not exist, anywhere (*).

We use cluster modelling to create our groups with a tolerance, no single person is exactly right in all areas and each of our customers will vary from the average in one or more ways.

(*) For a humorous look at this and many other assumptions and inaccuracies, take a look at Humble Pi (ISBN-13 : 978-0141989143), by Matt Parker, which is a very entertaining mix of mathematics and errors!

Propensity Scoring

A simple stat for you, that applies to every company, everywhere.

100% of your sales come from 100% of your customers.

If you look at your database, all your customers account for all your sales (either that or you have glaring database issues to investigate – that's for another book!)

However, most of us are also aware of the Pareto principle, that 80% of the benefit is achieved by 20% of the effort. Pareto identified that 80% of the peas in his garden came from 20% of the pea pods and later that 80% of Italian land was owned by 20% of the population. In technology, the rule of thumb is that 80% of the product complaints come from 20% of the product defects. In sales, 80% of revenue comes from 20% of the customer base. Whilst the ratio will vary by industry or organisation, such as 80:20, 70:30, etc, the principle applies, it is highly likely that a significant majority of your sales comes from a significant minority of your customer base.

I will caveat this by saying it varies massively by industry, for example, it is highly unlikely that manufacturers of commercial flight simulators will apply these activities in the same way that consumer telecommunications would.

What that means for us, as marketers, is that we don't have to talk to everyone all the time about every product, service or offering to maximise the benefit to the brand. The role of propensity modelling is to identify those that are most likely to respond positively to a communication as well as those that are less likely to respond positively (which is very different to responding negatively).

The most common approach to propensity modelling is to rate someone's likelihood to take a specific action on a decimal scale of 0-1, where 0 means no chance of performing an action and 1 means guaranteed to take the action. The "action" itself may vary, likelihood to buy, click or churn are all valid things to measure.

Data science teams spend a significant amount of their time building models and providing data that marketers can use to make better decisions about who to target, with what, and when. Marketing data should have numerous scores supporting the different business decisions you wish to make.

The biggest challenge we face is at what level to draw the line. Do we perform deep product level analysis for everything we sell? Or do we limit it to product categories? That is a business decision, based on the level of work involved, not to mention the vertical you are operating in.

For example, Stock Keeping Unit (SKU) level analysis for a grocery retailer would be very challenging to manage. Is it worth looking at a can of "Leading Brand" - Cola vs any cola vs carbonated soft drinks? There are always going to be die hard loyalist, who would buy nothing but that leading brand and those that would never buy it.

Creating models that are easy to deploy and relevant to your goals, are the key to helping future decisioning.

Predictions not facts

Past performance is not necessarily a predictor of future performance.

This statement can be seen in the collateral of financial institutions all the time, and it is a good mantra, showing that nothing is guaranteed, but it remains the best guide we have to the future. As marketers, we are constantly looking for good likelihoods, and as such great propensity models to learn and evolve, using new data to tweak understanding of how customers interact with us.

Once data has been generated by propensity models, it's time to look at their application in marketing decisions.

An important point to remember is that a propensity model is not a rank of one customer against another, but an indicator of how likely one customer is to perform a specific action. Therefore, this data can be used to make decisions in two primary ways.

1. Decisions based on score threshold
2. Decisions based on a target number of customers to contact

To explain the application of propensity score data, let's look at a sample data set of 1000k customers who have been modelled to purchase a widget with an average score of 0.2458 and a range of 0.0000 to 0.9303. Below is a selection of the data in this propensity score.

Top of the list		Bottom of the list	
Customer	Score	43886435	90.1112
88570	0.3395	43930720	0.3041
132855	0.1117	43975005	0.0429
177140	0.1942	44019290	0.9156
221425	0.0166	44063575	0.1718
265710	0.6514	44107860	0.1229
309995	0.3801	44152145	0.0002
354280	0.0137	44196430	0.4485
398565	0.0004	44240715	0.0672
442850	0.5613	44285000	0.0897
487135	0.115	44329285	0.0822

By using a score threshold, we could decide that a score of >=0.7 on our propensity to buy would be the minimum to accept, this would return 44k customers to contact.

Alternatively, we may have a campaign budget for 100k customer contacts, which would mean our propensity score drops to 0.5715.

One of the biggest challenges at this point is making the right decision for the business. Is it better to target fewer people for this campaign, reducing costs and increasing response rates or is it better to meet the needs of product category owners who want to maximise the total number of sales for their products & services. We will discuss options and concepts in this area in the next section (Marketing Optimisation).

Churn Analysis

Every organisation expects to lose some customers over the course of any given month. They can move to another brand, stop needing or using our products, or becoming disengaged with us. Our goal through Churn Analysis is to identify those that demonstrate the likelihood to leave this month in time to do something about it.

We use different business triggers and events to determine whether someone is at risk of leaving. Has their usage of our product/service reduced recently? Have they called to complain about something? Has a competitor launched in your industry or region? Have they opted out of a channel or otherwise reacted negatively to the brand?

Each of these is an indicator of whether someone could potentially move on. The question we ask ourselves at this point, is “are they valuable enough for us to try to retain them?”

Not every customer is as valuable as the other, so identifying those customers that are acceptable losses enables us to invest more into those customer relationships with a likelihood to grow and ultimately become advocates for our brands.

Next Best Action

Next Best Action (NBA) Modelling is a predictive method that drives a highly customer-centric approach. It encourages and drives the marketer to understand what the customer needs are, and by identifying at an almost individual level the most appropriate offers and communications to meet those needs. The key driver is to understand what the customer wants, and not what we want to send them. This highly customer focused approach to marketing can be a real challenge for product centred organisations to adapt to.

It's important to note that NBA is a multi-channel strategy, sitting across both live or real-time touch points as well as our offline and batch communication platforms. It uses deep understanding of the customer, their likely behaviours and potential lifetime values to grow the level of engagement.

Over the last few years, there has been an increase in NBA initiatives across many organisations, often through upsell and cross-sell opportunities. "People who bought X, also bought Y", pop-ups as we check out, are a simple, but effective method to generate additional business. NBA takes this much, much further.

NBA veers away from the traditional company focused "Customer Journey" approach where marketers define a specific set of events and communications, which are sent in a specific timescale, to create additional customer engagement, and instead switches to a responsive and reactive approach to communications. A downside of this approach is that marketers do not always feel like they have full control of what is being sent to whom and when and this lack of advanced visibility creates nervousness, especially amongst product marketers with conversion targets.

As the customer is at the centre of an NBA strategy, it's important to respond to their needs as they become apparent to us, not at some later date at our leisure.

This means making decisions in **Real-Time**, not in advance (when we don't know enough) or after the fact (when it's too late). As a result, NBA strategies do tend to require machine learning (or occasionally AI) platforms in order to get the best results, using past activities, context relevant information and potential future options, can require a high amount of data processing and often in a very short window of opportunity. For example, you may be presenting an offer just as the customer is interacting with your website.

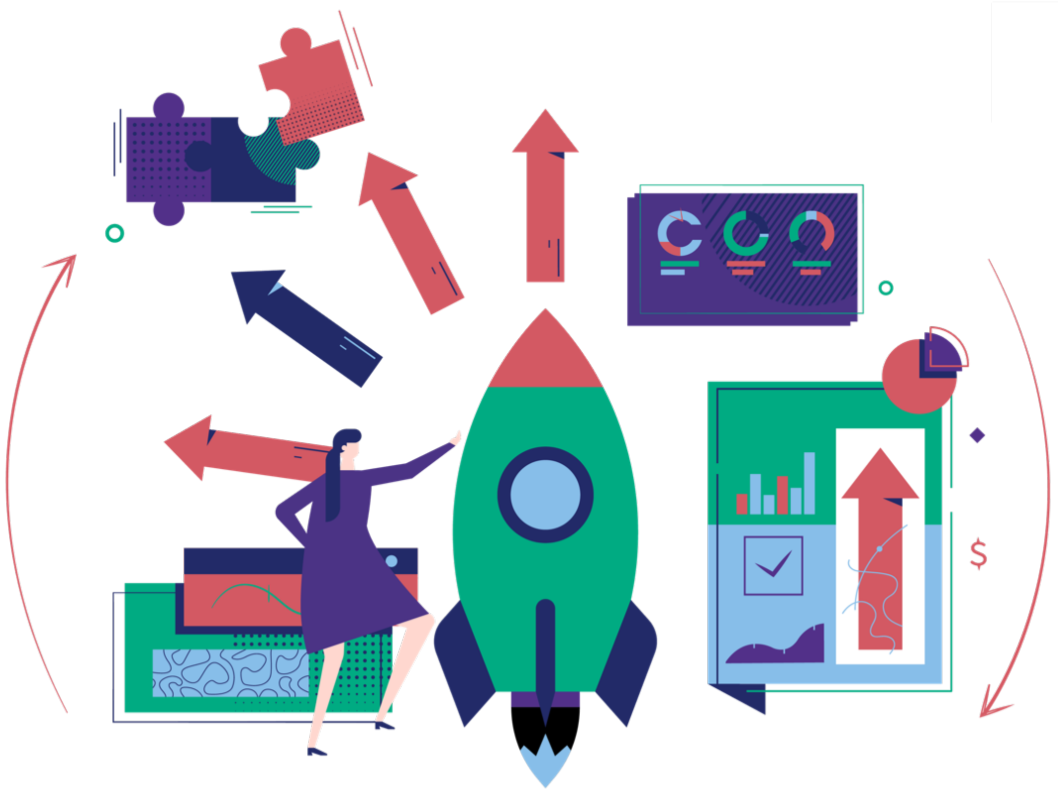

Okay, so now you know the key terms and concepts. Next, we'll jump into the bit that matters most to marketers – how to use all this analysis to optimise marketing campaigns!

Key Points: Predictive Marketing Analytics

◊ Cluster Modelling

- Identify groups of customers displaying similar behaviours

◊ Propensity Scoring

- Identify likelihood to respond/purchase
- Business defined score thresholds

◊ Churn Analysis

- Cheaper to maintain and grow customers than acquire
- Acceptable losses are OK

◊ Next Best Action

- Customer focused, not product led
- Multi-channel focus, inbound and outbound

3. Optimising our communications strategy

Once we have our clusters, predictive analytics scores and models ~~in place, we can now start to focus on the~~ application of these within customer communications, in short, who receives which offer and when.

Marketers are always seeking innovative ways to acquire, retain and grow customer value and engagement. The situation is exacerbated by a growing number of customer touch points, more granular audience segments and more product and offer permutations. The temptation is to send more messages to more customers in the hope of achieving better marketing results. This is a particular challenge for companies that have many competing communications (often originating from different departments) that they want to send to the same or overlapping sets of customers.

With each of these factors affecting marketing planning, marketers often find themselves struggling to reconcile the effect of each of these levers, requirements, or constraints across campaign planning. In the end priorities often depend on who submits their request first, who shouts loudest, who has the biggest budget, or what is subjectively perceived to be the next best offer.

This situation raises several business challenges:

With fixed channel capacity, what channels do I use to communicate with which customers and when? Our call centres can only handle a certain number of calls per hour, direct mail may stretch our budgets too far or we know we have limited opt-ins to our in-app messaging. Which channels and to whom can greatly impact our targeting decisions.

With limited opportunities to communicate, which offers do I give to which customers while managing contact fatigue? How can I choose the best campaign for a customer from among several that are overlapping?

Which customers get offer X, when there are only a limited number to distribute? How do I achieve the maximum expected return on my campaigns given finite resources?

Which offers conflict with each other? Which offers do I give only if another offer has already been presented? Can I avoid customer base cannibalisation by multiple campaigns run by different lines of business?

How do I allocate my money across different campaigns, offers, or customers? Can I design and proactively manage my customer contact strategy, rather than just letting it happen?

How do I ensure customers who elect to not receive email are not contacted by email? What is the best way to enforce cross-campaign policies (opt-outs, DNC, self-exclusion, etc.)?

There's maths in marketing too!

Marketing Optimisation is a mathematical approach for identifying the best combination of messages for each customer from among competing options, while complying with business objectives and constraints. All with the goal of targeting the right customer, with the most attractive offer, at the optimal moment, using the most appropriate channel.

This, as we have all experienced, is a complex problem to solve – but based on numerous success stories through the industry, if you can solve it, you will reap the rewards. The number of customers by channels, by offers, by time options, creates potentially billions of variable combinations. To add to the complexity, marketing organisations are also constrained by business goals and restrictions, such as revenue targets, contact fatigue policies and budget caps.

Under these pressures, organisations will use a variety of methods to get the best results, all of which are perfectly valid, but will often result in significantly different rates of success.

Whether you're dependent on external technology or have in-house capabilities around Predictive Analytics, how you use these insights will result in varying degrees of success.

Using a worked example, we'll walk through the most common approaches to marketing campaign optimisation. As you read through, you're likely to identify the approach your organisation currently uses.

A worked example

To demonstrate the pros/cons of some of these methods, I'm going to use a fictional organisation, the Purple Supplies Company (PSC). PSC doesn't have many customers (10) or many campaigns (3).

Using methods like those described in the earlier section, PSC data scientists have scored each customer on their likelihood (out of a maximum of 100) to take up the offer presented in each campaign as part of their daily data preparation activities. They have chosen a simple propensity score for consistency, but could have used any metric of success, such as potential revenue.

It is important to be consistent about the "value" used in the modelling to ensure a consistent measure. Try to avoid mixing and matching revenue/response/margin measures.

We will assess each campaign's potential effectiveness against each of the 4 common "prioritisation/ optimisation" methods and compare them using a per campaign (highlighted blue) and overall score (highlighted green).

We also have a very simple set of business constraints. These are applied to ensure we don't over market to an individual or leave any campaigns without enough recipients.

Business Constraints

1. One customer can receive one campaign only (this ensures contact fatigue rules are followed)
2. One campaign has a minimum of 3 and maximum of 4 recipients (this ensures a broad spread of relevant content to meet internal business needs)

The table below shows the score assigned for each customer by each offer/campaign.

		Campaign		
		A	B	C
Customer	1	56	44	44
	2	40	56	60
	3	48	56	52
	4	44	64	60
	5	60	48	40
	6	60	52	48
	7	64	56	60
	8	52	48	48
	9	64	88	60
	10	80	96	72

			Run Order

Campaign A Score	
Campaign B Score	
Campaign C Score	

Approach 1: First run, first out the door

This approach is effectively a random based selection on the first campaign to run that day. It may use elements of customer scoring but is most likely to choose the first data records it hits that meet the criteria in the database output.

		Campaign		
		A	B	C
Customer	1	56	44	44
	2	40	56	60
	3	48	56	52
	4	44	64	60
	5	60	48	40
	6	60	52	48
	7	64	56	60
	8	52	48	48
	9	64	88	60
	10	80	96	72
Run Order		2	1	3

Campaign A Score	220
Campaign B Score	184
Campaign C Score	180
	584

We've rated this group of campaigns with a combined effectiveness score of 584.

The resulting Campaign scores are arbitrary, and the selection is simply based on the first 3 or 4 customers available in each campaign and depending on the order the campaigns are run in the day, which will create significant variations. In this instance, Campaign C is left with the only customers not already contacted by Campaigns A and B.

Thankfully, this approach is increasingly rare, but we do still see it on occasion.

Approach 2: Campaign First

In this case the first step is to prioritise campaigns and then select the best customer(s) for each campaign in order. The target recipient for that campaign is the highest scoring customer within the campaign. Whilst this approach is likely to produce good results for the higher priority campaigns, it does result in poorer results for the lower priority communications as they will typically be left with the lower scored customers to choose from. i.e. the higher the campaign ranking the larger pool of customers they get to cherry pick.

		Campaign		
		A	B	C
Customer	1	56	44	44
	2	40	56	60
	3	48	56	52
	4	44	64	60
	5	60	48	40
	6	60	52	48
	7	64	56	60
	8	52	48	48
	9	64	88	60
	10	80	96	72

Campaign Priority

Campaign A Score	268
Campaign B Score	176
Campaign C Score	132
	576

We've rated this with a campaign combined effectiveness score of 576

In this instance campaigns are taken in the order A > B > C. In the example above, Campaign C can only be sent to customers 1, 5 and 8. However, as the better customers are often better across a range of product/offer areas, it is likely that this campaign will be left with a very marginal opportunity for success.

Approach 3: Customer First

This approach switches the emphasis to put the customer first. In this approach, each customer is evaluated row by row, evaluating each campaign to select the highest scored campaign for that individual. This can mean however that the higher up in the data the customer is, the more likely that they will get a campaign offer they are likely to respond to as they have wider pool of campaigns to select from.

		Campaign		
		A	B	C
Customer	1	56	44	44
	2	40	56	60
	3	48	56	52
	4	44	64	60
	5	60	48	40
	6	60	52	48
	7	64	56	60
	8	52	48	48
	9	64	88	60
	10	80	96	72

Customer Data Order

Campaign A Score	240
Campaign B Score	168
Campaign C Score	192
	600

We've rated this with a campaign combined effectiveness score of 600

As you can see in the data set above, customers 9 and 10 must receive Campaign C as they would otherwise breach the minimum number of recipients per campaign (min 3, max 4). The further down in the customer list an individual appears, the fewer campaign options they are left with.

Approach 4: Marketing Contact Optimised

This approach evaluates all of the scores, rules and constraints and reduces the bias of customer or campaign. Optimisation algorithms apply all the logic, across all the data, at the same time and are designed to identify the best combination of all the factors to provide the maximum output score. It can be difficult to determine which factors have greatest influence, and the introduction of one rule can significantly impact the output results, but always with the same goal.

		Campaign		
		A	B	C
Customer	1	56	44	44
	2	40	56	60
	3	48	56	52
	4	44	64	60
	5	60	48	40
	6	60	52	48
	7	64	56	60
	8	52	48	48
	9	64	88	60
	10	80	96	72

Campaign A Score	240
Campaign B Score	240
Campaign C Score	168
	648

Optimised Decision

We've rated this with a campaign combined effectiveness score of 648

Even using this simple model, this approach results in an 8% uplift on the best performing alternative method earlier (Approach 3: Customer First) and 12.5% uplift on the worst performing method (Approach 2: Campaign First). Converting this to the financial measures of your organisation it could result in many tens of thousands of pounds/dollars in incremental revenue.

Comparing the Results

Whilst in our example, the scores have been made up, what it demonstrates is that the way in which campaigns and customer communications are prioritised can have a significant impact on results, with a campaign combined effectiveness score range of 576-648.

	Campaign Score			
Approach	**A**	**B**	**C**	**Overall**
First run, first out the door	220	184	180	584
Campaign First	268	176	132	576
Customer First	240	168	192	600
Marketing Contact Optimised	240	240	168	648

The next question to ask yourself is, what would it mean if we scaled it up and each score point represented £10,000 in won or lost revenue? This could mean a £720,000 swing in total campaign attributable revenues for this campaign cycle.

IMPORTANT CONSIDERATION

When applying marketing optimisation techniques, you may suppress more records from a campaign than an individual campaign budget threshold will allow, for example you may have a "minimum constraint" for a campaign as well as a maximum. It is always worth selecting more records than needed within the targeting & selection phase to allow for suppression during the execution phase of your campaign, before performing a final "trim" at the end.

Marketing Optimisation Challenges

Whilst it is easy to see and evaluate the impact on three campaigns and 10 customers, it's much harder to do this with millions of customers, tens or even hundreds of campaigns and numerous channel constraints, all creating a significant number of permutations within a day or week. The situation is exacerbated by a growing number of customers touch points, more granular audience segments and more product and offer permutations. The temptation is to send more messages to more customers in the hope of achieving better marketing results, but this often has the opposite effect. This is a challenge for companies that have many competing communications they want to send to the same or overlapping sets of customers.

There are also further organisational challenges to implementing a more optimised approach;

Many marketing organisations we have worked with at Purple Square are swayed by the demands of product managers; their requirement is to push the highest scoring customers to the top of their campaigns (and rightly so, for them). However, this is likely to have a negative impact on other campaigns and the business overall. Changing the way product managers are recognised and rewarded is key to ensuring that the organisation benefits. It is important to set targets, but not at the cost of other business units.

Other organisations struggle to score their data at a meaningful level, whilst the most optimal approach is to score every offer, against every customer, this can create a significant burden on the data scientists. Starting at a high product/category level and working down through the product hierarchies will result in improving scores and optimisation over time.

It is tempting to try to create the perfect optimisation model for the organisation on day one, however this activity can take many months, if not years, to evolve and will never be completely "perfect". Most successful optimisation projects start with a limited scope to prove the concept and build confidence in the solution, followed by incremental deployment, adding new rules, campaigns and scoring as the business learns and develops.

Investing time and resources, addressing these challenges and embedding marketing optimisation into the marketer's business process will always lead to overall campaign effectiveness.

Key Points: Marketing Optimisation

- ◊ What do you want to optimise?
  - Revenue
  - Margin
  - Response rate
- ◊ Considerations
  - Channel Capacity
  - Contact Fatigue (over contacting)
  - Offer Constraints
  - Offer Conflicts
  - Budget Constraints
  - Channel and Customer Suppressions

4. Pulling it together

The end goal is to use the scored information generated in Predictive Analytics to maximise the overall effectiveness of all marketing communications, not just for an isolated campaign or communication. Together Predictive Analytics and Marketing Optimisation are the powerful duo that present the right offer to the right customer, at the right time.

Revisiting these two marketing analytical processes within a broad marketing campaign workflow:

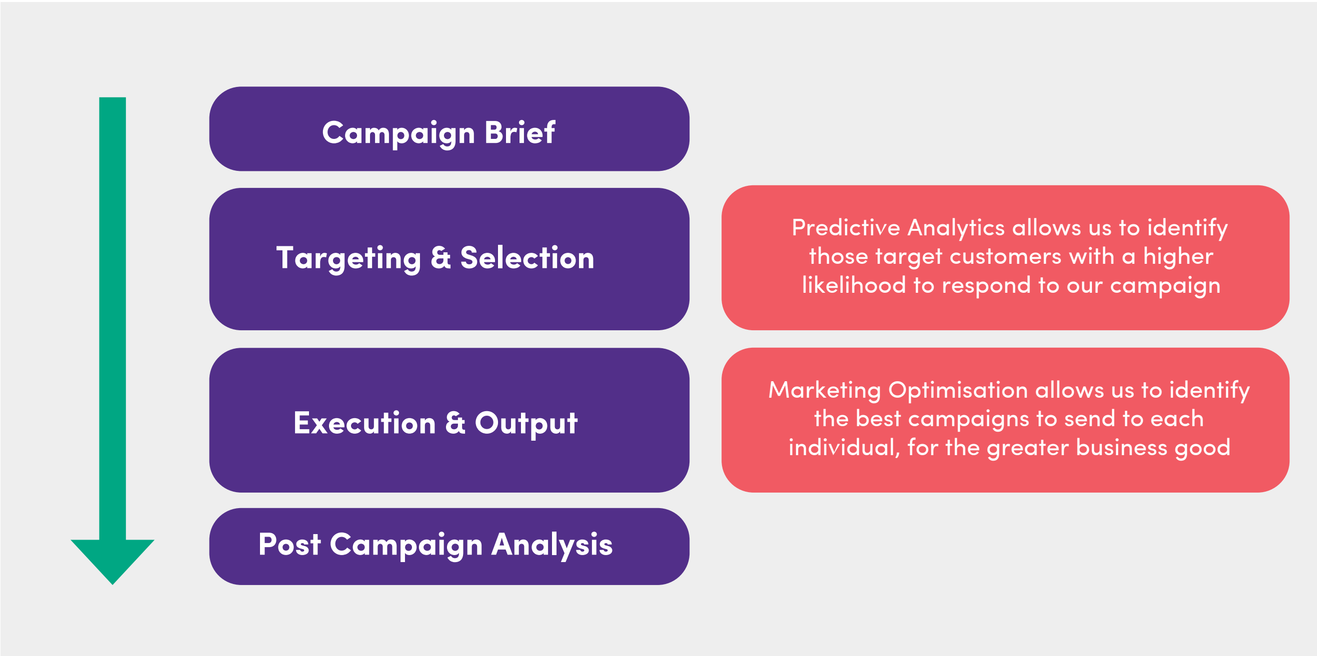

Predictive Analytics provides us with the scores and measures upon which to base decisions. Understanding who our customer segments (cluster analysis) are, how they are likely to respond to offers and communications (propensity modelling), whether they are at risk of leaving (churn analysis) and what is the right offer for them at this time (next best action) enables us to provide both a compelling and more importantly contextually relevant communication strategy for our customers.

When marketers combine Predictive Analytics data with our marketing optimisation programmes, they're able to widen the potential pool of contact/offer combinations and use the tools at their disposal to create bespoke and almost unique approaches to execution. All of this effort ultimately maximising the potential return on marketing campaigns for the greater [financial] good of an organisation.

About the Author

Andrew has been delivering CX and Marketing Automation solutions for over 25 years, working for software vendors, specialist consultancies and large enterprises throughout that time.

In 2011 he formed Purple Square to deliver independent advice, services and support to enable clients, technology vendors, systems integrators and business partners to make the best use of their marketing automation technology and deliver an ever improving customer experience.

When not working he spends most of his time with his family, playing golf (badly), watching rugby and visiting Italy whenever he can!

Linkedin: @andrewaddison

About Purple Square CX

Purple Squqre CX, a Customer Experience Advisory that offers a diverse range of services aimed at enhancing customer interactions for businesses. Our expertise lies in three key areas - CX advice and strategy, marketing automation, and customer data platforms (CDPs).

Our team comprises CX visionaries and strategists, engineers, architects, developers and builders, all focused on delivering against the five core principles of CX: Vision, Operations, People, Data and Technology.

We build long term partnerships with our clients, that deliver their Customer Experience goals, both short-term and long into the future.

PURPLE SQUARE